

The Spice of Life: Allowing for Heterogeneity in Macro Models

By Kartik B. Athreya and Renee Haltom

Macroeconomic models often make simplifying assumptions that suppress differences between people. Changing these assumptions and allowing for “heterogeneity” can enrich analysis of both observed data and policy trade-offs. Modern macroeconomics has made significant progress in this area, sometimes significantly altering economists’ views and policy prescriptions.

Until about 20 years ago, macroeconomic analysis often started with a model in which people had access to a rich array of competitive markets, particularly for insurance arrangements to protect themselves from any misfortune they might encounter. This is known as the “complete-markets” framework.¹ Together with other common assumptions, this framework produces two notable results: first, outcomes are “efficient” in the sense that there are no alternatives in which all households can be made better off; second, complete-markets models produce no inequality in the well-being of households that start from identical positions and have identical tastes, even in the wake of shocks such as disability, accidents, or even just car trouble.

Complete markets are, of course, far from realistic. Consider the resources that would be required simply to operate the markets needed to “fully insure” a population against the risk of just one event: getting sick in a setting where each person’s illness affects the total income of society. With two people, there are four possible outcomes for society—one for each combination of who ends up sick or healthy—and hence complete markets would require markets for four possible outcomes. But with just

1,000 people, there are 2^{1000} possible outcomes, making literally complete markets impractical.² Worse yet, even if markets were complete, some people may lack access to the markets relevant for them: think of someone living in a remote area or someone who will be born in the future but will be affected by decisions made now.

Fortunately, this degree of completeness and participation often is not required. Few people lay awake at night worrying that a distant neighbor might catch a cold, since this would only have a negligible effect on societal resources. In other cases, people may have family and social networks that approximate complete markets for current risks and adequately represent future generations.³ For other day-to-day risks, such as car trouble, commercial insurance markets can effectively pool risks and provide good protection.

Nonetheless, the complete-markets framework greatly simplifies economic analysis. It often allows a variety of questions about the macroeconomy to be explored using a “representative-agent” model, which includes one average person to represent the entire population. This simplification allows macroeconomists to introduce other realistic complications—such as intricate fiscal and monetary policies—into their

analyses of aggregate consumption, labor hours, and investment.⁴

However, there are many economic questions for which models with complete markets, even when good at describing aggregate outcomes, are simply not helpful. When economists are interested in the distribution of outcomes across different groups of people—or when there is reason to believe that the differences across people might affect macroeconomic outcomes—it is crucial to have models in which people differ, or are “heterogeneous.”

Economists generally consider two sources of heterogeneity. First are differences among decision-makers’ “endowments” and “preferences.” Endowments include firms’ technological know-how and individuals’ innate characteristics, such as intellect, gender, and health, and, in economic models, any ownership shares people hold in the economy’s firms. Preferences include people’s attitudes toward consumption, saving, and leisure. Both endowments and preferences often are modeled as unchangeable after birth.⁵ The second source of heterogeneity is uninsured risk, especially “idiosyncratic” risk, or risk that is not correlated among individuals, as opposed to aggregate risk, which affects everyone.⁶

If markets and participation were complete, everyone could simply insure themselves against idiosyncratic risk, and there would be no disparity in well-being among people with identical endowments and preferences. Once incomplete, however, these risks become a second source of heterogeneity. For example, if a set of otherwise identical households face the idiosyncratic risk of sickness, lack insurance, and have only a savings account for rainy days, then it is likely that those who get sick would enter retirement with smaller nest eggs than those who don’t get sick.⁷

In the past two decades, macroeconomists have made significant strides in allowing for heterogeneity, especially the kind arising from uninsurable idiosyncratic risk. This *Economic Brief* describes some areas where this development has enriched policy analysis.

Asking Different Questions

Economists have long documented that even people of the same gender and similar age and education tend to experience starkly different levels of income, consumption, and wealth—an observation that is hard to reconcile with complete-markets models and suggestive of uninsurable idiosyncratic risk. Complete-markets models were also unable to address the important presence of governments in insurance provision, such as unemployment benefits and Social Security. In models that assume market completeness and full participation, there would not be an obvious role for either.

Starting in the late 1980s, economists began to incorporate incomplete markets into models, which allowed them to better understand the heterogeneity they saw and to explore new policy questions, particularly those related to social safety nets.⁸ Indeed, a common argument for public insurance programs is that they provide protection that is difficult to purchase in the private sector. For example, some determinants of labor income—like effort—are difficult for would-be insurers to observe, and hence to price and insure. In addition, some forms of insurance are difficult to obtain before the risks actually materialize, which is at birth in many cases, such as children born into poverty.

