The Peopling of Macroeconomics:
Microeconomics of Aggregate Consumer Expenditures*

BY SATYAJIT CHATTERJEE

Since the 1950s economists have been building a theory of aggregate consumer spending, seeking to understand how individual households choose to spend and how their choices change when interest rates, the unemployment rate, and other economic indicators change. Before that time, economists looked for "economic laws" that would explain the connection between one set of economic aggregates and another, without considering the decisions of individual households. Although the process of connecting macroeconomic aggregates to individuals' behavior is far from complete, predictions of aggregate consumer spending are now rooted in predictions of individual behavior. In this article, Satyajit Chatterjee takes readers through a brief historical survey from the early work on the consumption function to the theory of aggregate consumer spending in modern macroeconomic models.

Consumer spending is the largest single expenditure category in the final demand for goods and services, accounting for more than two-thirds of

Satyajit Chatterjee is a senior economic advisor and economist in the Research Department of the Philadelphia Fed. This article is available free of charge at www. philadelphiafed.org/research-and-data/ publications/.

gross domestic product (GDP). A clear understanding of the underpinnings of consumer spending is a valuable asset for central bankers and policymakers. Since the 1950s, macroeconomists have been engaged in building a theory of aggregate consumer spending from the bottom up.[1] In this approach, macroeconomists first seek to understand

how individual households choose to spend and how their choices change when interest rates, the unemployment rate, and other indicators of overall economic activity change. The relationship between aggregate consumer spending and indicators of economic activity is then obtained by aggregating the predicted changes in the spending choices of individual households with respect to changes in indicators of overall economic activity.

It was not always so. In the early years of macroeconomics, scholars looked for enduring empirical relationships ("economic laws") that connected one set of macroeconomic aggregates to another without explicit reference to the individual decisions that would make sense of such connections. This was because economists hadn't fully worked out how a household acting rationally in the face of uncertainty would behave over time — the sort of knowledge needed to meaningfully connect macroeconomic aggregates to the millions of individual choices that make up those aggregates. But as economists began to acquire this knowledge, the process of connecting macroeconomics to individuals' behavior started in the 1950s and gathered steam in the 1970s and 1980s. Although the process of integration is far from complete, predictions of aggregate consumer spending are now rooted in predictions of individual behavior.

The attempt to predict aggregate consumer spending by first predicting what individual households would do is

[1] Aggregate consumer spending is total consumer spending in the economy.

what I mean by "the peopling of macroeconomics." The aim of this article is to give an account of this now half-century-long intellectual endeavor. It is meant to be a (quick!) historical survey that takes the reader from the early work on the consumption function to the theory of aggregate consumer spending in modern macroeconomic models.

GENESIS OF THE CONSUMPTION FUNCTION

The origin of macroeconomics as a distinct sub-field of economics is often traced to John Maynard Keynes's *General Theory of Employment, Interest, and Money*. Published in 1936, the book sought to explain the reasons for the economic depression that gripped the industrialized world after 1929. In the course of doing so, Keynes introduced a theoretical construct he called the consumption function. According to Keynes, the consumption function was the causal relationship between annual aggregate disposable (or after-tax) income and annual aggregate consumer spending.

Keynes asserted that this relationship looked like the brown line shown in Figure 1. Aggregate consumer spending was directly and linearly related to aggregate disposable income. The point at which the brown line crosses the black line gives the income level at which consumer spending is equal to income. To the left of this point, spending exceeds income, and to the right of this point, spending is less than income. Importantly, the relationship between income and spending was a nonproportional one, with higher incomes associated with a smaller ratio of spending to income. To see this, consider the points marked X and Y on the brown line. At point X income is $40,000 and spending is $34,000; at point Y, income is $80,000 and spending is $58,000. Thus, a doubling of income leads to less than a doubling of spending, which means that the ratio of spending to income declines as incomes rise.

Because the consumption function was central to Keynes's analysis, the construct attracted a great deal of attention and soon became the focus of controversy. The problem was that Keynes did not explain how the consumption function could arise from the choices of individual households acting rationally. Instead, he defended his construct as a "psychological law" that accorded well with common sense. In Keynes's favor, the construct seemed to accord with some facts as well: Household-level incomes and expenditures were (roughly) related as shown in Figure 1, with higher income households spending more than lower income households but spending proportionately less of their income than lower income households.

