

# Mystery Solved?

*A Minneapolis Fed economist offers a promising solution to a persistent puzzle. But work remains ...*

## Douglas Clement

Editor

The equity premium—the higher return from stocks than from bonds—is an entrancing puzzle for economists. Standard theory suggests that stockholders should receive perhaps a 1 percent greater return from stocks than from safer bonds, to compensate for the larger risk inherent to equity investing. But the historical reality is dramatically different. Between 1926 and 1999 (to choose a period without recent booms and busts), stocks returned nearly 7 percent more per year than bonds.

This discrepancy between theoretical prediction and empirical fact—revealed by Rajnish Mehra and Edward Prescott in their famous 1985 article, “The Equity Premium: A Puzzle”—has drawn attention from the discipline’s brightest minds for over two decades, but it remains unsolved, one of the most significant and stubborn questions in asset pricing theory. In a 1996 article reviewing efforts to solve the puzzle, the Minneapolis Fed’s Narayana Kocherlakota observed, “The large equity premium is still largely a mystery to economists.”

Subsequent efforts have made headway—Harvard’s Robert Barro, for example, has proposed that aversion to the extreme risk of rare events such as large wars may be sufficient to generate the premium (see the September 2005 *Region* interview)—but no single explanation has been widely accepted. Indeed, not only has the equity premium puzzle confounded repeated efforts at solution, but it also has generated a proliferation of linked riddles. Financial economists have noticed that the equity

premium can be predicted by dividend yields and other variables, though conventional theory suggests that stock returns are too volatile to forecast. And why does the expected equity premium move countercyclically—that is, rising during recessions and declining in booms? Economists don’t know. The volatility of the premium is also countercyclical, but economic theorists have not been able to explain that relation. The Sharpe ratio (the ratio of expected premium to its volatility) is inexplicably countercyclical as well. A riddle wrapped in an enigma.

Still, there is always hope, and a recent effort by Fatih Guvenen, a Minneapolis Fed visiting scholar and assistant professor at the University of Minnesota, has both intuitive appeal and quantitative power—it makes sense and matches reality.

In simplest terms, Guvenen suggests that the equity premium is generated because stockholders require high equity returns to motivate them to assume economic risk from non-stockholders. The latter tend to be lower-income people dependent largely on labor income, which falls dramatically during recessions. To guard against potential collapse in their consumption, they insure themselves by shifting risk—through purchase of low-yield, risk-free bonds—to high-income people who do hold stocks, but the latter demand a high premium for holding volatile equities, selling risk-free bonds (via their firms) and enduring the ups and downs of business cycles.

Does the equity premium result from

the exchange of risk for security?



**Risk**  
shifts to stockholders.  
Non-stockholders are less willing to postpone consumption. They insure themselves from business cycles by purchasing low-yield, risk-free bonds.

**Security**  
shifts to non-stockholders.  
Stockholders can defer consumption, but they require a high premium for holding volatile equities, selling risk-free bonds and enduring the ups and downs of business cycles.



Guvenen's paper, "A Parsimonious Macroeconomic Model for Asset Pricing," SR 434, published as the lead article in the November 2009 issue of *Econometrica*, is thorough, creative and intriguing. Without question, it is not the last word in this line of research. But because it provides a fresh approach to a classic theoretical question, it offers perspective on the power of new techniques to address old problems—with promising results.

### Pillars of the model

Guvenen's explanation for the equity premium and associated puzzles rests on just two features (thus "parsimonious" in the paper's title): "limited participation in the stock market and heterogeneity in the elasticity of intertemporal substitution in consumption (EIS)."

The first is intuitive and empirically evident. Only a relatively small portion of an economy's population holds significant levels of equities, particularly in the pre-1990s time period that Guvenen explores.

The second feature is cloaked in jargon but also rather easy to accept. "Intertemporal substitution in consumption" means shifting expenditure from one time to another or, stated more simply: waiting until tomorrow to buy something you want today.

### In Brief

#### Premium puzzle

- Economists have been unable to explain why stock returns are far higher than bond returns; theory suggests just a 1 percent premium.
- In recent research, Fatih Guvenen argues that the actual 7 percent premium is generated because stockholders require high returns for assuming risk from non-stockholders. The latter tend to be lower-income people dependent on labor income. They insure themselves with risk-free bonds sold by higher-income people who own stocks. The equity premium is necessary to motivate stockholders to assume aggregate risk and endure business cycle volatility.
- Guvenen's model, while not perfect, has intuitive appeal and a strong match to data.

