Inequality, Heterogeneity, and Consumption in the *Journal of Political Economy*

Greg Kaplan

August 2017

Today, inequality and heterogeneity are front-and-center in macroeconomics. Most macroeconomists agree that the distribution of household-level variables, in particular consumption and wealth, matter for the dynamics of macroeconomic aggregates and that macroeconomic shocks affect the distribution of consumption and wealth across households. However, it was not always this way. Getting to this point has been a long road, along which papers published in the JPE have been essential guideposts.

I will discuss six influential papers from the JPE that have helped shape the way that inequality is studied by economists today. I have organized my discussion along three strands that have each contributed to the introduction of heterogeneity into macroeconomic models: the microeconomics of consumption behavior; the use of structural models of precautionary savings for insights about policy; and finally, the development of general equilibrium models with heterogeneous households and aggregate shocks. Within each of the three strands, I have chosen two defining JPE papers that roughly correspond to the trajectory that this intellectual journey has followed.

First I will examine two empirical papers — Zeldes (1989) and Attanasio and Weber (1995) — which provided new insights into the microeconomics of consumption behavior. Both are empirical analyses that are closely guided by theory. These and other related empirical papers paved the way for a class of structural models of consumption — precautionary savings models — that later become the workhorse models of heterogeneous agent macroeconomics. I will next highlight two early examples of how structural models of consumption can be used to explore the ways that economic policies shape the distribution of household outcomes. Hubbard, Skinner, and Zeldes (1995) and Rosenzweig and Wolpin (1993) illustrate two alternate ways in which structural models of consumption can be disciplined by data and turned into quantitative laboratories; the first calibrated their model of the United States using mostly external sources, while the second estimated their model of rural India using maximum likelihood. In both models, precautionary motives drive household responses to changing government policies. Finally, I will discuss two early heterogeneous agent models with aggregate shocks: Imrohoroglu (1989), one of the earliest partial equilibrium models, and Krusell and Smith (1998), a now-seminal general equilibrium model. These took previously standalone precautionary savings models and cast them in an equilibrium framework exposed to aggregate shocks, thus opening the door for fully-fledged macroeconomic models of inequality.

Zeldes (1989) was one of the first papers to provide convincing evidence using micro panel data that liquidity constraints are indeed important for household-level consumption. The importance of his contribution is reflected in the fat that today, virtually every paper being written about consumption, either at the individual or aggregate level, somewhere addresses the implications of binding liquidity constraints. The paper is notable because rather than simply rejecting the implications of consumption models that abstract from liquidity constraints, as had been done by much of the pre-existing literature, he carefully derives testable implications of a model that includes liquidity constraints. Zeldes (1989) suggests three novel tests for the importance of these constraints. The first test is beautiful for its simplicity. Because the model with liquidity constraints predicts that consumption growth should be sensitive to income for low wealth households but not for high wealth households, Zeldes (1989) splits his sample between these two groups of households and measures the sensitivity of consumption growth to income growth for each. Using data on food consumption from the Panel Study of Income Dynamics (PSID) he confirms the predictions of the model – that low wealth households have a much lower sensitivity of consumption growth to income than high wealth households. Refined versions of this test for the presence of liquidity constraints are still the go-to approaches in empirical analyses of consumption behavior today. Many of these recent studies (and there are many) essentially repeat the analysis of Zeldes (1989) using better-identified income shocks and much larger and higher quality datasets, reaching the same conclusion.

Zeldes (1989) also derives two additional tests implied by the consumption model with liquidity constraints: (i) that Lagrange multipliers on liquidity constraints should be positive for constrained households; and (ii) that the multipliers should be negatively related to current income. He estimates Lagrange multipliers for low wealth households by using the residuals from their Euler equations when the remaining parameters are estimated in the sample of high wealth households.

He finds that the estimated multipliers are indeed positive and are negatively related to income, as predicted. One leaves the paper with the sense that any model of household consumption should treat liquidity constraints seriously. The structural literature that followed did.