In the case of unemployment insurance (UI), the standard complete-markets model for studying business cycles—the real business cycle model—was concerned mostly with explaining fluctuations in aggregates such as output, investment, consumption, and total labor hours. In many versions of that model, representative agents respond to shocks by increasing or decreasing work hours; there is no role for the painful severing of household-level labor market relationships that we associate with the term “unemployment.” Yet understanding and evaluating the tradeoffs involved with a publicly funded UI system require models that can capture such phenomena.

In one class of incomplete-markets models of the labor market, agents are modeled as engaged in “search” to “match” with employers, where the search process is costly in terms of time and effort, and

there is no private unemployment insurance. In this setting, unemployed individuals can be unsuccessful in their job search, while the employed can lose their jobs. By modeling the job-search process as dependent on time and effort, such models make it possible to evaluate the provision of UI in terms of the incentive structures it creates. This makes a host of analyses possible: Should unemployment benefits completely offset lost wages no matter how long workers remain unemployed? How much pickier would higher benefits make workers about job opportunities? Would it be good for the productive potential of the economy if skilled workers, in the absence of UI, felt desperate for work and were not picky enough?

The resulting body of research has helped economists quantify the incentive and insurance effects of the current UI system. The seminal early work of economists Gary Hansen and Ayse Imrohorglu (1991), for example, shows that UI benefits can improve welfare. However, their work also quantified ways in which the current system may distort incentives, especially by making it less essential for workers to engage in costly job searches.⁹

With respect to Social Security, a common result in models in which all idiosyncratic risks are insurable is that shutting down the current retirement insurance system provided by Social Security would be beneficial after winners (current retirees) have compensated losers (those who have already paid into the system but not yet received benefits). For example, Shinichi Nishiyama and Kent Smetters (2007) calculate that privatization would be worth \$18,100 of extra resources to each household.¹⁰ A key mechanism in their model is that the system's tax on labor supply means that the returns on Social Security contributions are less than what agents could obtain in private, complete markets.

However, systems such as Social Security also provide a means of risk sharing among households that is not always available in private markets, including insurance against wage shocks and the "risk" of longevity, since benefits are paid until recipients die. In their model, Nishiyama and Smetters find that when

these risks are not insurable via private markets, privatization reduces well-being by \$2,400 per household.¹¹

As these examples show, capturing heterogeneity—in these instances by allowing for uninsurable idiosyncratic risk—allows economists to better account for the data and evaluate policy tradeoffs likely present in the "real world" but absent under complete markets.

Changing Views of Current Policy

In addition to allowing economists to consider some questions that otherwise would be difficult to pose, incorporating heterogeneity can sometimes suggest policies that look more like what we see in the real world.

Consumer bankruptcy law is one such area. The availability of unsecured credit to households expanded considerably in the 1980s and 1990s, and as household indebtedness rose, so did personal bankruptcy rates. This led to questions about how easy it ought to be for people to eliminate their past debts. The benefits and costs of default and bankruptcy are intuitive: they allow households to smooth consumption following negative shocks, but the option to do so can reduce credit access by making lenders hesitant to lend.¹² Yet if markets are complete, the ability to escape obligations would only hinder households' ability to buy insurance in the way they would wish to. Therefore, a complete-markets perspective would prescribe a policy of severe punishments to deter default and bankruptcy entirely.

Clearly this is not what we see in the real world. In particular, households must usually use debt to finance large purchases instead of the more flexible forms of finance that complete-markets would allow, and Chapter 7 of the U.S. bankruptcy code provides a means for households to have debts expunged and start afresh. It also appears that insurance is not always available for the life events that commonly coincide with default and bankruptcy, such as divorce, job loss, and lawsuits. Once we allow for these realities via models with incomplete markets, the intuitive benefits of default and bankruptcy can

be compared with their costs. In a prominent example, Igor Livshits, James MacGee, and Michele Tertilt (2003) compare the U.S. bankruptcy system with a European-style system in which it is more difficult for individuals to escape past debts.¹³ Their model considers households that become heterogeneous in the face of idiosyncratic risks of income loss, divorce, unplanned children, and uninsured medical bills. They find that for the risks Americans face, a more “American” system of debt forgiveness achieves higher welfare, where the extent of those gains depends on the persistence of shocks, the size of the household, and the way earnings evolve with age. This lends logic to the U.S. bankruptcy code.

