Finance for the Rest of Us

James J. Choi
Yale University and NBER

Talk delivered at The University of Texas at Austin, McCombs School of Business
November 2, 2022

Nearly 40 years ago, Apple introduced the Macintosh computer to the world with the tagline, "The computer for the rest of us." This was the computer that didn't require you to be a highly trained computer geek to use. This was a computer for ordinary people. Apple has ridden this philosophy of making products for ordinary people to becoming one of the most valuable companies in the world.

A common theme running through the rise of behavioral finance since the 1980s and household finance since the early 2000s is the ambition to bring the academic discipline of finance closer to ordinary people—to create a finance for the rest of us. Rather than studying hyper-rational agents who control large corporations and have enough wealth to move security prices, behavioral and household finance are all about agents with human foibles, agents who aren't corporate titans or Wall Street traders or billionaires, but rather are people with modest or little wealth.

As I think about where behavioral and household finance could fruitfully go in the next decade, I believe we would do well to push our field to become even closer to ordinary people.

I specifically see two ways we could do that: first, by listening to what ordinary people have to say, and second, by making finance more practically useful for ordinary people. I'll illustrate what I mean by this using some of my own recent work. But my main ambition for this talk is not to tell you about my research, but to spark new research ideas for you. So throughout the talk, I'm going to mention research questions that my recent work on these two fronts has caused me to become curious about, but which I haven't had time to pursue or which I don't know how to pursue. My hope is that some of you will have your curiosity piqued, and you'll write papers that answer some of these questions.

Before I continue, let me go on a brief digression about choosing research topics, since I've been told that many of you in the audience are Ph.D. students.

Richard Hamming was a 20th century applied mathematician who won the Turing Award, which is considered the Nobel Prize of computer science. He gave a wonderful talk called "You and Your Research," the transcript of which you can find online by Googling, or by going to the Teaching tab of my webpage. I re-read this talk about once a year to reorient my research agenda.

Hamming tells a story that took place in the Bell Labs lunchroom. One day, he decided to sit at the chemistry table. He says this:

> I went over and said, "Do you mind if I join you?" They can't say no, so I started eating with them for a while. And I started asking, "What are the important problems of your field?" And after a week or so, "What important problems are you working on?"" And after some more time I came in one day and said, "If what you are doing is not important, and if

> you don't think it is going to lead to something important, why are you at Bell Labs working on it?" I wasn't welcomed after that; I had to find somebody else to eat with! That was in the spring.
>
> In the fall, Dave McCall stopped me in the hall and said, "Hamming, that remark of yours got underneath my skin. I thought about it all summer, that is, what were the important problems in my field. I haven't changed my research," he says, "but I think it was well worthwhile." And I said, "Thank you Dave," and went on. I noticed a couple of months later he was made the head of the department. I noticed the other day he was a Member of the National Academy of Engineering. I noticed he has succeeded. I have never heard the names of any of the other fellows at that table mentioned in science and scientific circles. They were unable to ask themselves, "What are the important problems in my field?"
>
> If you do not work on an important problem, it's unlikely you'll do important work. It's perfectly obvious. … The average scientist, so far as I can make out, spends almost all his time working on problems which he believes will not be important and he also doesn't believe that they will lead to important problems.

What are the important problems in finance? What important problems are you working on?

Larry Summers famously said that it takes just as much time to write an unimportant paper as it does to write an important paper, so you might as well work on important topics. I'll add to that that being convinced in your own mind of your paper's potential importance makes it a lot easier to endure the hardships of research and the publication process. So work on papers you think could be important!

The wonderful thing about working in household finance is that it's easy to see the connection between your research and issues that matter to real people, which means that it's easy to think about how your research could potentially be important.

It's that desire for a connection to real people that caused me to recently write a couple of papers that directly asked people about the motivations and beliefs that drove their personal financial decisions.