But the household-level evidence was not definitive because Keynes's consumption function was supposed to hold for *aggregate* consumer spending and *aggregate* disposable income measured at *different points in time*. The issue remained unsettled because data on aggregate consumer spending and aggregate income for different years were not readily available. When the data were eventually assembled, they showed a relationship like the brown line shown in Figure 2. Over a long period of time, the relationship between consumer spending and income was proportional. As illustrated by the points X and Y, a doubling of (per capita) income from $40,000 to $80,000 leads to a doubling of (per capita) consumer spending from $32,000 to $64,000 — something that is not true of the consumption function in Figure 1. Although one might be tempted to gloss over this difference, the difference was important: Keynes's theory assumed that the consumption function looked like the one in Figure 1, not like the one in Figure 2.

The puzzling difference between consumption-income relationships "across households" (*cross-section*) and "across time" (*time-series*) became the focus of macroeconomic research in the 1940s and 1950s. By that time, many economists had accepted Keynes's *General Theory* as being essentially correct, and it became a matter of some urgency to understand why these relationships differed and how both could be true at the same time. Progress came in the form of two studies that pretty much set the stage for research on the aggregate consumption function for the next 30 years. One was by economist Franco Modigliani and the other by economist Milton Friedman. Both contributions earned their progenitors Nobel prizes: Friedman in 1976 and Modigliani in 1985.

RATIONAL CHOICE: AN ENGINE FOR PREDICTION

Both Friedman and Modigliani focused on understanding the relationship between spending and income at the household level, and both sought
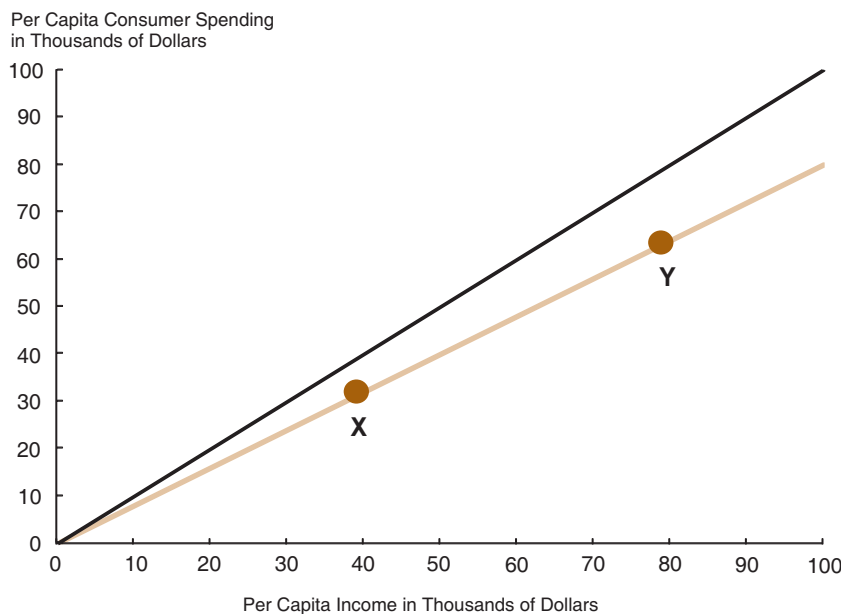
> According to Keynes, the consumption function was the causal relationship between annual aggregate disposable (or after-tax) income and annual aggregate consumer spending.

FIGURE 1

Keynesian Consumption Function

Consumer Spending
in Thousands of Dollars

Income in Thousands of Dollars

FIGURE 2

Relationship Between Spending and Income Over Time

Per Capita Consumer Spending
in Thousands of Dollars

Per Capita Income in Thousands of Dollars

to achieve this using the model of rational choice. Rational choice — the dominant paradigm for thinking about behavior in economics — posits that people make decisions to maximize their well-being subject to the limitations imposed on them by the finiteness of resources. For instance, the theory of rational choice prescribes how a family should allocate its finite income among competing uses in order to generate the maximum possible well-being for the family.