Another example in which heterogeneity might change the policy prescription concerns the taxation of capital income, such as direct taxes on dividends and capital gains. One focus of tax theory is how to raise revenue by causing the fewest distortions to the economy. A very robust result arising from complete-markets models is that capital taxes should eventually be set to zero. This result may seem surprising, given the fact that capital taxes are extremely common across developed countries. It raises the question of whether policymakers are leaving a clear improvement on the table.

Recent work suggests instead that a model without heterogeneity may not be appropriate for this question. Carlos Garriga (2001) and Andrés Erosa and Martin Gervais (2002) show that once households’ productivity varies with age (with the middle-aged being the most productive, in line with stylized facts), taxes that depend on a household’s age are the optimal way to raise revenue.¹⁴ In this light, positive capital income taxes become a useful tool because they approximate an age-specific tax. A 2009 study by Juan Carlos Conesa, Sagiri Kitao, and Dirk Krueger goes further by adding uninsurable risk in ways that capture many relevant dimensions of U.S. household heterogeneity (income, consumption, and wealth). They too find that a capital income tax is a relatively efficient way of raising revenue.¹⁵ In these examples, allowing for heterogeneity helps economists understand the widespread use of policies we commonly observe.

Heterogeneity, Inequality, and Bad Luck

The introduction to this *Economic Brief* stated that complete markets imply no inequality in well-being among identical people: everyone is fully insured against idiosyncratic risk, equalizing negative shocks that might otherwise leave some groups with greater wealth, consumption, or income. In that setting, the only source of inequality would be innate differences between people’s initial positions, such as their preferences, intellect, and health status. If idiosyncratic risks are not insurable, however, those shocks would be a second source of inequality. Thus, heterogeneity via uninsurable risks also helps us think differently about why inequality arises.

One goal of incomplete-markets research is to distinguish between inequalities resulting from differences in people’s unchangeable characteristics established at birth versus inequalities created by the inability to insure against shocks as life unfolds. That is, models incorporating heterogeneity via incomplete markets help us better understand the extent to which relatively poor people are unlucky rather than simply unambitious or impatient.

This is important because different sources of heterogeneity may imply different potential policy responses. If people’s initial conditions are the ultimate sources of inequality, resources targeted at early life, when those factors are more easily modified, may be in order. If idiosyncratic risks, such as labor market shocks, are more relevant, then resources aimed at retraining or unemployment benefits to smooth consumption may be more effective at reducing inequality.

Moreover, society might place different values on the usefulness of devoting public resources to addressing inequality based on its source. Economic models often treat people’s preferences—such as how patient they are regarding current consumption versus future consumption—as constant over time, but people’s preferences may change with age and education. For that reason, it may be harder to justify the use of public resources to reduce inequality that results from people’s deliberate behavior. (One important exception may occur when people’s prefer-

ences produce spillovers for society, such as through the increased uptake of social programs, in which case it might be possible to use policy to “nudge” people toward a more efficient outcome.) On the other hand, people are sometimes powerless against hard-to-insure risks, such as business closures and disability. It is no surprise that society has tended to devote more resources to buffering such shocks via insurance programs like unemployment insurance, Social Security, and laws like those defining personal bankruptcy.

In arguing that models with heterogeneity via incomplete markets provide more useful information when making policy decisions, a caveat is in order: analyses where the underlying reason for market incompleteness is not explicitly modeled or is a feature of the environment that a policymaker cannot overcome will tend to overstate the benefits of policy interventions. For example, some insurance policies entail a high deductible for making a claim, a construct meant to overcome moral hazard when the insurer cannot perfectly monitor policy holders. As a result, claimants will see their resources depleted relative to those who have been lucky enough to avoid the shock. Yet the incompleteness of insurance in this instance cannot be usefully overcome by government provision of insurance unless the government is in a better position to monitor policy holders. Therefore, policy interventions to address heterogeneity arising from market incompleteness must be approached with care. ■

Kartik B. Athreya is a senior economist and research advisor and Renee Haltom is a writer in the Research Department at the Federal Reserve Bank of Richmond.

Endnotes

¹ This framework is often referred to as the Arrow-Debreu model after economists Kenneth Arrow and Gerard Debreu.

² This example is provided by Andreu Mas-Colell, Michael Whinston, and Jerry Green, *Microeconomic Theory*, Oxford University Press, 1995.

³ For a review of the literature on ways in which people attempt to self-insure through the purchase of assets and reliance on means of “partial insurance” through social networks, see Fatih

Guvenen, “Macroeconomics with Heterogeneity: A Practical Guide,” Federal Reserve Bank of Richmond *Economic Quarterly*, Third Quarter 2011, vol. 97, no. 3, pp. 255–326.