Asking people directly about why or how they make their economic decisions has had a bad rap in economics ever since Milton Friedman told his famous—or infamous—parable about the expert billiards player in 1953. You know, the one where he argues that the best theory for predicting an expert billiards player's shots is one that assumes that the player acts as if he knows the mathematical laws of physics. Friedman writes, "The billiards player, if asked how he decides where to hit the ball, may say that he 'just figures it out' but then also rubs a rabbit's foot just to make sure." According to Friedman, for the purpose of creating and testing theories, it's useless to ask economic agents about how or why they made their decisions. The only thing that matters is the accuracy of the theory's predictions.

And yet despite economists' officially stated aversion to asking people how and why they do the things that they do, we seem to find the answers that people give when we *do* ask them to be pretty darned fascinating. Ironically, our *revealed preference* is that these direct questions are useful. Lintner (1956) remains the paper that has had the most influence on our profession's understanding

of corporate dividend policy. What did he do? He interviewed executives at 28 companies on how they set dividend policy! More recently, Graham and Harvey (2001) surveyed 392 CFOs about how they set their firms' financial policies, and that paper has been cited over 8,000 times. In the daily lives of economists, we're asking people around us why they did things all the time. Why did you decide to study finance? Why did you move to Austin? Honey, why are you upset? We believe the answers to these kinds of questions to be quite informative. It's strange that the answers to the same kinds of questions would suddenly become completely uninformative when they are asked for the purpose of an economics research paper.

We're also coming to a better understanding of how revealed preferences are often not nearly as revealing and informative as we would like. The modern standard for empirical identification is to exploit exogenous variation, ideally truly randomized variation. It's usually not possible to run randomized experiments in economics. Even if we are able to run a randomized experiment, or to find a very convincing instrumental variable that provides exogenous variation, we now understand that the parameter that's being estimated in these very special setups is not what we would usually like to know, which is the average treatment effect in the entire population. We get the average treatment effect in the special population that happened to be subject to the randomized experiment, or we get a local average treatment effect on compliers to the instrumental variable.

The critique, "What about external validity," is kind of a cheap shot because it can be leveled at *any* empirical study. The 18th century Scottish philosopher David Hume noted that even though the sun has risen in the morning a million times before, there is no strictly rational basis for expecting that it will rise again tomorrow. So every claim to external validity is subject to a leap of faith, an extrapolation.

Nonetheless, cheap shot or not, the question of whether a finding has external validity is a reasonable one. This is one reason why there has been a growing hunger in empirical economics research to identify not just treatment effects, but mechanisms. We can maybe assess the extent to which a given estimated treatment effect will port over into a new setting by assessing whether the posited mechanism is operable in the new setting.

But mechanisms are hard to identify! We're so hungry to identify mechanisms that we're willing to lower our empirical identification standards to measure them. It's not uncommon to see researchers argue for the existence of a mechanism by examining the heterogeneity of a treatment effect within a population. But that's not identification by exogenous variation of the hypothesized mechanism's strength! The variable along which the heterogeneity is measured could covary with not just the hypothesized mechanism but a million other things.

Finally, most outcomes can be "explained" by multiple combinations of preferences, beliefs, and/or constraints. By assuming rational expectations, we can cut down on the number of combinations that fit the data, but I think it's getting harder and harder to really believe that people—especially ordinary people with no financial expertise—have rational expectations in the financial domain. And even if we stay in the rational expectations sandbox, people can suffer from various reasoning fallacies, have exotic preferences, or have hidden constraints. So we still have a lot of flexibility in "explaining" what we see in the data.

A defender of Friedman's position might say, so what? If multiple theories can explain the observed outcomes, then each one is as good as the other, so we can feel free to choose the one

that is most aesthetically pleasing to us. The problem with this is that when the circumstances facing the agents change to something not already in the data, it's unlikely that all of the theories that fit the old data will have the exact same prediction of how the agents will act in the new circumstances. The theory whose assumptions are closest to the true beliefs, decision processes, and constraints of the agent doesn't logically *have* to be a better predictor out of sample, but it feels to me like it would have a better chance of being the winner.