Modigliani's Work. Modigliani and his student Richard Brumberg began by studying a very simple individual choice problem. They imagined a young adult starting out on his working life at age 20 (say) knowing (with perfect certainty) that he would live up to a given age (say, 90 years), knowing how many of those years he would work (say, 40), and knowing how much he would earn in each period of his working life (say, $35,000 each year). Modigliani and Brumberg assumed that the young adult obtained the same physical and psychological benefit (or utility, as economists call it) from any given amount of spending in any given year. They also assumed that, as is customary in economics, the benefit obtained by the adult from an additional dollar of spending declines with the amount already spent that year: That is, the first dollar spent in any year gives more benefit than the second dollar spent in that year and the second dollar spent in that year gives more benefit than the third dollar spent in that year and so on. Finally, they assumed that the adult could borrow or save at a bank at a zero interest rate.

The question they asked was: What is this individual's best lifetime spending plan? The answer is that the individual should spend his average *lifetime* income of $20,000 each year, where $20,000 is the sum of his in-

come over his working years ($35,000 multiplied by 40) divided by the number of years he will live (70, which is 90 less 20). Because the individual gets the same benefit from spending in each year of his life and because every dollar spent gives less benefit than the previous dollar spent that year, it is best for the individual to spend the *same* amount every year.[2] And if he is to spend the same amount every year and live within his means, he must spend his average lifetime income each year.

Even though the example was highly unrealistic, it served to show that Keynes's consumption function (or "psychological law") had no obvious basis in rational choice. If we could observe this hypothetical individual over time, we would see his income change from $35,000 to zero when he retires and yet we would see his spending stay unchanged at $20,000 a year. Contrary to Keynes's assertion, an increase or decrease in income need not be accompanied by an increase or decrease in spending. In a rational choice context, current spending need not respond to a change in current income if that change is fully anticipated in a previous period.

While these findings raised doubts about Keynes's "psychological law," they did not resolve the issue of the de-

scriptive realism of the "law." Perhaps the "law" was a better description of reality than rational choice. To be truly convincing, proponents of rational choice had to show that their theory explained the facts better than other alternatives. To prove their point, both Friedman and Modigliani concentrated on reconciling the differences between cross-section and time-series consumption-income relationships.

> Contrary to Keynes's assertion, an increase or decrease in income need not be accompanied by an increase or decrease in spending.

Modigliani and Brumberg's simple model is consistent with the differences seen in the data between time-series and cross-section consumption-income relationships. In their model, economic growth causes everyone's average lifetime income to grow over time. Since everyone spends their average lifetime income, economic growth also causes aggregate spending to grow at the same rate as average lifetime income. Therefore, spending and income grow in proportion to each other. In contrast, the relationship between income and spending across people alive at any point in time will be necessarily nonproportional because even people without any income (retirees) spend a positive amount (for more details on this point, see *Reconciling Secular and Cross-Section Consumption Functions*).

Of course, this is a simple example, and one might wonder whether the rational choice paradigm would predict these relationships in more

realistic situations. The answer to this question is a resounding yes, and the person most responsible for showing why was Milton Friedman.

FRIEDMAN'S PERMANENT INCOME HYPOTHESIS

In 1957 Friedman published a monograph titled *A Theory of the Consumption Function*. As an enduring example of the interplay between economic theory and facts, the treatise has few equals.[3] Friedman distinguished between a household's permanent income and its actual income and — with the help of rational choice theory and empirical facts — argued that a household tends to spend its permanent income.[4]

Friedman defined permanent income as the amount a household could spend and still maintain its wealth. To understand what this definition means, it is helpful to think of some simple examples. First, imagine a household, such as a new retiree, that in terms of resources has only financial wealth. Suppose that a household has a million dollars in the bank, and the interest rate available at the bank is 5 percent. Then, this household's annual *perma-*

---

[2] To see why, imagine that the individual plans to spend $50,000 in 2008 and only $40,000 in 2009. Because the additional benefit from each dollar spent is declining with the total amount spent, the benefit obtained from spending the 40,001st dollar in 2008 is more than the benefit obtained from spending the 50,000th dollar in 2008. Since the benefit obtained from spending the 40,001st dollar in 2008 is the same as the benefit obtained from spending the 40,001st dollar in 2009, the individual can increase his total benefit by reducing his expenditures by $1 in 2008 and increasing it by $1 in 2009: The loss in benefit in 2008 will be more than compensated by the gain in benefit in 2009. This sort of logic can be applied repeatedly to conclude that the best the individual can do is spend the same amount each year.