Elasticity in intertemporal substitution is how flexible someone is about waiting to buy, depending on trends in interest rates. If interest rates were high, would you be more willing to delay a purchase—that is, to save? The last bit: heterogeneity just means that people differ in their sensitivity to changing interest rates.

So, in sum, "heterogeneity in the elasticity of intertemporal substitution in consumption (EIS)" means that as interest rates rise, some people are more willing than others to put off buying. For them, the possibility of consuming even more in the future is more attractive than consuming a lesser quantity right now. Other people have a low elasticity: They are unwilling to delay consumption, regardless of interest rate trends.

These two features—and their interaction—are the pillars of Guvenen's explanation for the equity premium, so he devotes some time to demonstrating their empirical validity.

First: limited participation in the stock market. Until the 1990s, more than two-thirds of American households held no stocks at all, while the richest 1 percent held almost half of all U.S. equity. Over the past two decades, more workers established 401(k)s, IRAs and other investment accounts; by 2002, Guvenen reports, half of U.S. households held some stock, but still the vast majority of all equity remained in relatively few hands. Nonetheless, Guvenen studies data prior to 1992 to abstract from the recent rise in participation.

Second: EIS heterogeneity. In other articles, Guvenen has focused specifically on research about elasticity of substitution, and he summarizes here that "these studies find that, by and large, non-stockholders (and the poor in general) have an elasticity of substitution that is very low—close to zero—while stockholders (and the wealthy in general) have an EIS that is higher." Some studies find that the highest-income households have an EIS that is three times that of the lowest-income homes. Others report that the average EIS is below 0.2, but the richest households' EIS is above 1—a factor five times that of the average family. So the data substantiate the notion of EIS heterogeneity.

What explains the heterogeneity by wealth?

Fatih Guvenen



PHOTOGRAPH BY STEVE NIEDORF

Guvenen provides a plausible explanation with empirical backing. For the wealthy, luxuries—nonessential purchases—constitute a high percentage of total expenditure—higher than for the less well off, at least. Therefore, the overall consumption bundle of rich individuals includes many items whose purchase is easy to postpone. Poor people live nearer subsistence: Delaying consumption cuts close to the bone. Therefore, EIS—flexibility in timing of purchases—increases with wealth. “Since stockholders are substantially wealthier than non-stockholders,” writes Guvenen, “this also implies heterogeneity in the EIS across these two groups.”

### Stacking the deck, not

With these pillars in place, Guvenen builds his model economy. To demonstrate that he’s not gaming the outcome—indeed, that he’s willing to stack the deck against it—he chooses a real business cycle (RBC) model, a framework that tends not to be good at generating accurate asset pricing results. If he can obtain strong results with such an unsympathetic model, the explanation would appear to hold considerable power.

In Guvenen’s RBC model, there are two types of economic actors: households that hold stock and those that do not. Non-stockholders generate income through their labor, but that income is volatile. To smooth volatility in consumption, they can purchase risk-free bonds from stockholder households; these bonds provide a steady flow of income even in bad economic times. (At the level of abstraction assumed in the model, Guvenen points out, the “bond market” includes other low-risk assets in addition to actual bonds, such as checking and savings accounts, money market accounts, CDs and the like.) And consistent with evidence, non-stockholders are assumed to have a low EIS while stockholders have a high EIS.

To preview his conclusion, Guvenen uses this model to explore asset price behavior and finds that it does a remarkably good job at matching empirical reality—not perfect, but better than many other contenders. Again, the mechanism that Guvenen says may explain the mystery: In a nutshell, the

equity premium compensates stockholders for taking on aggregate economic risk from non-stockholders who wish to even out fluctuations in their income.

Here’s a fuller explanation. Non-stockholders receive labor income, which fluctuates with business cycles, the general state of the economy, good luck and bad luck. Given their low EIS, non-stockholders have a strong desire to smooth the ups and downs in their consumption, so they buy risk-free bonds from stockholders that provide income even when times are bad. Stockholders are willing to provide that flow of income to non-stockholders because they have a higher EIS—they’re more willing to delay purchases (after all, many of those purchases are luxury items)—plus they have stock dividends to smooth their consumption in the face of labor income fluctuation.

However, this bond market does not eliminate overall economic risk; it simply reallocates it from non-stockholders to stockholders. That risk is the crux of the puzzle’s solution. “In equilibrium,” writes Guvenen, “stockholders make payments to non-stockholders in a countercyclical fashion, which serves to smooth the consumption of non-stockholders and amplifies the volatility of stockholders, who then demand a large premium for holding aggregate risk.”

Thus is born the large (and heretofore puzzling) equity premium.