Another beautifully executed empirical analysis of consumption behavior that is also carefully motivated by theory — this time without liquidity constraints — is Attanasio and Weber (1995). Using the Consumer Expenditure Survey (CEX), they pioneered the approach of working with synthetic cohorts (i.e. data moments for groups of households from the same birth cohort with similar demographic characteristics) to overcome the limitation that the CEX has only a short panel component, unlike the PSID. This enabled them to exploit the comprehensive consumption data in the CEX, rather than the PSID, which at the time contained data on only food consumption.

Attanasio and Weber (1995) make three points. First, they show that looking only at food consumption can be misleading because preferences are non-separable between food consumption and other consumption categories. Partly for this reason, the vast majority of the structural consumption literature that has followed favors comprehensive measures of consumption or explicitly models this non-separability. Second, they illustrate the pitfalls of using aggregate data to test models of heterogeneous households. By aggregating micro data in exactly the way prescribed by theory they show that there can be large differences between the dynamics of the log of mean consumption (the focus of representative agent models and what is measured in aggregate data) and the dynamics of the mean of log consumption (the focus of heterogeneous agent models and which can only be constructed with household-level data). Third, they show that it is easy to spuriously reject frictionless life-cycle models of consumption if one ignores predictable changes in either household composition or labor supply of individual household members. They explain how the hump-shaped age profiles for family size and female labor supply would lead to a hump-shaped age profile for consumption – an observation that had frequently been cited as evidence against the frictionless model.

All three points lead one to re-think the numerous previous studies that had seemingly shown deviations from frictionless consumption models, including Zeldes (1989). Attanasio and Weber (1995) never actually claimed that liquidity constraints and precautionary motives were not important — only that the then-existing tests were much more fragile than one might have thought. Reading both of these papers today, one is struck by the careful connection between empirics and theory. Both papers carefully explain their null and alternative hypotheses, build up their estimating equations from precisely specified models, and go to great lengths to spell out the assumptions required to go from their model to their regression equations. These are classic qualities of empirical analyses in the JPE.

Attanasio and Weber (1995) essentially highlights the limitations of examining theories of consumption through the lens of only a small subset of a model's predictions. By exposing these limitations, they drove later papers to use a larger set of predictions that are obtained by explicitly computing consumption and savings decisions under alternative parameterizations. These structural papers that followed moved beyond testing models to using models to quantify the effects of public policies on household consumption and savings behavior. Two JPE-published papers represent some of the best early examples of how to utilize a quantitative structural model of consumption for effective policy analysis.

Hubbard, Skinner, and Zeldes (1995), is a classic example of the power of a calibrated structural model. They observe that many households with low lifetime incomes accumulate little or no wealth over their lifetimes. Even just before retirement, when life-cycle models of precautionary savings predict that households should hold substantial wealth, many such households are essentially hand-to-mouth. According to life-cycle consumption theory, having low lifetime income, even in the presence of liquidity constraints, is no excuse for not saving for retirement — households should smooth consumption, albeit at a low level. While this might explain why households do not borrow, it does not explain why they do not save. Hubbard, Skinner, and Zeldes (1995) suggest a possible reason for the lack of saving: asset-based means-tested public insurance programs, which they model as a consumption floor, reduce the incentives for households to save. The presence of a consumption floor not only reduces households' exposure to consumption fluctuations — lowering their incentive to save for precautionary reasons — but also implies an effective tax rate of 100 percent on assets in the states of the world where the consumption floor binds.

The authors use a simple two-period model as an elegant theoretical proof-of-concept. But the calibrated life-cycle model, which is the meat of the paper, provides two additional benefits. First, it acts as a quantitative proof of concept, which, in my opinion, is one of the most valuable benefits of quantitative structural analysis. It is one thing to show that asset-based means-tested public insurance programs can distort savings decisions; it is another to show that in empirically plausible settings these distortions are large enough to have an economically important effect on observed savings. To do so, the authors choose parameter values that they argue reflect U.S. data, of which the most important are the stochastic processes for earnings risk and medical expense risk, the level of the consumption floor and the degree of risk aversion. They then simulate their model economy and show that it generates the aforementioned patterns of lifecycle wealth accumulation by lifetime income. Second, the calibrated model can be used to evaluate the implications of alternative versions of means-testing public programs for wealth accumulation, which the authors show can be substantial.