[3] To quote Friedman's Nobel citation: "From a purely scientific viewpoint, one of Friedman's most important contributions is his reshaping of consumption theory with the help of the hypotheses about 'the permanent income', in place of current annual income, as a decisive factor in determining total consumption expenditure. Here an extremely fruitful distinction is made between households' temporary income and more permanent income; Friedman shows that a substantially larger part of the former income is saved than of the latter. Friedman has carefully tested this theory on comprehensive statistical material and gained interesting results. Friedman's version of the consumption function has had a lasting effect both on theory and on empirical research."

[4] This is simplifying matters somewhat. Friedman's permanent income hypothesis is the assertion that a household's planned level of spending will be some proportion of its permanent income, where the proportion could fluctuate around unity over time.

nent income is $50,000 — the amount the household would earn in interest and therefore could spend without reducing or augmenting its (financial) wealth.

The example above imagined a household, such as a new retiree, with only financial wealth. What about a young household that has no financial wealth but expects to earn income for many years into the future? Suppose a household expects to earn $40,000 in each of the next 20 years and $60,000 in each of the following 20 years (after which it retires). Suppose it can borrow from the bank against this income stream at an interest rate of 5 percent. Then it is as if this household has financial wealth of (roughly) $820,000 in the bank today — which is the discounted value of the household's stream of future earnings.[5] Then the same logic as above applies, and the household can spend about $41,000 annually — which is (roughly) the annual interest earned on $820,000 — and still maintain its wealth.

Why would a household wish to spend its permanent income? Note that when thinking about how much a household should spend from one month to the next, it is fine to imagine that a household's circumstances are similar from one month to the next. Thus, all else being the same, the household should spend the same amount each month. Second, note that even though a household will exist for a finite length of time, for practical purposes it is fine to imagine that there is no natural end to the household's planning horizon. This

may be because the end is really far away or because the household cares about its descendants and its descendants' descendants and so on, so that there is literally no end to its planning horizon. Thus, a household's decision problem is to use a finite amount of wealth to provide for spending over infinitely many future months. The only way this household can spend the same amount each period forever is to spend the constant interest earned on

> A household's decision problem is to use a finite amount of wealth to provide for spending over infinitely many future months.

its financial wealth each period, that is, spend its permanent income.

One can see how Friedman's permanent income theory could account for the proportional spending-income relationship over time and the nonproportional relationship across households at a point in time. If people's perceptions of their permanent incomes rise with the general rise in living standards, everyone's spending will rise in proportion to the rise in living standards. But if we look at households at a point in time, there will be some households whose income is temporarily above their permanent income, and those households will save most of the additional income; there will also be households whose income is temporarily below their permanent income, and they will draw down their savings to maintain their consumption. Therefore, the relationship between in-

come and spending across households at a point in time will naturally tend to be nonproportional.

Friedman was aware that this approach to consumer spending needed to be amended when uncertainty about future earnings is taken into account. Because a household cannot perfectly forecast its future earnings and because banks do not lend against the promise of uncertain future earnings, there is no way for a household to actually convert its future income stream into an equivalent amount of financial wealth. Nonetheless, Friedman maintained that there must be some notion of permanent income to which household spending is adapted. The level of this permanent income will be household specific and will depend on such things as the household's expected earnings and the household's perception of future earnings risk as well as the household's stock of financial assets.

Although uncertainty about future earnings played a key role in Friedman's theory, the implications of such uncertainty for rational choice were only dimly understood at the time. Friedman did not provide a rigorous foundation for his ideas. The result was that while macroeconomists quickly accepted the distinction between actual and permanent income, they ignored Friedman's assertion that permanent income was something not directly observable. Instead, they took permanent income to mean the annual interest earned on the sum of financial and human wealth, where human wealth was calculated as the present discounted value of current and future *expected* earnings. A key reason behind the adoption of this particular definition was the discovery — made in the 1960s — that under certain conditions the theory of rational choice implied that households should set current

---

[5] Discounted (or present) value refers to an amount of money today that will become a given amount at a stated point in the future, depending on the interest rate. For example, if the interest rate is 10 percent, $100 today will be worth $110 one year from now. So the present value of $110 one year from now (when the interest rate is 10 percent) is $100.

