FEDERAL RESERVE BANK of ST. LOUIS
CENTRAL TO AMERICA'S ECONOMY®
WWW.STLOUISFED.ORG

Income Inequality and Monetary Policy: A Framework with Answers to Three Questions[1]

James Bullard, President and CEO

C. Peter McColough Series on International Economics
Council on Foreign Relations
New York, N.Y.
June 26, 2014

Three Provocative Questions

Concerns about economic inequality have been voiced throughout history—Thomas Malthus, David Ricardo and Karl Marx were among the first but they had little systematic data with which they could work. In the 20th century, the advent of national income, tax return and other economic data have allowed for a more rigorous analysis of the issues surrounding inequality. Generally, the focus has been on income inequality, but wealth and consumption inequality are of much interest as well—consumption might ultimately be a more useful variable for assessing economic well-being. Estimates suggest that wealth is much more concentrated than income in the U.S. Consumption inequality is generally thought to be less than income inequality. So the ranking seems to be: The wealth distribution is the most unequal, the income distribution is somewhat less unequal, and the consumption distribution is even less unequal.

Virtually all research shows that U.S. income inequality has increased over the past three decades, but there is much disagreement over the extent of the increase. Major disagreements and controversies arise from different income measurements—pre-tax vs. after-tax vs. after-tax plus in-kind benefits; annual vs. lifetime earnings. A clear message is that measurement matters.

Research also shows that income inequality across countries is considerably more pronounced than within the U.S. Moreover, over the past 50 years, income inequality across countries has declined if one weights countries by their populations. Rapid growth in China, India and elsewhere has reduced global income inequality and lifted many millions out of poverty. For today, I will focus on wealth, income and consumption inequality in the U.S., which is where much of the recent debate has centered.

[1] I thank Cletus Coughlin for assistance in preparing these remarks. Any opinions expressed here are my own and do not necessarily reflect those of others on the Federal Open Market Committee or the Federal Reserve System.

According to a January 2014 Gallup Poll,[2] two of three Americans were either somewhat or very dissatisfied with the distribution of income and wealth in the U.S. This dissatisfaction has led to opinions that government should pursue policies to reduce the income gap between rich and poor. A recent CNN/ORC International Survey[3] found that nearly 70 percent of respondents felt that government should work to substantially reduce the gap.[4]

What might these policies look like? What role might monetary policy play in this debate?

To focus our attention, I thought I would outline three provocative questions concerning monetary policy and income inequality that have repeatedly been asked in the rousing public debate over monetary policy options in the past five years. To keep suspense at its very peak, I do not plan to provide my answers to these provocative questions until the very end of the talk.

Here are the three provocative questions: (1) Does the Federal Reserve's quantitative easing program exacerbate income inequality in the U.S. by putting upward pressure on equity prices? (2) Would a higher inflation target in the U.S. help or hurt the poor? (3) Does current monetary policy hurt savers?

Interesting questions indeed. We need a simple way to think about these issues before some tentative answers can be provided.

Framework

My preferred framework to approach these questions is a simple modification of a life cycle economy, and so I plan to talk through some of the nice features of thinking of the macroeconomic world using this approach. The life cycle model is a workhorse within modern macroeconomics, although it has been less popular in the past three decades than its single household cousin, the representative agent model. The chief advantage of the life cycle framework is that, like the real world, it has plenty of heterogeneity—many different households making many different economic decisions. It also provides a natural and realistic setting for household borrowing and lending, an essential feature if we are to understand the impact of monetary policy on credit markets.

For our purposes here, I can describe the basic outline of this famous framework in just a few sentences. The life cycle concept is that people begin to enter the part of their lives where they

[2] See http://www.gallup.com/poll/166904/dissatisfied-income-wealth-distribution.aspx.
[3] See http://cnnpoliticalticker.files.wordpress.com/2014/02/rel3d.pdf.
[4] This response finding seems to take for granted existing policies, such as progressive taxation, which have been designed to help mitigate income inequality.

make independent economic decisions in their late teens or early 20s. They then live quarter by quarter, making economic decisions about how much to work, consume, borrow and save. They do this until death, which in the U.S. averages around age 80. When people die off, they are replaced in the economy by new entrants, in such a way that, in the simplest versions, the total population remains constant.