Another example of how a quantitative structural model of precautionary savings can be used to evaluate public policies is Rosenzweig and Wolpin (1993). They consider the consumption-savings problem of farmers in India whose only mechanism for smoothing consumption is through the accumulation of bullocks. Bullocks are also an input used in agricultural production, making this paper one of the first examples of a structural model in which households save in a productive asset in the face of idiosyncratic risk. The authors confront their model with panel data from the International Crops Research Institute for the Semi-arid Tropics (ICRISAT). They construct the likelihood over sequences of farmers' assets and profits and use a two-stage maximum likelihood procedure to estimate preference parameters, prices of bullocks and other inputs and production parameters. Even today, Rosenzweig and Wolpin (1993) remains one of the few examples of maximum likelihood estimation of a precautionary savings model of consumption using micro data. Their parameter estimates imply under-investment in bullocks on the part of farmers, as a result of borrowing constraints and the inability of farmers to accumulate precautionary savings in a financial asset. Through a series of counterfactual experiments, the authors evaluate the relative merits of alternative interventions. They find that the provision of actuarially fair-weather insurance would have little effect on farmer welfare, whereas access to assured income streams would have a large effect on welfare. These are quantitative conclusions that can only be obtained with a suitably parameterized model.

There are important senses in which neither the models of Hubbard, Skinner, or Zeldes (1995) nor Rosenzweig and Wolpin (1993) are 'macroeconomic.' First, in neither model is the return on savings determined as an equilibrium outcome. Second, in neither paper do the authors explore how aggregate disturbances affect the economy. I will finish by discussing two papers in the JPE that contributed to the transition towards developing realistic models of consumption that are macroeconomic in this sense.

Imrohoroglu (1989) was a significant early paper that recognized the potential importance of precautionary motives in the presence of aggregate shocks. Her paper was motivated by Lucas's famous costs of business cycles calculation (Lucas 1987). He had shown that in representative agent economies the welfare costs of business cycles are small both because fluctuations in aggregate income are themselves small and because these fluctuations have only a second-order effect on welfare. It was natural to conjecture that in heterogeneous agent economies with incomplete markets this quantitative conclusion might be overturned — both because fluctuations in individual income can be substantial and because the presence of liquidity constraints means that for some households these fluctuations have a first-order effect on welfare.

Imrohoroglu (1989) set out to evaluate this conjecture. She examines a consumption-savings model with liquidity constraints in which households face unemployment risk that varies stochastically with macroeconomic conditions. It is interesting to note how our understanding (and expectations) of what it means for a macroeconomic model to be labeled as 'general equilibrium' have evolved. Despite describing her environment as general equilibrium, most macroeconomists today would describe the model in Imrohoroglu (1989) as a partial equilibrium environment because all prices — interest rates, wages, job destruction rates and job finding rates — are exogenous. She finds that when aggregate shocks change the extent of unemployment risk faced by households, the welfare cost of business cycles can be four to five times larger than in a corresponding representative agent economy.

Perhaps the most influential macroeconomic model with heterogeneous agents and incomplete markets is Krusell and Smith (1998). They study an infinite-horizon consumption-savings problem in which ex-ante identical households are subject to idiosyncratic unemployment risk. As in the other models I have discussed, households can self-insure this risk through a single risk-free asset. Krusell and Smith (1998)'s innovation was to embed this precautionary savings problem in a stochastic version of the Neoclassical growth model. As in Aiyagari (1994), they interpret the savings instrument as capital that is used by a representative firm as input to a constant returns to scale production function. The interest rate earned by households is thus determined in equilibrium as the marginal product of capital. However, they differ from Aiyagari (1994) in that they allow for the possibility that the production function is disturbed by exogenous stochastic productivity shocks.

Krusell and Smith (1998) wanted to understand how the equilibrium business cycle dynamics of macroeconomic variables in this heterogeneous agent economy compare to those in a corresponding representative agent economy — an important open question at the time. If the macroeconomic dynamics of the two economies were not too different, it would provide some justification for the common practice of studying macroeconomics through the lens of a single representative agent. Answering this question, however, required solving their model, which raised substantial challenges. Even before Krusell and Smith (1998), it was well understood that the relevant state variable in this type of economy is an infinite-dimensional object – the endogenous cross-sectional distribution of households' employment states and holdings of capital.