Reconciling Secular and Cross-Section Consumption Functions

S uppose that each individual lives for two years. An individual works for the first year of his life and enjoys retirement in his second and final year. Each year, one one-year-old is "born" and one two-year-old "dies," so that the total population is always constant at two. There is growth in incomes over time: Every year, newborns earn 20 percent more than the previous year's newborns.

TABLE

Income and Spending in a World of Overlapping Generations

	Year 0		Year 1		Year 2		Year 3	
	Inc.	Spend.	Inc.	Spend.	Inc.	Spend.	Inc.	Spend.
Gen 0	100	50	0	50				
Gen 1			120	60	0	60		
Gen 2					144	72	0	72
Average	…	…	60	55	72	66	…	…

The table records the relevant data for this hypothetical economy. In the table, columns represent either income (Inc.) or spending (Spend.) for a particular year. The generation born in year 0 is denoted Gen 0, the generation born in year 1 is denoted Gen 1, and so on. Thus, under the income column for year 0, there is an entry for 100 in the row representing Gen 0 because that is what the person born in year 0 earns in that year. Moving across the same row, the entry under the spending column in year 0 is 50 because that is what Gen 0 spends in year 0 (the rest of his or her earnings are saved). Continuing to move across, the corresponding entries for year 1 are 0 and 50, respectively, because Gen 0 retires in year 1 and earns nothing but spends 50 in year 1 (this spending is financed by savings accumulated in year 0). Finally, there are no entries for Gen 0 for years 2 and 3 (and beyond) because Gen 0 is not alive in those years. Moving down to Gen 1, there are no entries for year 0 or year 3, since Gen 1 is not alive in those years. For year 1, the entry under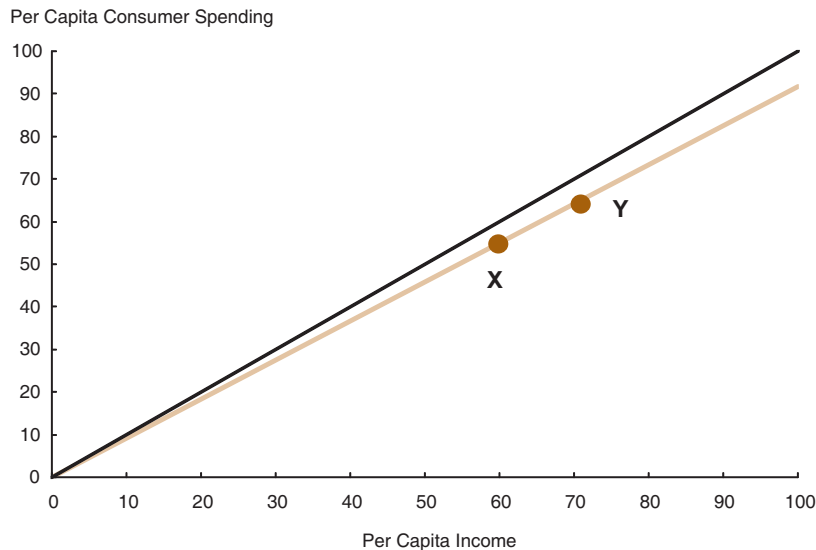 the income column is 120 because Gen 1 earns 20 percent more than Gen 0. Gen 1 spends 60 in year 1, and this is recorded under the spending column for year 1. For year 2 the corresponding entries for income and spending are 0 and 60, respectively. The situation is similar for Gen 2. Gen 2 earns 20 percent more than Gen 1 in year 2 and spends half of his earnings in year 2 and the remaining half in year 3. Naturally, there are no entries for Gen 2 for years 0 and 1.

We can use the snapshots of the overlapping generation world displayed in the table to reconcile the shapes of the consumption functions across households (cross-section) and across time (time series). First, let's look at how aggregate per capita income and spending evolve in this economy. The bottom row of the table reports the average income and spending in each of the years for which these averages can be computed from the information reported in the preceding rows.