⁴ Economists have identified some conditions under which it may be appropriate to use the representative-agent framework. For example, see George M. Constantinides, “Intertemporal Asset Pricing with Heterogeneous Consumers and Without Demand Aggregation,” *Journal of Business*, April 1982, vol. 55, no. 2, pp. 253–267, and Per Krusell and Anthony A. Smith Jr., “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, October 1998, vol. 106, no. 5, pp. 867–896.

⁵ In one example of how to decide what aspects should be modeled as endowments, when studying the consumption and savings behavior of adults, economists often model people as having different levels of formal education at the outset. But education is, of course, hardly immutable. However, when such decisions are not central to the question under study, treating education as an endowment allows the model to be enriched elsewhere (for example, allowing for richer modeling of aspects of the tax code likely to be important for savings decisions). By contrast, it would not make sense to model education as an endowment in a study of educational investment decisions.

⁶ Therefore, idiosyncratic risks are insurable in principle, while aggregate risks are not since they affect even people selling the insurance.

⁷ Once markets against idiosyncratic risk are incomplete, aggregate risk can create further differences across groups of people. Think of a subset of Americans who experience layoffs due to a plant’s closure, followed by a severe recession that hits the long-term unemployed especially hard. The recession may well hurt the laid-off group more than others who enter the recession employed.

⁸ Models with incomplete markets also have allowed economists to study two key considerations for policy: how policies affect the welfare and decisions of specific groups (for example, the healthy versus the sick) and how their responses to those policies affect the aggregate economy. In both cases, allowing for heterogeneity helps economists better identify who is likely to win and lose from a given policy, which informs the democratic process.

⁹ “See Gary D. Hansen and Ayse Imrohorglu, “The Role of Unemployment Insurance in an Economy with Liquidity Constraints and Moral Hazard,” *Journal of Political Economy*, February 1992, vol. 100, no. 1, pp. 118–142. Later work by Eric Young suggests that caution is warranted because the presence of a UI system lowers the savings of households in ways that lower aggregate labor productivity (people have less capital equipment to work with) and leave them worse off. See Eric R. Young, “Unemployment Insurance and Capital Accumulation,” *Journal of Monetary Economics*, November 2004, vol. 51, no. 8, pp. 1683–1710.

¹⁰ See Shinichi Nishiyama and Kent Smetters, “Does Social Security Privatization Produce Efficiency Gains?” *Quarterly Journal of Economics*, November 2007, vol. 122, no. 4, pp. 1677–1719.

- ¹¹ "Well-being" here measures the gain as perceived by someone who doesn't yet know how much they will earn or how long they will live. Also, for a recent analysis of Social Security under incomplete markets, see Mark Huggett and Juan Carlos Parra, "How Well Does the U.S. Social Insurance System Provide Social Insurance?" *Journal of Political Economy*, February 2010, vol. 118, no. 1, pp. 76–112.
- ¹² For example, see Narayana Kocherlakota, "Implications of Efficient Risk Sharing Without Commitment," *Review of Economic Studies*, October 1996, vol. 63, no. 4, pp. 595–609.
- ¹³ See Igor Livshits, James MacGee, and Michele Tertilt, "Consumer Bankruptcy: A Fresh Start," Federal Reserve Bank of Minneapolis Working Paper No. 617, January 2003. A later version of this paper was published in March 2007 in the *American Economic Review*.
- ¹⁴ See Carlos Garriga, "Optimal Fiscal Policy in Overlapping Generations Models," Universitat de Barcelona Working Papers in Economics No. 66, January 2001, and Andrés Erosa and Martin Gervais, "Optimal Taxation in Life-Cycle Economies," *Journal of Economic Theory*, August 2002, vol. 105, no. 2, pp. 338–369.
- ¹⁵ Their paper is straightforwardly titled, "Taxing Capital? Not a Bad Idea After All!" *American Economic Review*, March 2009, vol. 99, no. 1, pp. 25–48. The title is an amusing answer to the previous conventional complete-markets wisdom embodied by "Taxing Capital Income: A Bad Idea," by Andrew Atkeson, V. V. Chari, and Patrick J. Kehoe, Federal Reserve Bank of Minneapolis *Quarterly Review*, Summer 1999, vol. 23, no. 3, pp. 3–17.

This article may be photocopied or reprinted in its entirety. Please credit the authors, source, and the Federal Reserve Bank of Richmond and include the italicized statement below.

Views expressed in this article are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