The key aspect of the framework for our purposes is the following: Labor productivity varies over the life cycle. We can think of each person as entering the economy with a given life cycle productivity profile which is initially near zero, rises to a peak in the middle of adult life, near age 50, and then declines again to a value near zero. Each person can sell the productivity they have at a particular point in the life cycle in a labor market at the competitive wage per productivity unit, producing income. However, those at the beginning and the end of the life cycle will have very little productivity to bring to the market and hence will have low incomes, while those in the middle of life have a lot of productivity to bring to the market and thus have relatively high incomes. This latter group will be in their "peak earning years." Given these basic features, we will necessarily observe income inequality.

One hardly needs a background in economic theory to accept the basic outline I have just given. Indeed, nearly all participants in the U.S. economy understand at an intuitive level that their ability to earn income will vary substantially as they age.

Income and Wealth Inequality

Very simple versions of this type of model can generate substantial income and wealth inequality without adding anything further to the analysis. Consider the case where the productivity profile begins at zero, rises linearly to a peak at one, and then declines linearly to zero.[5] In this special case, 50 percent of the population would earn 75 percent of the income, that is, there would be a lot of income inequality as an ongoing feature of the economy. In addition, only 25 percent of the population would hold 75 percent of the net assets as an ongoing feature of the economy. Fifty percent of the population—the relatively young—would hold no net assets at all, but would instead be net debtors. Wealth inequality would therefore be substantial and would be even greater than income inequality.[6]

[5] To be more specific, I would have to list many additional assumptions. Those interested in more details may wish to consult Bullard (2014).

[6] This statement equates the wealth distribution with financial asset holding. This will keep the discussion in this speech consistent with popular discussions of wealth. In macroeconomics, the "wealth of the nation" is the value

These types of statistics have a broadly similar flavor to the ones discussed in the contemporary income and wealth inequality debate in the U.S. Yet, while all the figures I cite above are true, there would actually be no income inequality in this economy at all. People are at different stages of the life cycle, and taking a picture of income earners at a point in time—as the figures cited above do, or as a Gini coefficient does—reflects the different productivity inherent in the life cycle. For 20-year-olds their peak earning years are ahead, for 50-year-olds the peak earning years are at hand, and for 80-year-olds the peak earning years are in the past. These people have different incomes today. But looking at their lifetime as a whole, these three groups have exactly the same income if they have exactly the same lifetime productivity profile.[7]

Benign Income and Wealth Distributions

The point of this is to say that the simplest life cycle framework will naturally generate relatively benign income and wealth distributions. These distributions will reflect variable labor productivity over the life cycle, and not more malevolent forces at work. This raises the question of whether the entire observed level of income and wealth inequality in the U.S. could be due to this benign force at work. In other words, can a life cycle model like the one I have described generate income and wealth inequality on the scale observed in the U.S. economy today?

The answer is that the plain vanilla versions of the model I have described cannot give a satisfactory explanation of the observed income and wealth distribution in the U.S. A textbook calculation due to Heer and Maussner (2009) is a sophisticated attempt to find out what a realistic version of this framework has to say about income and wealth inequality.[8] Their calibration of the model generates an income Gini coefficient of about 0.42. A Gini coefficient is a number between zero and one indicating the degree of inequality, with zero indicating perfect equality and one indicating perfect inequality. We want to compare this number with what other researchers think the income Gini is based on U.S. data alone. For this we can consider estimates by Budría Rodríguez et al. (2002), who suggest the U.S. income Gini is about 0.55. We conclude that the model falls short of explaining observed U.S. income inequality. Similarly, Heer and Maussner (2009) find that the wealth Gini generated by the calibration of

---

of the physical capital stock, or, in more sophisticated versions, the value of the physical and human capital stocks added together.

[7] This statement assumes no ongoing economic growth. If there were ongoing growth, the person born later is richer, but most of the contemporary discussion of income inequality is not about this type of inequality.