Let's look at year 1. Aggregate per capita income in year 1 is simply income averaged over the two individuals alive in year 1. The two individuals alive in year 1

FIGURE A

Per Capita Spending and Income in the Overlapping Generations Example

Per Capita Consumer Spending

[Graph showing two lines from origin, with points X at approximately (60, 55) and Y at approximately (70, 64). Y-axis labeled 0 to 100 in increments of 10, X-axis labeled Per Capita Income from 0 to 100 in increments of 10.]

Per Capita Income

FIGURE B

Spending and Income Across Households in the Overlapping Generations Example

Household Spending

[Graph showing a steep line through origin and a nearly flat line with point X at approximately (0, 51) and point Y at approximately (120, 59). Y-axis labeled 0 to 120 in increments of 10, X-axis labeled Household Income from 0 to 120 in increments of 10.]

Household Income

are Gen 0 and Gen 1. Gen 0 has no earnings in year 1 (because he or she is retired) and Gen 1 earns 120 units. So, the average income in year 1 is 60 units (the sum of 0 and 120 divided by 2). Similarly, the aggregate per capita spending in year 1 is 55 units (the sum of 50 and 60 divided by 2). Thus, aggregate per capita consumer spending in year 1 is 11/12 of aggregate income in year 1.

In year 2, the two individuals alive are Gen 1 and Gen 2. Aggregate per capita income is 72 units (the sum of 0 and 144 divided by 2), and aggregate per capita spending is 66 units (the sum of 60 and 72 divided by 2). Once again, aggregate per capita consumer spending is 11/12 of aggregate per capita income. Figure A plots aggregate per capita spending and income at successive points in time for this economy. As is evident, income and spending grow in proportion to each other over time exactly as found in the data.

Next, let's look at the cross-section consumption-income relationship in this economy. Let's pick year 1. Gen 0 has no income and spends 50 units, and Gen 1 earns 120 units and spends 60 units. Therefore, the cross-section consumption-income relationship for year 1 looks like the one in Figure B. This relationship is clearly not proportional and, in fact, resembles the consumption function in Figure 1 in the text. If we were to pick a different year, say, 2, we would get a similar nonproportional relationship except that it would be shifted upward because of income growth.
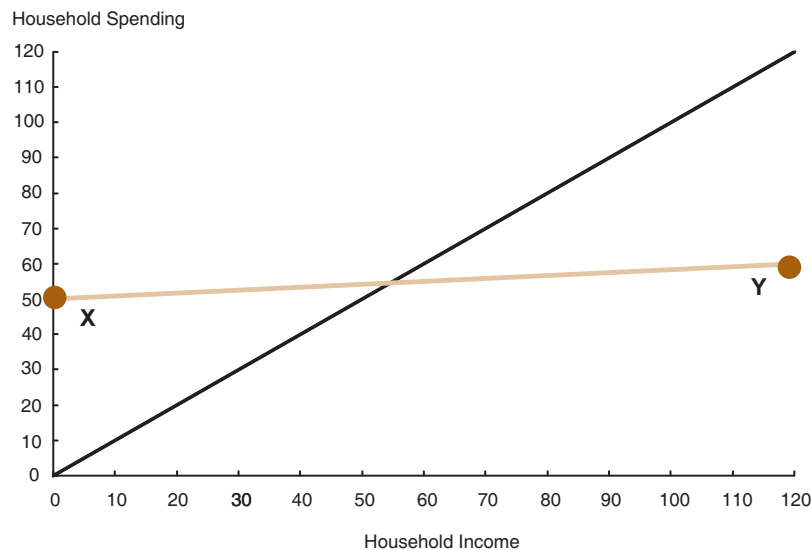
spending equal to permanent income calculated in this way. This resulted in this particular definition becoming commonly used, and eventually, the very idea of permanent income became associated with this particular definition.