### Under the hood

Now for some details. Guvenen builds three variations of the same fundamental model, with each variation successively introducing greater flexibility as to how much labor people supply given current wages. He calibrates these model variations by choosing specific values for a number of basic variables, such as capital’s share of income and the discount rate, that are pretty standard for macroeconomic models. His choices for the nonstandard variables—EIS of stockholders (0.3) and non-stockholders (0.1), and the fraction of total population participating in the stock market (20 percent)—are crucial and carefully researched.

The model constructed, he then runs it through

Table 1: Predictions from Different Economic Models, Compared to U.S. Data

	U.S. Data	Benchmark model	Cobb-Douglas model	GHH model
	Column 1	Column 2	Column 3	Column 4
Equity premium (%)	6.17	5.46	2.65	4.21
Volatility of EP	19.4	21.9	15.4	17.4
Sharpe ratio (EP/EP volatility)	0.32	0.25	0.17	0.24
Risk-free interest rate (%)	1.94	1.31	2.87	1.42
Volatility of interest rate	5.44	6.65	4.91	4.10
Price/dividend ratio	22.1	27.2	25.7	24.7
Volatility of log of price/dividend ratio	26.3	26.6	13.6	17.8

Source: Guvenen, Fatih. "A Parsimonious Macroeconomic Model for Asset Pricing," SR 434, Federal Reserve Bank of Minneapolis, August 2009

a computer with data from 1890 to 1991 and examines his models' predictions about the equity premium (and related variables) relative to actual U.S. data on those variables. The closer the fit between model and data—the better theory matches evidence—the stronger the chance that his explanation is what truly lies behind the equity premium.

The benchmark model is one that holds labor supply fixed—regardless of wage rates or income trends, workers provide the same overall level of labor—and with this strong simplification, the benchmark does quite well at matching several features in the data. (See Table 1, columns 1 and 2.) The data show an equity premium of 6.2 percent and the model delivers 5.5 percent, a close match considering that economic theory would predict just 1 percent. Volatility of the equity premium is 19.4 percent in the data and 21.9 percent in the model. "Therefore," writes Guvenen, "the model generates an equity premium with a mean and volatility that are in the right ballpark compared to the data." Other important data points, such as the

Sharpe ratio, risk-free interest rate, price/dividend ratio and consumption volatility, are also matched well by the benchmark model.

Again, a key innovation in this paper is its use of EIS—flexibility in spending versus saving—and so to gauge how important EIS is in generating his good results, Guvenen tries a few experiments. In one trial, he sees what happens when he uses the same EIS figure for both stockholders (the wealthy) and non-stockholders (average folks) rather than giving a far higher EIS to stockholders. The results are considerably worse. The equity premium generated by this variation drops to 2.4—far below the actual figure of 6.2.

In another test, he doubles the non-stockholder figure for risk aversion—people's desire to avoid risk—while holding EIS figures at their preset 0.1 and 0.3 levels and finds that this dramatic change has a minor impact on outcomes, suggesting that the demand of non-stockholders for bonds is largely determined by their EIS, not their risk aversion.

*Continued on page 48*

## Mystery solved *from page 35*

### Elastic labor

The good results from the benchmark model come with a serious caveat: That version of the model assumes inelastic labor supply. But “labor supply choice is central for any serious macroeconomic analysis,” notes Guvenen, so he then presents results from model variations that relax the fixed labor assumption. Columns 3 and 4 of Table 1 on page 35 show the figures generated by two model variations that allow people to vary their labor supply according to prevailing conditions.

The first version, a classic “Cobb-Douglas” model, does a poor job. The equity premium figure is about 2.7 percent, less than half of the result from the benchmark fixed labor model and far below the historical number. Similar poor results are generated for EP volatility, Sharpe ratio, risk-free interest rate and volatility of price/dividend ratio. “Overall,” he concludes, “I view these results as a step backward compared to the model with inelastic labor supply.”

However, another variation that allows for more refined estimates of labor supply “delivers a more respectable equity premium level of 4.2%, with a volatility of 17.4%,” Guvenen writes. (He labels it the “GHH model” after the economists who developed it, Jeremy Greenwood, Zvi Hercowitz and Gregory W. Huffman.) Other figures from the GHH model (and Guvenen tests two versions of it) are generally better than the Cobb-Douglas model, if not as good as the unrealistic fixed labor model.

Guvenen draws attention to one key variable—the volatility of the risk-free rate of interest (essentially, fluctuation in bond yields)—to note that the results of his GHH model are considerably better than those generated by earlier economic models. One well-known model generates volatility of nearly 25 percent, almost five times higher than reality. Guvenen’s model generates a volatility prediction of 4.1 percent, much closer to the actual historical data point of 5.4. “This low volatility,” he observes, “is an important improvement of this model over earlier production economy models.”