[8] See Chapter 10.2.2., p. 540 in Heer and Maussner (2009).

their model is about 0.58. Budría Rodríguez et al. (2002) estimate the actual U.S. wealth Gini at 0.78. Thus the model falls short on this dimension as well. One evidently needs something else, something beyond the simple life cycle framework, to explain the levels of income and wealth inequality we observe in the U.S. There are many candidates for this "something else," so I will leave it to you, dear listener, to insert your favorite villain here.[9]

Still, let's not be too dismissive. The basic life cycle model evidently explains an important fraction of the observed U.S. income and wealth Gini coefficients. If you will permit taking ratios of Gini coefficients, the relatively unadorned life cycle model accounts for something on the order of 75 percent of the story of measured income and wealth inequality in the U.S., according to the estimates above.

One might want to think of the level of inequality generated by the life cycle model, as well as closely related estimates, as the natural or ordinary level of income and wealth inequality to be expected in a large capitalist economy with relatively smoothly functioning markets and stable policy. One may want to be especially careful not to disturb this portion of income and wealth inequality through tax policy or monetary policy.

Why do we want to be careful about this?

Shocking Secret

It is because this model also has a shocking secret—shocking at least to the uninitiated. The secret is that smoothly functioning credit markets work to fix the income inequality problem I am describing. If everyone in this economy were to simply consume according to their income—if there were no credit markets—people would consume very little early and late in the life cycle and live like kings in the middle. This means there are powerful incentives for the relatively young—those in their 20s and 30s, say, to take on debt in order to smooth lifetime consumption. There are also powerful incentives for households in their peak earning years to save in order to move income into their retirement years. This happy coincidence creates a market, a fact that forms the foundation of U.S. household credit markets.

How large is this market in the actual U.S. data? According to Mian and Sufi (2011), the household debt-to-GDP ratio in the U.S. has ranged from about 1.15 to 1.65 in recent years. In

[9] One intriguing candidate for a villain has recently been put forward by Greenwood et al. (2014). They investigate how assortative mating—that is, highly educated people marrying other highly educated people—has contributed to increased household income inequality in the U.S. during the post-war era.

today's dollars, this would amount to something on the order of $19 trillion to $28 trillion.[10] That's trillion with a capital "T." So these markets seem to be large indeed, much of it mortgage debt being incurred by the relatively young in order to move housing services consumption forward in the life cycle. This borrowing simultaneously helps peak-earning saver households move income into retirement years where they will need it.

The secret really hits home if you are willing to make enough simplifying assumptions to really get to the core of what this model says about income inequality: In the simplest and most transparent version of the model,[11] all households alive at a point in time would consume exactly the same amount, even though their incomes are radically different. A smoothly functioning credit market would completely solve the income inequality problem I am describing. Consumption inequality would be zero, and so the consumption Gini would be zero. This would be about the best outcome one could hope for, because it would mean that even though income varies widely by household, and even though asset holding differs even more widely by household, actual consumption would even out completely. To the extent that credit markets are doing their job reasonably well, one would not want to distort this life cycle allocation process, and hence one might want to be very careful in trying to design fiscal or monetary policies that might impact U.S. credit markets.

All very well in theory, you say, but is this really what is going on in the U.S. economy? Certainly not in the very extreme form I have described. Still, the life cycle model does tend to predict a lower consumption Gini coefficient relative to the income or wealth Gini, which is true in the U.S. data. This suggests that the framework has some merit. Observed credit markets are surely facilitating considerable consumption smoothing over the life cycle.

In the beginning of this talk, I said that income inequality has been rising over time in the U.S. Could this also happen in a life cycle framework? It certainly could. One might think that those at the very beginning or end of the life cycle are relatively unproductive today, and this situation will not change much over the next 50 or 100 years. For peak earners, however, new technology will likely increase productivity, leading to even higher life cycle peaks in income than we see today. In other words, future technological change will likely benefit the highest income earners rather than the lowest, increasing income inequality. Variations on this theme

[10] For background on how household balance sheets were affected by the financial crisis and related issues, see the St. Louis Fed's Center for Household Financial Stability: http://www.stlouisfed.org/household-financial-stability/. For additional discussion on income inequality, see an upcoming article by Chris Waller and Lowell Ricketts in the Federal Reserve Bank of St. Louis' *The Regional Economist*.