But this interpretation of permanent income turned out to be inconsistent with the evidence. As more extensive aggregate and household-level data became available for macroeconomists to analyze, it was found that consumer spending responded too much to transitory deviations in income from permanent income defined in this way to be consistent with the theory's predictions.[6]

UNCERTAINTY, BUFFER-STOCK SAVINGS, AND SPENDING DYNAMICS

While macroeconomists were busy testing the permanent income theory against aggregate and household data and finding it wanting, others were concentrating on working out the implications of rational choice for decision-making over time when the future could not be perfectly forecast. The big hurdle here was that it was not easy to divine the full implications of rational choice theory because the theory's predictions could not be reduced to a simple formula. Consequently, it was not easy to figure out if some version of rational choice theory could explain the data on household spending and income better than the permanent income theory.

Two key developments eventually allowed progress to be made. The first development was something not intrinsically connected to economics. It was the increasing availability of

(and access to) high-speed computers on university campuses and the concurrent rapid development and standardization of computer languages designed to express and solve difficult numerical problems. Along with a deeper understanding of the nature

of rational choice over time, the rapid improvement in the hardware and software for numerical computations permitted macroeconomists to pose, solve, and simulate rational choice problems on the computer.

The other, more important development was connected with the progress of economics as a discipline. To solve rational choice problems on the computer, one must specify the problem in exact numerical form. For instance, it is no longer sufficient to assert (as Modigliani and Brumberg did) that the benefit from an additional dollar of spending declines with the amount already spent; it is necessary to specify how much it declines at any given level of spending. In other words, the computer needs to know the exact numerical relationship between the benefit from an additional dollar of spending and the level of spending.

It took macroeconomists decades to gather this kind of knowledge. The process was helped by the fact that the rational choice paradigm had become

the *lingua franca* of economics — the common language economists use to make sense of behavior in diverse branches of economics. For instance, one source from which macroeconomists learned of the numerical relationship between the benefit from

> While macroeconomists were busy testing the permanent income theory against aggregate and household data and finding it wanting, others were concentrating on working out the implications of rational choice for decision-making over time when the future could not be perfectly forecast.

an additional dollar of spending and the level of spending was researchers trying to understand how fluctuations in expected rates of return on financial assets affected the growth rate of consumer spending.

Through this process, it became possible for macroeconomists to explore the implications of rational choice for consumer spending using computer simulations. Christopher Carroll and Angus Deaton were among the pioneers of this research effort. Their simulations revealed that households that start without any financial assets initially consume less than their earnings in order to accumulate a *buffer* stock of savings.[7] They do so because savings can protect the household from temporary shortfalls in earnings. Since earnings in any period (a year or a month) are uncer-

---

[6] See the article by Robert Hall and Frederic Mishkin.

[7] The discussion in the rest of this section draws on Christopher Carroll's article.

tain, there is value to having such a buffer stock of savings. Then, once the household accumulates its target buffer stock of savings, it acts to maintain that stock over time. Unexpected increases in income are initially saved but then gradually spent to bring the stock of savings down to its target level. Similarly, an unexpected decline in earnings is initially met by a reduction in the stock of savings (as the household tries to maintain spending), but then the resulting deficit in its stock of savings is gradually made up over time.[8] Finally, starting at about age 50, behavior undergoes a significant change: While still working, households rein in their spending and begin to accumulate additional savings to provide for their retirement.

The behavior revealed by these computer simulations has a simplicity to it that gives it a ring of truth. But what makes these predictions compelling for macroeconomists is that the simulations also explain why spending's response to transitory fluctuations in earnings can be larger than that predicted by the permanent income theory. When households are working toward accumulating their target level of buffer-stock savings, their spending is depressed. The simulations reveal that in these circumstances a household that receives an unexpected transitory increase in income has an incentive to boost spending from its depressed level. This happens because the extra income is used to augment the household's savings, and therefore, the household gets closer to (or

achieves) its target level of buffer-stock savings. Consequently, the incentive to curtail spending in order to get to the target level of buffer-stock savings is attenuated or eliminated, and spending responds strongly to a transitory increase in income. This effect is absent in the permanent income theory because households do not curtail their spending in order to accumulate a buffer stock of savings.