The nice thing about asking respondents about their decision-making processes is that we can get some direct information on mechanisms. We don't necessarily have to 100% believe what they tell us, but if a ton of people say that they made a decision in a particular way, it feels hard to not have your priors moved towards thinking that this is an important mechanism. Second, getting more free-form responses from people can inspire new hypotheses about which beliefs, preferences, and constraints matter. We become less constrained by our imagination as economists who, frankly, have been a bit brainwashed into a certain way of thinking. Third, we can ask respondents to do the casual inference for us. How important was this factor for the action you took action last quarter? Instead of relying upon exogenous variation, somebody who has the blueprints of a machine has a reasonable chance of knowing what would happen if a certain gear in the machine were changed or removed. I think psychologists and economists have good reason to be somewhat skeptical about asking people to predict what they will do in the future under certain circumstances. But asking why people did something in the past is an easier question to answer. Parker and Souleles (2019) find that when survey respondents were asked how their spending responded to the 2008 economic stimulus payments, the responses corresponded well with econometric estimates of the spending response that exploits the quasi-random timing of the payments. So it's not crazy to think that people do a reasonable job of explaining why they did something in the past.

Even if the "as if" assumption ends up holding, so that an individual or an organization acts as if it is doing something quite different from what the individual or organization believes it's doing, isn't it really interesting to understand how an individual or an organization gets to that "as if"? Economists have been oddly not curious about this. The process by which somebody or a group of somebodies gets to "as if" might give hints about transition dynamics if the environment changes. It could even be that there's some sort of path dependence in how agents got to "as if" that would prevent them from adequately adapting to certain types of changes in the environment.

I think it's also interesting to know what individuals and organizations believe about themselves. If they're deluding themselves about their true nature and what they end up actually maximizing, that seems to be a social science phenomenon worthy of study.

I don't have hard numbers to back me up on this, but I do feel like there's been a growing acceptance within economics of research papers that directly ask economic agents about their motives, beliefs, and decision processes.

Asking people can't provide definitive evidence. At the end of the day, I stand with the traditional economist view that actions are what ultimately matters most, not words. So survey responses aren't going to get us all the way there. But what we've learned over the decades is that observational data rarely get us all the way there either. There are, as the epistemologists say, different ways of knowing, and we can learn from multiple channels.

Brogaard, Engelberg, and Van Wesep (2018) is a paper published in the *Journal of Economic Perspectives* titled, "Do Economists Swing for the Fences After Tenure?" One of the defenses of tenure is that it allows scholars to take risks they otherwise wouldn't take with less job security. But Brogaard et al. conclude that no, economists don't take more research risks with home run potential after tenure. We just get less productive and have fewer right tail outcomes as a percent of total production.

With that critique haunting me, a few years ago, I decided to start a project with Adriana Robertson—then a J.D./Ph.D. student at Yale, now on faculty at the University of Chicago—where we asked a representative sample of U.S. adults how well leading academic theories describe their concerns, preferences, and beliefs in the domains of the percent of their portfolio they held in stocks, actively managed mutual funds, and the cross-section of stock returns. I thought, this is a little unusual methodologically, and maybe it'll have a hard time getting published, but this is the type of risky paper I should try writing post-tenure.

To my surprise, I have never had an easier time getting a paper published than with this paper. Choi and Robertson (2020) was solicited by the *Journal of Finance* and accepted with minimal revisions. I joked to Adriana that I was afraid she was getting an unrealistically rosy picture of the academic life, and she needed to go out and experience some soul-crushing rejections soon in order to debias her expectations.

I want to talk just briefly about a few findings from this paper on how people chose the fraction of their portfolio to hold in stocks.

One of the factors that people say was most important for determining their portfolio equity share is the fear of an economic disaster like the Great Depression. I had long been skeptical of rare disaster explanations of financial phenomena because—rare disaster being rare—the theory is conveniently hard to test in data. But this survey finding has caused me to become much more sympathetic to the possibility that the fear of rare disasters is really quite important. And these fears need not be rational, nor do agents have to properly work through in their minds what an equilibrium outcome of such fears should be.