The best word to describe Krusell and Smith's (1998) approach to this challenge is 'chutzpah'. Perhaps, they thought, all the information contained in the distribution of household wealth is overkill. What if we look for an equilibrium in a smaller space by summarizing the distribution with only a finite-dimensional set of moments? What if we use just one moment, the mean? Lo and behold, it worked, in a very precise sense. They showed that using only the mean of the distribution of capital holdings, households could forecast future interest rates, which are what matter for consumption decisions, extremely accurately. Thus, Krusell and Smith (1998) could approximate the equilibrium with a much smaller and computationally feasible set of state variables.

Krusell and Smith (1998) labeled this finding ‘approximate aggregation’. It arises because in precautionary savings models optimal savings decisions are extremely close to linear, except for households with very little capital. But since the savings decisions of households with little capital matter little for the dynamics of aggregate capital, the dynamics of aggregate capital (and hence the interest rate) depends approximately on only the level of aggregate capital, not the distribution of capital across households.

Using this computational strategy Krusell and Smith (1998) simulate the dynamics of aggregate output, consumption, and investment in a plausibly calibrated version of their model. They find that the dynamics of these variables are virtually indistinguishable from the dynamics of a similarly calibrated representative agent economy. It is important to remember that their finding of indistinguishability between the aggregate dynamics of the heterogeneous agent and representative agent economies is conceptually different from their finding of approximate aggregation. It is relatively easy to construct economies in which approximate aggregation holds but in which aggregate dynamics look different in the corresponding heterogeneous agent and representative agent economies. For example, they show that when the model is modified to better match the empirical distribution of wealth (in part by exploiting the ideas in Hubbard, Skinner, and Zeldes (1995)) the co-movement of consumption and income look very different from the corresponding representative agent economy.

The lasting influence of Krusell and Smith (1998) is remarkable. It has turned out that approximate aggregation is far more applicable than one might have thought and has been used in a number of other contexts in papers published in the JPE. For example, a variant of the Krusell and Smith (1998) algorithm was used by Khan and Thomas (2013) in the context of an economy with heterogeneous firms and by Favilukis, Ludvigson and Van Nieuwerburgh (2017) in the context of a model with fluctuating aggregate house prices.

The JPE has played an essential role in fostering the growth of the study of macroeconomics with heterogeneity. I hope, and predict, that the journal will continue to play such a role in the future.

References

Aiyagari, S. Rao. 1994. “Uninsured Idiosyncratic Risk and Aggregate Savings.” *Quarterly Journal of Economics*. doi:10.2307/2118417.

Attanasio, Orazio P., and Guglielmo Weber. 1995. “Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey.” *Journal of Political Economy* 103 (6): 1121–57. doi:10.1086/601443.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh. 2017. “The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium.” *Journal of Political Economy* 125 (1): 140–223. doi:10.1086/689606.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. “Precautionary Saving and Social Insurance.” *Journal of Political Economy* 103 (2): 360. doi:10.1086/261987.

Imrohoroglu, Ayse. 1989. “Cost of Business Cycles with Indivisibilities and Liquidity Constraints.” *Journal of Political Economy* 0022–3808 (6): 1364–83. doi:10.1086/261658.

Khan, Aubhik, and Julia K. Thomas. 2013. “Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity.” *Journal of Political Economy* 121 (6): 1055–1107. doi:10.1086/674142.

Krusell, Per, and Anthony A. Jr. Smith. 1998. “Income and Wealth Heterogeneity in the Macroeconomy.” *Journal of Political Economy* 106 (5): 867–96. doi:10.1086/250034.

Lucas, Robert E. Jr. 1987. *Models of Business Cycles*. London, U.K.: B. Blackwell.

Rosenzweig, Mark R., and Kenneth I. Wolpin. 1993. “Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India.” *Journal of Political Economy* 101 (2): 223–44.

Zeldes, Stephen P. 1989. “Consumption and Liquidity Constraints: An Empirical Investigation.” *Journal of Political Economy* 97 (2): 305–46. doi:10.1086/261605.