Interestingly, the simulations also reveal that the predictions of the permanent income theory continue to be relevant once a household reaches its target level of savings. As the buffer-stock of savings is approached, households act more like future uncertainty does not matter – just as the permanent income theory had assumed. Of course, the household behaves this way because it has accumulated a buffer stock of savings to counter the risk of lost earnings. That being said, it is important to note that because households are continually buffeted by shocks to earnings and are therefore accumulating or drawing down financial assets, the fraction that behaves according to the permanent income theory in the simulations is a minority.

CYCLICAL IMPLICATIONS OF BUFFER-STOCK SAVINGS MODELS

Macroeconomists and policymakers are interested in what Deaton's and

Carroll's spending models have to say about aggregate consumer spending and savings as well as the movement of these aggregates over the course of business cycles. As one might suspect, the only way to get answers to these questions is by computer simulations. But the simulations are no longer about the behavior of a typical household but the aggregate behavior of an entire *ensemble* of households,

an ensemble whose summed behavior has measurable effects on the cyclical behavior of market prices and interest rates. Since the cyclical behavior of market prices and interest rates, in turn, affects the behavior of each household in the ensemble, the challenge for the simulation is to properly account for the feedback from behavior to market prices and back to behavior.

The "feedback" problem prevented macroeconomists from analyzing the business-cycle implications of buffer-stock savings behavior until, in an important paper, Per Krusell and Anthony Smith showed how the problem could be solved. They developed a procedure for reliably compressing the amount of information required by the computer to keep track of feedback effects. With this innovation, macroeconomists are now able to simulate the behavior implied by rational choice of a large ensemble of interacting households living through expansions and recessions.

The simulations reveal that cycli-

> When households are working toward accumulating their target level of buffer-stock savings, their spending is depressed.

---

[8] It is worth noting that macroeconomists were aware that rational choice theory was consistent with households' accumulating assets in order to meet a potential shortfall in earnings in the future. What the simulations revealed — and this came as a surprise — was the centrality of precautionary or buffer-stock savings in the household's spending decisions.

cal fluctuations in aggregate consumer spending and aggregate income are more tightly linked than the permanent income theory implies. This makes intuitive sense: The tighter link is a consequence of the fact that those households whose spending is depressed because they are in the process of accumulating their target level of buffer-stock savings will increase their spending more when income is temporarily high (as it is in an expansion).

CONCLUSION

Macroeconomics studies the structure and performance of an economy as a whole. Although the founding documents of economics have a decidedly macroeconomic

focus — Adam Smith wrote about the wealth of nations, after all — the development of economics as a modern discipline has been a long and arduous effort to understand and predict the behavior of individual decision-making units, such as households and business firms.

For historical reasons, macroeconomics began life with a rather tenuous connection to the principles of rational choice, in part because John Maynard Keynes explicitly rejected rational choice – and its correlate of competitive markets – as a framework unsuitable for explaining the Great Depression. But it was also because of the broad scope and general complexity of the subject matter; it is a field

that invites theorizing at the macro rather than at the micro level.

But fortunately for the development of macroeconomics, there was one very important point of contact between macro- and microeconomics, namely, the consumption function. To make sense of this function, macroeconomists had to think seriously about individual behavior. And so began the "peopling of macroeconomics." The process has gone on now for more than 50 years, and, to quote Angus Deaton, it has "generated some of the best science in economics." This article has endeavored to give a glimpse of this fascinating and ongoing intellectual journey. ⑱

REFERENCES

Carroll, C. "Buffer Stock Savings and the Life-Cycle/PIH.," *Quarterly Journal of Economics*, 112 (1997), pp. 1-56.

Deaton, A. "Savings and Liquidity Constraints," *Econometrica*, 59 (1991) pp. 1221-48.

Friedman, Milton. *A Theory of the Consumption Function*. Princeton: Princeton University Press, 1957.

Hall, R.E., and F.S. Mishkin. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, 50 (1982), pp. 461-81.

Krusell, Per, and Anthony Smith. "Income and Wealth Heterogeneity and the Macroeconomy," *Journal of Political Economy*, 106:5 (1998), pp. 867-97.

Modigliani F., and R. Brumberg. "Utility Analysis and the Consumption Function: An Interpretation of Cross-section Data," in Kenneth K. Kurihara, ed., *Post-Keynesian Economics*. New Brunswick, N.J.: Rutgers University Press, 1954, pp. 388-436.