Giglio et al. (2021) measure subjective disaster probabilities using investor surveys. They find that subjective expected returns of the stock market are *negatively* related to beliefs about how likely a disaster is. In other words, when people think that a disaster is more likely, they think the expected return on the market going forward will be lower. This holds both across investors and within investors over time, and it's completely the opposite of the economists' story of why rare disasters are supposed to matter. When rare disasters are thought to be more likely, investors should demand a higher premium to hold the stock market, so their expectation of the market's return should be higher, not lower.

Giglio et al. also find that subjective disaster probability is more related to investors' asset allocations than investors' beliefs about stock return variance.

It's hard to see how we could have discovered these sorts of stylized facts without asking investors.

Another set of factors that are highly rated in importance are factors related to time until expenditure—time left until retirement, the need to finance routine expenditures, the time remaining until a non-retirement expenditure. I'll come back to this topic shortly.

An interesting factor that bubbles up to the top ten is what we call finance phobia. Specifically, we described this factor to survey respondents as the statement, "I don't like to think about my finances." About 35% of stock market non-participants say that not liking to think about their finances is a very or extremely important reason why they don't invest in the stock market.

I was recently talking to a rather prominent economics professor who was a refugee as a child, and therefore grew up with a lot of financial insecurity in her family. The result of that early-life experience is that she still feels a lot of anxiety about financial decisions, which she deals with by just not acting. She was mistakenly put into the wrong target date retirement fund in her retirement savings plan, so that she has a very conservative asset allocation, and she still hasn't acted to correct her asset allocation because of her anxiety around finance.

When Adriana and I put this factor into our survey, we thought it would be easy to find papers to cite about finance phobia, because it intuitively feels like such a common phenomenon. Instead, we had a hard time, and what we ultimately found were a couple papers in relatively obscure outlets. Why are economists not studying this? Ordinary people are telling us that this is really important.

The last factor I'll mention is religion. About a quarter of respondents say that their "religious beliefs, values, and experiences" were very or extremely important in determining the fraction of their portfolio invested in equities. As a religious man myself, my reaction to this high percentage is, "What? Why? It would never occur to *me* that my religious beliefs would influence the fraction of *my* portfolio that's invested in stocks."

Now there is a large body of empirical evidence that Catholics are more risk-seeking than Protestants. The explanation in the literature for why this is is that some Protestant denominations prohibit gambling. But gambling at a casino is quite a different thing than investing in the stock market. I'd love to know more about what this 25% of people are thinking when they say their religious beliefs, values, and experiences played an important role in determining their equity share. To start, somebody—maybe in this audience—should just ask a large number of this 25% to give a free-form response to how exactly their religion played this role. That would provide the specific hypotheses for a broader empirical investigation. We want to be investigating important questions, and there are few social forces as powerful as religion.

We can listen directly to people. We can also listen to the people whom the people are listening to. This was the idea behind Choi (2022), where I surveyed the advice being given by the 50 most popular personal finance books, written by people like Dave Ramsey, Suze Orman, and the founder of Vanguard, Jack Bogle.

These authors are probably way more influential than professional economists. Take Dave Ramsey. He's sold 1.5 million copies of his book *Total Money Makeover*. 18 million people listen to his radio show every week. The uploads of his radio show episodes to YouTube have been viewed over 600 million times. A recent job market paper, Chopra (2021), exploits the staggered geographic diffusion of the Dave Ramsey radio show to estimate that exposure to the show and its

anti-consumption, pro-saving message reduces household retail spending by at least 5.4%, which is a huge effect.

What do these authors advise? I'll highlight just a few things that I think are the most promising to investigate in future research.

When it comes to making asset allocation decisions, investment horizon is king for these authors. There's a firm conviction that stocks become less risky as the holding period increases. Dave Ramsey writes, stocks "are lousy short-term investments because they go up and down in value, but they are excellent long-term investments when leaving the money longer than five years" because "one hundred percent of the fifteen-year periods in the stock market's history have made money."

Jack Bogle writes, "I have almost no idea how to forecast these short-term swings in investor emotions… [but] largely because the arithmetic of investing is so basic, I have been able to forecast the long-term economics of investing with remarkably high odds of success." Four books in my sample say that stocks are "on sale" after a large price decline, explicitly indicating a belief that stock market returns are mean-reverting.