### Asset price dynamics and macro variables

How well do these models replicate facts about asset price changes during business cycles? In reality, the

price/dividend ratio is procyclical: It rises when business is good. However, most economic models have had a hard time duplicating this reality. Another historical fact: Stock returns and the equity premium tend to revert to their mean values over time, contrary to conventional economic theory.

Unlike earlier attempts at reproducing these realities, Guvenen’s model works—it generates results consistent with, though not perfectly matching, historical data. In his model, the price/dividend ratio is tightly correlated with economic output, as it is in reality. In addition, the model’s statistics on predictability of stock returns and of the equity premium closely match U.S. numbers, especially over longer time horizons. Still, he admits, the model falls short on some points. “The model is qualitatively consistent with excess return probability, but does not quantitatively capture the total magnitude observed in the U.S. data.”

Guvenen also examines his models’ ability to generate figures consistent with various important macroeconomic variables, such as output, consumption, income and labor.

All three variations do a reasonable job of matching volatility in output (see Table 2). All models, especially the Cobb-Douglas version, overstate volatility in consumption. Investment volatility and labor volatility are underestimated, especially labor volatility from Cobb-Douglas. “To sum up,” he writes, “the macroeconomic implications of GHH preferences are tolerable ... significantly better than the CD model.”

### Delving deeper

In the hunt for an answer to the high equity premium puzzle, economists have generated many hypotheses. Guvenen’s model and the mechanism it embodies suggest that the answer lies in differences among economic actors in their willingness to delay consumption expenditure and in their participation in the stock market. To better understand the forces at work—to see what truly fuels the mechanism—he focuses on the dynamics of consumption by stockholders and non-stockholders, and of financial exchanges between them.

By breaking down consumption and income

Table 2: Predictions for Volatility of Macroeconomic Variables, Compared to U.S. Data

	U.S. Data	Cobb-Douglas model	GHH model 1 (baseline)	GHH model 2 (low labor-supply elasticity)
Output	1.89	1.97	1.95	1.96
Consumption	0.40	0.92	0.78	0.75
Investment	2.39	1.38	1.76	1.86
Labor	0.80	0.07	0.50	0.32

Note: Volatility for consumption, investment and labor are all normalized by that of output.

Source: Guvenen, Fatih. "A Parsimonious Macroeconomic Model for Asset Pricing," SR 434, Federal Reserve Bank of Minneapolis, August 2009

flows, he isolates the two primary sources for consumption and therefore for volatility in consumption: (1) wage and capital income, and (2) net interest payments from stockholders to non-stockholders. (These payments are made indirectly as interest on bonds from the companies whose equity the stock owners hold.) Non-stockholders derive income only from labor wages and bond interest; by definition, they don't own any capital. Stockholders derive income from all sources: wages from their labor, dividends from their equity holdings and (negative income) from the payments they make on bonds.

Analyzing this quantitatively, Guvenen finds that while the vast majority of consumption for both stockholders and non-stockholders is accounted for by the first category of income: wage and/or capital income, the *volatility* in consumption growth that stockholders face is minimally touched by fluctuations in that income source. Rather, their consumption trends are affected mostly by changes in the payments they make to non-stockholders. "The main source of volatility for stockholders comes from the bond market: interest payments," he writes.

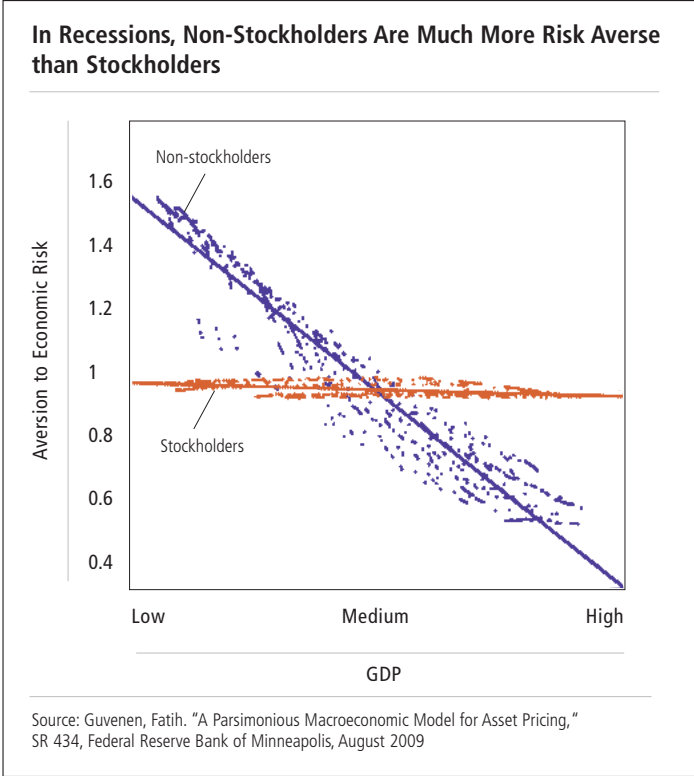
But he points out that "what is really crucial for this extra volatility is the *timing* of the interest payments." They are made precisely when overall economic output (that is, gross domestic product)

