

My research interests are broadly in macroeconomics and household finance. My current research aims to understand how households make portfolio choices between liquid and illiquid assets over the lifecycle. Recent work has shown that the way in which households allocate their savings across assets matters for fiscal policy, as individuals who are wealthy but hold few liquid assets may find themselves constrained and therefore willing to spend more out of a temporary stimulus payment. Further, this distinction between illiquid and liquid assets is important as two of the most important assets in the median household's balance sheet, housing and retirement accounts, are arguably primarily characterized by their illiquidity. To the extent that one would want to understand the potential effects of policies relevant to these assets—such as changing the level of subsidies built into the tax code for retirement accounts or implementing new regulations on lending practices to home buyers—it is important to have an understanding of the trade-off households face between liquid and illiquid savings throughout their lifetime.

My job market paper "Learning and Asset Acquisition Over the Lifecycle" seeks to understand how income risk affects this margin. The typical way of thinking of exogenous income in the literature has been as a combination of a deterministic trend and an ARMA(p,q) process. For tractability this stochastic component is often modelled as a combination of a very persistent AR(1) and white noise. In other words, any earning deviations are a result of shocks that are equally unanticipated by both the individual and an outside observer. An alternative, perhaps more intuitively appealing, modelling choice is to allow for a heterogeneous level of innate skill between individuals, which affects initial earnings and drives long run income growth. Because skill is intangible, it must be inferred over time from repeated income realizations. This way of thinking about income has been supported in the literature through examinations of income data, consumption data, and data on household equity holdings.

I incorporate this richer income process into a two-asset lifecycle model. Individuals save both for a lengthy retirement as well as for precautionary reasons. They allocate savings between a low-return, freely adjustable asset and a high-return asset which is fully irreversible until retirement. As they gain experience, they update and tighten their beliefs about their future earning potential. I compare the lifecycle choices individuals make in this environment to behavior in the standard model, and calibrate each income process to parameters estimated from identical income data.

I show that there are stark differences in savings behavior between the two models, both in quantity and timing. Relative to the benchmark model, individuals in the learning model save more overall throughout the entire lifecycle. They also hold on average a higher level of liquid assets throughout the lifecycle as well, and these holdings tend to peak a few years earlier than in the benchmark model. Moreover, these individuals delay illiquid saving, holding relatively less illiquid assets early in life and more illiquid assets later in life, as compared to their benchmark counterparts. These results arise due to the way learning impacts both uncertainty and the mean forecasts of individuals over the lifecycle.

The potential implications of this paper are numerous. The consumption response of an individual to a one-off stimulus payment depends crucially on the composition of that individual's assets. Individuals with positive savings will only consume large amounts out of a temporary transfer if those savings are difficult to access. What my paper suggests is that by merely modelling income in an alternative, but empirically justifiable, way, individuals are less likely to be constrained as they hold a higher portion of their savings in liquid accounts. Therefore the stimulative effect of any government transfer will be diminished.

Further, illiquid assets are dominated by housing and, to a lesser extent, retirement savings. These classes of assets are continually the subject of policy discussions and homeownership behavior in particular has been of interest since the Great Recession. My research suggests that any counterfactual policy exercise affecting these types of assets will not be invariant to the underlying exogenous income process.

For example, suppose that one wishes to consider the effect of lowering down payment requirements on home purchases. If individuals want to immediately accumulate illiquid assets, as in the standard model, the only constraint from doing so is an inability to cover the down payment. Thus, individuals will buy into the illiquid account as soon as they have built up enough liquid assets to afford the down payment and lowering this requirement would have a substantial effect. However, if instead individuals delay illiquid accumulation due to their beliefs and uncertainty about their future earning potential, this change in lending practices will not have as large of an effect. I plan to study the implications of this scenario and similar counterfactuals in this environment in greater detail.

One other promising avenue I am beginning to explore is the way in which the addition of illiquid assets to the individual's choice set affects our understanding of relevant parameter values. Typically, estimates of household parameters assume only non-durable consumption and one completely liquid asset. Since the major component of both durable goods and illiquid assets is housing, by ignoring durable expenditures these approaches undercount the level of actual household savings as well as consumption. Further, to the extent that there are non-separabilities between durable and non-durable consumption, or features of illiquid assets (such as transaction costs) affect the consumption-savings decision of the individual, inclusion of a durable good/illiquid asset into these models may affect parameter estimates. This could potentially yield new insights into savings behavior over the lifetime, the amount of implicit insurance individuals have against income shocks, and the amount of knowledge individuals have about their future income potential.

Finally, I am pursuing a joint work with Vasco Botelho titled "Involuntary Retirement and the Fall in Consumption Growth." Still in its early stages, the paper explores the relationship between hastened retirement and the decline in consumption growth immediately following the Great Recession. Using the Current Population Survey (CPS) and the Panel Study of Income Dynamics (PSID), we seek to investigate the ways in which retirement patterns have changed since the Great Recession and how these outcomes relate to individual labor and income conditions and balance sheets. We aim to determine the extent to which this change in retirement behavior, particularly the increase in involuntary retirement, dampened aggregate consumption growth in the aftermath of the recession.