What really concerns popular authors is the possibility of *selling* an investment at a loss. Vicki Robin and Joe Dominguez write, "The only days you care about an investment's value are the day you buy it and the day you sell it." If your investment is underwater but you can hold onto it, it's not a problem, or at least much less of a problem.

A consequence of this belief is that popular authors bucket money by when it might be spent. Money that might be spent soon should be held *entirely* in cash. This encompasses emergency savings, where the typical recommendation is to hold three to twelve months of living expenses in an emergency account. Some books also include non-emergency savings that will be needed in the near-term as money that should be held entirely in cash. "Near-term" encompasses a wide range of horizons, up to ten years, but the median horizon mentioned is five years—any money you might spend in the next five years, you should hold entirely in cash. It's only money that you won't spend in the "near term" that you should consider investing in equities.

To take a detour into what financial economic theory has to say about this matter, Samuelson (1969) and Merton (1969) proved that if stock market returns are identically and independently distributed, and you have constant relative risk aversion utility and no labor income, your asset allocation should not depend on your investment horizon. A property of i.i.d. returns is that the annualized variance of log returns does not vary with the holding period. If you look at the data, the annualized variance of log U.S. stock market returns is pretty flat with respect to holding period.

So where does this popular belief that the stock market is less risky as investment horizon increases come from? I think it comes from popular authors thinking about risk differently than economists think about risk. Popular authors think about risk in terms of binary outcomes. What is the probability that stocks will underperform bonds over the investment horizon? What is the probability that stocks will have a negative return over the investment horizon? These probabilities decrease with investment horizon for any risky asset whose expected return is positive and above the riskfree asset's return. Economists think about risk in terms of variance because we think that the magnitude of underperformance or outperformance matters, not just this binary outcome.

The very common advice that any money that might be spent soon should be held entirely in cash provides a new explanation for the stock market non-participation puzzle. Under *any* expected utility function, in the absence of frictions, an agent should hold at least a small positive amount in every positive expected value gamble—the intuition being that any smooth utility function is locally linear, so agents should be risk-neutral over small gambles. Stocks are clearly a positive expected return gamble, so the fact that so many households have zero stock market exposure is a puzzle. It's been long known that stock market participation increases with wealth, which has caused a fixed cost of stock market participation to become a leading explanation for non-participation. The idea is that if you only have $100 to invest, and the fixed cost of stock market participation is $100 a year, then it's not worthwhile for you to invest in the stock market, but if you have $100,000 to invest, paying that $100 fixed cost becomes worthwhile.

But if you think that any money that you might use in the short-term should be held in cash, then low rates of stock market participation that rise with wealth arise naturally, because people with very little money have a high likelihood of spending any money they do have in the short-term. As far as I can tell, this explanation for stock market non-participation is completely absent from the academic literature. This should be investigated empirically. The tricky thing is finding truly exogenous variation in the timing of expected future expenditures.

Turning to saving, the popular authors emphasize three themes: save early, save consistently, save more over time. In other words, smooth your savings rates over time. Economic theory instead emphasizes smoothing consumption. But popular author David Chilton explicitly writes against consumption smoothing: "Strangely, a few economists and mathematicians have been pushing the idea of intentionally not saving in your early working years because your income is low and your starting-out-in-life costs are high. They advocate ramping up efforts big time in your middle years… Do not heed that advice… it seldom works in the living room. First, costs have a funny way of never stabilizing. Second, most people aren't going to be able to transition from setting aside nothing to being supersavers at the flip of a switch. Psychologically, that's just not realistic. Finally, I can't get the numbers to work anyway."

Tony Robbins writes, "Whatever that [savings percentage] number is, you've got to stick to it. In good times and bad. No matter what. Why? Because the laws of compounding punish even one missed contribution. Don't think of it in terms of what you can afford to set aside—that's a sure way to sell yourself short. And don't put yourself in a position where you can suspend (or even invade) your savings if your income slows to a trickle some months and money is tight."