drops. "The payments received by the non-stockholders increase exactly when aggregate income falls, i.e., in recessions," Guvenen writes. "Therefore, consumption smoothing for non-stockholders comes at the expense of large fluctuations in stockholders' consumption."

Without these payments, non-stockholders would face a far rougher road. Variance in their consumption "would be 3.13 times higher were it not for the consumption smoothing provided by stockholders through the bond market (that is, if the bond market were shut down)."

Empirical evidence supports this theoretical suggestion of higher consumption instability for equity owners. A 1991 study of U.S. food expenditures found that stockholder consumption was 60 percent more volatile than non-stockholder consumption. A 2002 study in the United Kingdom estimated that expenditures on nondurables and services were 50 percent to 100 percent more volatile for stockholders than non-stockholders. Guvenen also cites research showing that expenditure on luxury goods (expensive French wine, for example, and rents on luxury condominiums in Manhattan) is five to 10 times more volatile than overall spending on nondurables, indicative of "highly volatile expenditures by the very wealthy."





### Why is it countercyclical?

Another piece of the equity premium puzzle is its negative correlation with the business cycle. According to Guvenen’s model, the expected equity premium is 68 percent higher during recessions than in booms. What explains this mystery? Guvenen suggests that differential wealth levels are at the heart of the countercyclical movement.

Stockholders—empirically and in his model—hold far more wealth than non-stockholders do. During recessions, as labor income is battered by layoffs and wage pressures, non-stockholders, with little wealth to fall back on, become very risk averse and place greater value on the income stream provided by their risk-free bonds. “A relatively small change in their wealth holdings changes their attitudes toward risk,” writes Guvenen, “and, consequently, changes their precautionary savings demand significantly.” He illustrates this with a graph that shows risk aversion

levels for stockholders and non-stockholders over the course of a business cycle (see above).

While stockholders’ attitudes toward risk vary little with levels of economic output, “non-stockholders’ risk aversion increases significantly in recessions (and consequently their demand for insurance)”. This increasing risk aversion generates demand for the safety of risk-free interest payments. During recessions, then, non-stockholders buy more bonds from stockholders. Now combine this with the source of the equity premium: stockholders demanding higher equity return to compensate them for smoothing consumption for non-stockholders. The result is that when recessions hit, and non-stockholders increase their demand for the risk-free asset (bonds), stockholders in turn demand a higher premium for holding the economy’s risky asset (stocks). The equity premium is countercyclical. Puzzle solved. Q.E.D.

## Game not over

Well, no. Guvenen's theory about the equity premium and the mechanism he models show considerable promise, but he is the first to admit that it falls short in explaining several aspects of the puzzle.

His model does a good job of matching high equity premium levels, smooth interest rates, procyclical stock prices and countercyclical movement in the equity premium and its volatility, and it proves consistent with real world data on the very small share of wealth held by non-stockholders.

Nevertheless, there are several weaknesses. "A limitation of the present model is that it abstracts from long-run growth," he acknowledges. In addition, his model's predictions for consumption volatility and investment volatility are too high and too low, respectively. "Therefore," he writes, "the present model makes only modest progress in tackling these well-known shortcomings shared by many macro-asset pricing models."

Guvenen's primary research agenda is to explore the nature of labor income uncertainty, "arguably the most significant economic risk that an individual faces during his/her lifetime," he says. In two other recent staff reports (SR 426 and SR 427), he looks at related aspects of labor markets and income patterns—what opportunities and frictions arise when married couples search for job opportunities, and how income distribution evolves over time when people differ in learning ability.

Work on the equity premium puzzle evolves from this broader agenda on income uncertainty—after all, a key part of the mechanism is economic actors reacting to fluctuating income streams. Nevertheless, as he reviews the weaknesses that remain in his equity premium model, he acknowledges that while he has made progress in solving it, the mystery lives on. There remain, he writes, "important areas for potential improvement in this framework." ■