[11] See Bullard (2014).

go by the name of skill-biased technical change in the macroeconomics literature. Recent research by Lansing and Markiewicz (2014) provides a detailed model of how skill-biased technological change can explain increasing income inequality in the U.S. in recent decades. Interestingly, the model suggests all households benefit from the skill-biased technical change, not just those who enjoy higher incomes.

Non-Life Cycle Households

I said that one needs more than the unadorned life cycle model to understand income and wealth inequality in the U.S. What might we add to the simplest versions of the model? There are many possibilities. Decisions to acquire human capital, for instance, would be an excellent addition to the model. We could understand how and why the relatively young might or might not invest in education and thereby increase income (or not) in their peak earning years. In addition, actual borrowing and lending goes through intermediaries, and the U.S. intermediation system has been rocked with controversy since the financial crisis of 2007-2009. Surely a realistic intermediation sector, with all its many dimensions, is important.

But let's focus.

For the purposes of this talk, I want to stress just one addition. It is that not all households in the U.S. are likely to be well-described by the "work every day," "plan-out-your-life" aspects of the life cycle model. Many households instead struggle with attachment to the labor force, working only intermittently, and earning income where and when they can. These households generally tend to have lower incomes, and tend to suffer longer and more frequent bouts of unemployment. Their life cycle plans can frequently be derailed. This group of people tends to rely much more on cash than the life cycle group. Yes, life cycle borrowers and savers use cash and other forms of money, but their most important transactions are accomplished through credit markets. The non-life cycle group uses cash to get by every day. We might proxy this group by the unbanked. According to some accounts, the percent of U.S. households that are unbanked is perhaps near 10 percent, and the nearly unbanked may add to this for a total of as much as 30 percent.[12] This is essentially a relatively poor group of households that is heavily reliant on cash.

[12] See FDIC (2012).

Suppose we add this group to our model. Now we can answer the three provocative questions posed at the beginning of this talk.[13]

Answers to the Provocative Questions

Does quantitative easing exacerbate income inequality in the U.S. by encouraging savers to move into riskier assets, such as equities? Many have suggested that the FOMC policy of buying U.S. Treasury securities and mortgage-backed securities has depressed real yields on relatively safe assets and thus encouraged movement into equities, raising equity prices. It is often said that only 50 percent of households hold equities in the U.S., and they tend to be the wealthiest households; so this policy is making the wealth distribution more unequal.

The life cycle model gives us some perspective on this type of thinking. The framework indeed suggests that relatively older households—only half the population—should hold the lion's share of assets, including equities. In my opinion, equity prices have indeed been influenced by quantitative easing. But I would stop short of saying that this has made wealth inequality worse. The relatively old are going to have to be the domestic holders of the capital stock of the U.S., and they will sell this ownership on to the next generation as they exit the economy. Ideally, when each generation is holding the capital stock, they do so at "normal prices," neither too high nor too low. Actual equity prices were well below normal by conventional valuation metrics in 2008 and 2009, and they have recently returned to more standard valuations. To me, this suggests that quantitative easing had no medium-term implications for the U.S. income or wealth distribution—it is only as good or bad as it was before the crisis.[14]

How about the second question: Would a higher inflation target help or hurt the poorest segment of society? For this question, recall that I added a non-life cycle group to the economy in the previous section. These households rely on cash for much or all of their financial life. They tend to have lower incomes than the life cycle households. Higher average inflation is going to damage the well-being of these households directly. They are holding all of their income each year in the form of cash, unprotected from inflation. A higher average inflation rate directly reduces the value of their financial wealth. While it is true this part of the population tends to have longer and more frequent spells of unemployment, monetary policy cannot influence the average unemployment rate in the medium- or long-term. The answer to

[13] For more perspectives on the intersection of monetary policy and income inequality, interested readers may wish to consult Coibion et al. (2012), Romer and Romer (1998), Gornemann et al. (2012), Airaudo and Bossi (2014) and Gottlieb (2014).

[14] For a sophisticated variant of this thinking generally supporting quantitative easing, see the life cycle analysis of Glover et al. (2011).

this question is that a higher average inflation rate would hurt this poorest group in the economy.