Dave Ramsey writes, "The worst time to borrow is when times are bad… What would be difficult is to make the payments and even pay off the debt if you don't get that new replacement job." The funny thing is that economists would say that the *best* time to borrow is when times are temporarily bad! Dave Ramsey's way of thinking might explain a phenomenon that Ganong and Noel (2019) document: "new credit card borrowing finances less than 0.5 percent of monthly consumption during unemployment… We think that understanding the limited use of credit card spending during unemployment is an important area for future research."

Popular books frequently give advice on how to get out of debt. Economists don't have much to say about this subject, but for sure, we'd say that except in some special cases where you're close

to defaulting on your debt, you should prioritize paying down the debt with the highest interest rate.

About half of popular authors agree with the economists' advice on this. But nearly as many advocate following the debt snowball strategy instead, which prioritizes paying off your smallest-balance debt first. Dave Ramsey writes, "People sometimes say, 'But Dave, doesn't it make more sense mathematically to pay off the highest interest rates first?' Maybe. But if you were doing math, you wouldn't have credit card debt, would you? This is about behavior modification. You need some quick wins or you will lose steam and get discouraged… every time you cross a debt off the list, you get more energy and momentum…"

A bunch of books say that if you're going to get out of debt, it's important to establish a firm personal rule that you will not borrow any additional amounts. For example, Senator Elizabeth Warren and her daughter write, "This is the moment to look at yourself in the mirror and say out loud, 'The debt stops here.' Every morning tell yourself, 'I will not take on more debt today.'"

The existence of this belief may shed new light on the co-holding puzzle, which is the fact that about a third of households who borrow on their credit cards and pay interest on that debt hold at least one month of income in liquid assets. It seems like these households would be strictly better off by using some of those low-interest-earning assets to pay down their high-interest-rate credit card debt.

Economists have come up with a few explanations for the co-holding puzzle: the utility of cash for transactions, strategic maneuvering before bankruptcy, attempts to reduce household spending by reducing unused credit capacity, or insuring against the risk that one's credit limit will be reduced.

Perhaps surprisingly, many popular personal finance books *endorse* co-holding, but none cite the academic justifications for co-holding. Instead, the most common reason for endorsing co-holding is that it allows you to not borrow additional amounts while you're on your debt repayment plan. Senator Warren and her daughter write, "This [emergency savings buffer] is the money that will keep you from sliding back into the credit card trap when something goes wrong." Dave Ramsey writes that he used to recommend that people use all of their savings to pay down their debt, but he stopped doing that because "If you use debt after swearing off it, you lose the momentum to keep going."

Another group of books say that building assets while paying down debt builds motivation. David Bach writes, "If you were to direct all of your available cash flow to debt reduction…it might literally be years before you could begin saving for the future. This is too negative—so negative, in fact, that many people who follow this path get discouraged, give up early, and never get to the saving part."

And this brings me to the second big theme in my talk, which is that we should try to make finance more practically useful for ordinary people.

I was recently on the Freakonomics Radio podcast to talk about my paper on popular financial advice. The host also interviewed Morgan Housel, the author of the popular book, *The Psychology of Money*. Housel said, "How many people has Dave Ramsey helped out of debt versus the average academic economist? It's a million to one."

And you know what? He's right. But does it have to be this way? Isn't part of our job as economists to help society—and therefore real, ordinary people—manage economic matters better?

Our normative advice at the personal financial level tends to assume that people have infinite ability to stick to a plan, infinite self-control. It's as if we were giving diet and exercise advice, and the only advice we gave was how to eat and train like an Olympic athlete.

In fact, even Dwayne Johnson, also known as The Rock, who is famous for his physique, says he has to have his cheat meals. "For me, the cheat meals—it's like church. You work out hard and once a week, you treat yourself."

In personal finance, what are the cheat meal equivalents that enable us to stick to a plan and have scientific evidence backing up their efficacy? Does the debt snowball actually work to keep you motivated to stay on your debt repayment plan? Does co-holding actually help you pay off your debt?

Imagine if diet recommendations for those struggling with their body weight were only available from TV and radio hosts, popular book authors, and Instagram influencers, and that health professionals had no advice or research in this area. That's kind of the situation we're in with personal finance. The economists have nothing to say, so people turn to Dave Ramsey and Suze Orman.

The best diet is a reasonable one that you can stick with. Health researchers have written papers about what kind of diet is easiest to comply with. Finding: It's not the Atkins diet. Economists should write papers about what kind of reasonable financial plan is easiest to comply with.

We should study what motivates people to stick to a financial plan. How should financial plans be adapted in light of motivational constraints? Popular authors recommend smoothing savings rates because they view savings as a habit, a virtue that's built up by doing it consistently. Is that right? If it is, how does that change optimal lifecycle savings and investment strategies?

There's emerging evidence that living paycheck to paycheck is stressful in a way that affects labor productivity. I imagine that living paycheck to paycheck also feeds the finance phobia I talked about earlier. If your financial situation is bad, you don't like to think about it, which makes you not manage your finances, which makes your financial situation worse.

Sergeyev, Lian, and Gordnichenko (2022) model financial stress. Utility consists of two parts: CRRA utility from consumption, and disutility from labor. But the smaller your cash on hand is, the more financially stressed out you are, so the more disutility you get from labor. They impose a functional form on this and estimate parameters from a survey. Their aim in the paper is to do positive economics—describe what's happening. But suppose we flip the aim and use this framework to write a normative paper about lifecycle consumption. If financial stress is a real thing, to what extent would our normative models stop recommending that young people save very little in order to smooth consumption over time, but instead start recommending that everybody start saving from a young age, like the popular authors do?

Three years ago, I started teaching a personal finance class at Yale. Suddenly, I needed to have an opinion about how my students should be figuring out their optimal personal finance strategy. It was this experience that really drove home for me the fact that most of our normative models in economics are not very useful for ordinary people, even ordinary people who are in the upper percentiles of numeracy like our MBA students.

Take asset allocation recommendations. We have a beautiful formula for the fraction of your portfolio that should be invested in the stock market, which applies for people who have zero labor income and zero non-financial income. In other words, we have a beautiful formula that applies to almost nobody.

Or we have models that apply to more realistic settings, but they're solved numerically. Authors write down a Bellman equation, specify parameters for preferences and the environment, and give us graphs of the solution for this parameter set, and maybe a handful of other parameter sets. But what if these aren't my parameters? What do I do? Unless I write my own code to solve this dynamic programming problem and insert my own parameters, I'm just out of luck. So this class of papers is also not useful for ordinary people.

My talk in the finance seminar tomorrow will propose a middle-ground approach that my coauthors and I call "practical finance," which creates easily computed approximations to optimal choice in realistic settings (Choi, Liu, and Liu, 2022).

Teaching personal finance has also made salient to me that we still don't have good guidance on many first-order personal finance questions. If I have some extra money sitting around, should I use that to pay my mortgage more quickly than the regular payment schedule, or should I invest that money in other financial assets? Amromin, Huang, and Sialm (2007) have a solution for the case where your marginal dollar of mortgage interest is tax deductible, in which case there can be a tax arbitrage to not prepay your mortgage, but to contribute those dollars to a 401(k) instead. But nowadays, with the cap on mortgage interest deductibility, a lot of people aren't in this situation. What should they do?

What are the true financial risks of home ownership? If my house price goes up 10%, I'm not necessarily better off by 10% of the old value of my house, because it may just be that the price of my housing services has gone up 10%, and if I sell my house, I'll just end up having to pay 10% more for my next house. On the other hand, if I move to a geography whose house prices have low correlation with my current house's price, I could indeed be 10% better off. How could an ordinary person go about figuring this out?

How should my financial portfolio's asset allocation change because I own a home?

I find that the savings rate recommendations coming out of lifecycle consumption models are hard to take seriously, because a first-order consideration in real life is taste shocks. If you're getting married next year, an extra $40,000 of expenditure may buy a lot more utility than that same $40,000 spent the following year. Our models almost always assume that your utility function is the same from period to period. How should I figure out how much I should save today given the existence of taste shocks? How do we as econometricians even measure legitimate taste shocks and distinguish them from unwise, out-of-control spending or paranoid miserliness?