The final provocative question is: Does current monetary policy hurt savers? Many have argued that FOMC policy over the past five years has been to keep real interest rates low, and that these low real yields have impaired the returns of those saving for retirement or in retirement. I have saved this question for last because I think it is the most difficult of the three I have posed here today. In my opinion, Fed policy generally and quantitative easing in particular have influenced the real yield earned by savers. The question is then whether the Fed helped or hurt the situation by pushing real yields lower during the past five years. This hinges on whether credit markets have been functioning smoothly during the period when quantitative easing has been a popular policy. If credit markets were working perfectly or nearly perfectly, then the Fed intervention to push real yields lower than normal was unwarranted and the low real yields were indeed punishing savers. My University of Chicago economics instincts give some credence to this view. At the same time, it seems odd to argue that credit markets were working perfectly or nearly perfectly over the past five years, in the aftermath of one of the largest financial crises the country has ever experienced, and one that was largely driven by mortgage debt run awry. The policy of the FOMC has been that, on balance, low real yields will help repair the damage from the crisis more quickly, and I have largely sided with the Committee in this judgment. As time passes, however, it becomes more and more difficult to argue that credit markets remain in a state of disrepair, and thus harder and harder to justify continued low real rates.

I hope these answers are as provocative as the questions. I appreciate your kind attention and I look forward to taking your questions.

Thank you.

References

Airaudo, Marco and Bossi, Luca. "Trickle-Down Consumption, Monetary Policy, and Inequality," unpublished manuscript, University of Pennsylvania, 2014.

Budría Rodríguez, Santiago; Díaz-Giménez, Javier; Quadrini, Vincenzo and Ríos-Rull, José-Víctor. "Updated Facts on the U.S. Distributions of Earnings, Income, and Wealth," Federal Reserve Bank of Minneapolis *Quarterly Review*, 2002, *26*(3): 2-35.

Bullard, James. "Discussion of Kevin Sheedy, *Debt and Incomplete Financial Markets*," Brookings Papers on Economic Activity, forthcoming, 2014. The related presentation is available at http://research.stlouisfed.org/econ/bullard/pdf/BullardBrookingsDiscussion21March2014_Final.pdf.

Coibion, Olivier; Gorodnichenko, Yuriy; Kueng, Lorenz and Silvia, John. "Innocent Bystanders? Monetary Policy and Inequality in the U.S.," NBER Working Paper No. 18170, 2012.

Federal Deposit Insurance Corporation. *2011 FDIC National Survey of Unbanked and Underbanked Households*, available at http://www.fdic.gov/householdsurvey/, September 2012.

Glover, Andrew; Heathcote, Jonathan; Krueger, Dirk and Ríos-Rull, José-Víctor. "Intergenerational Redistribution in the Great Recession," NBER Working Paper No. 16924, 2011.

Gornemann, Nils; Kuester, Keith and Nakajima, Makoto. "Monetary Policy with Heterogeneous Agents," Federal Reserve Bank of Philadelphia Working Paper No. 12-21, 2012.

Gottlieb, Charles. "On the Distributive Effects of Inflation," unpublished manuscript, Oxford University, 2014.

Greenwood, Jeremy; Guner, Nezih; Kocharkov, Georgi and Santos, Cezar. "Marry Your Like: Assortative Mating and Income Inequality," *American Economic Review (Papers and Proceedings)*, 2014, *104*(5): 348-53.

Heer, Burkhard and Maussner, Alfred. *Dynamic General Equilibrium Modeling: Computational Methods and Applications,* Second edition. Berlin: Springer-Verlag, 2009.

Lansing, Kevin and Markiewicz, Agnieszka. "Top Incomes, Rising Inequality, and Welfare," Federal Reserve Bank of San Francisco Working Paper No. 2012-23, 2014.

Mian, Atif and Sufi, Amir. “House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis,” *American Economic Review*, 2011, *101*(5): 2132-56.

Romer, Christina D. and Romer, David H. “Monetary Policy and the Well-Being of the Poor,” in Proceedings of the Jackson Hole, Wyo. Symposium *Income Inequality Issues and Policy Options*, 1998.