How much does the marginal utility of expenditure diminish in deep old age? My 85-year-old self is a stranger to me. Will he want to spend a lot of money, or is he going to be past the point where spending money is enjoyable? This is a super-important question as I consider how much I should be spending early in retirement, or frankly, how much I should be saving leading up to retirement.

These are questions whose answers matter a great deal to people. Shouldn't the amount of effort devoted to answering them by economists be commensurate with their importance?

I've argued that there are two potentially fruitful approaches to go forward in household and behavioral finance: listening to what ordinary people have to say, and making finance more practically useful for ordinary people. I've thrown out a bunch of research ideas along the way, which I summarize here. Let's study finance phobia. Can we empirically document a link between time to expenditure and asset allocation, particularly as it relates to stock market non-participation? What do people have in mind when they say their religion affected their asset allocation? Do people conceive of risk in terms of binary outcomes of outperforming a certain benchmark, and not in terms of variance? Should we smooth our savings rates instead of our consumption because saving is a virtue that needs to be exercised consistently to be built? How do we stay motivated to stick to a financial plan, and do some of these weird strategies like the debt snowball and co-holding actually motivate people enough to outweigh the real monetary costs they impose? What are the properties of psychological stress from finance, and how should that change our normative personal finance recommendations? And finally, there are so many personal finance questions that we face that we don't have answers to. If you can think of a personal finance question your friend or family member might ask you, thinking that you would know the answer as an economist, and you realize you have no idea, then that very well may be a worthy research question to pursue.

What are the important problems of your field? What important problems are you working on?

References

Amromin, Gene, Jennifer Huang, and Clemens Sialm, 2007. "The tradeoff between mortgage prepayments and tax-deferred retirement savings." *Journal of Public Economics* 91, pp. 2014-2040.

Brogaard, Jonathan, Joseph Engelberg, and Edward Van Wesep, 2018. "Do economists swing for the fences after tenure?" *Journal of Economic Perspectives* 32, pp. 179-194.

Choi, James, 2022. "Popular personal financial advice versus the professors." *Journal of Economic Perspectives* 36 (Fall), pp. 167-192.

Choi, James, Canyao Liu, and Pengcheng Liu, 2022. "Practical finance: An approximate solution to life-cycle portfolio choice." Working paper.

Choi, James, and Adriana Robertson, 2020. "What matters to individual investors? Evidence from the horse's mouth" *Journal of Finance* 75, pp. 1965-2020.

Chopra, Felix, 2021. "Media persuasion and consumption: Evidence from the Dave Ramsey Show." Working paper.

Ganong, Peter, and Pascal Noel, 2019. "Consumer spending during unemployment: Positive and normative implications." *American Economic Review* 109, pp. 2383-2424.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Steven Utkus, 2021. "Five facts about beliefs and portfolios." *American Economic Review* 111, pp. 1481-1522.

Graham, John, and Campbell Harvey, 2001. "The theory and practice of corporate finance: evidence from the field." *Journal of Financial Economics* 60, pp. 187-243.

Lintner, John, 1956. "Distribution of incomes of corporations among dividends, retained earnings, and taxes." *American Economic Review* 46, pp. 97-113.

Merton, Robert, 1969. "Lifetime portfolio selection under uncertainty: The continuous-time case." *Review of Economics and Statistics* 51, pp. 247-257.

Parker, Jonathan, and Nicholas Souleles, 2019. "Reported effects versus revealed-preference estimates: Evidence from the propensity to spend tax rebates." *American Economic Review: Insights* 1, pp. 273-290.

Samuelson, Paul, 1969. "Lifetime portfolio selection by dynamic stochastic programming." *Review of Economics and Statistics* 51, pp. 239-246.

Sergeyev, Dmitriy, Chen Lian, and Yuriy Gorodnichenko, 2022. "The economics of financial stress." Working paper.