

EconomicDynamics Newsletter

Volume 6, Issue 2, April 2005

The Research Agenda: Dirk Krueger and Fabrizio Perri on Risk Sharing across Households, Generations and Countries

Dirk Krueger is Professor of Economics, especially Macroeconomics at Goethe University Frankfurt (Germany). Fabrizio Perri is Associate Professor of Economics at the Stern School of Business, New York University and currently visiting the Research Department at the Federal Reserve Bank of Minneapolis. They have both worked, often in collaboration, on issues of consumption risk sharing, incomplete markets and distributions of income and consumption. Krueger's [RePEc/IDEAS entry](#). Perri's [RePEc/IDEAS entry](#).

Risk is pervasive in macroeconomics and the question that our research has focused on most is whether, how and to what extent this risk is shared across households or groups of households. Since the risk that a typical household in the macro economy faces is large the welfare impact of sharing it can be substantial. We now briefly describe our research in this area, carried out jointly and with separate coauthors.

Risk sharing across households

It is a well known fact that the distribution of earnings across households is very dispersed. For us, it is crucial to understand whether these earnings differences across households are completely determined at the beginning of the working life of household members (say by their education, skill or endowments) or whether they are the results of idiosyncratic earnings shocks realized during the working life of the members of the household. Recent research (see for example Storesletten, Telmer and Yaron, 2004) seems to indicate that these types of shocks (which we will call earnings risk) are persistent, large and they can be responsible for as much as 50% of the cross sectional variation in earnings. To get a sense of their magnitude, note that household earnings shocks have the same order of persistence as business cycles shocks, but that their percentage volatility has been estimated to be roughly 20 times as large as the percentage volatility of business cycles shocks! Given the sheer size of household earnings risk it is relevant to understand how and to what extent this risk can be shared across households, or to what extent households are at least self-insured.

Theory

One useful benchmark model to assess the extent of risk sharing is the Arrow-Debreu complete markets model. In that model households have access to a full set of state contingent securities for every possible realization of their income so they can fully insure against earnings risk. Several authors have argued (see for example Attanasio and Davis, 1996) that this model overstates the actual risk sharing possibilities available to households, by showing that the complete markets model cannot explain the joint distribution of household earnings and consumption observed in US cross sectional data. Therefore our research on this issue has focused on two popular classes of models that imply only partial risk-sharing or self-insurance of earnings risk.

In the first model (which we refer to as the standard incomplete markets model, SIM) households cannot explicitly share risk with one another, but rather can only self-insure by trading a single, uncontingent bond, potentially subject to borrowing constraints. The second model (which we refer to as the debt constraint model, DCM) follows the work of Kehoe and Levine (1993) and has been further developed by Kocherlakota (1996) and Alvarez and Jermann (2000). In this framework a full set of state contingent contracts is available to all agents, but that intertemporal contracts can only be enforced by exclusion from future intertemporal trade. Since exclusion from credit markets is not infinitely costly, in some states of the world agents might find optimal not to repay their debts and suffer the consequences of exclusion from financial markets. This possibility endogenously restricts the extent to which each contingent asset can be traded and thus limits risk sharing. This is an appealing feature as the extent of risk sharing is not exogenously assumed but depends on the fundamentals of the model (i.e. preferences and technology); for some fundamentals the DCM model generates complete risk sharing, while for different fundamentals the model generate only partial or no risk sharing at all.

Our main theoretical contribution has been the analysis of the DCM model with a continuum of agents. In Krueger and Perri (1999) we show how to characterize and compute stationary equilibria of such a model, using the dual approach developed by Atkeson and Lucas (1992) for private information economies. The consumption dynamics

mirrors the two main assumptions of this model: a complete set of contingent consumption claims and constraints on allocations that require agents to weakly prefer continuation in the market to reverting to autarky. Since it is agents with currently high income whose constraint is binding, agents with high income growth exhibit strong consumption growth, whereas agents with low income are unconstrained and have their consumption decline at a common rate as implied by a perfect consumption insurance Euler equation (we show in the paper that the equilibrium interest rate lies strictly below the time discount factor, making consumption drift down over time when unconstrained).

The crucial friction in this model is the inability of households to commit to repay their state-contingent debt, leading to endogenously determined borrowing constraints whose size depends on how the consequences of default are determined. In the standard limited commitment model this is specified as having to consume the autarkic allocation from the point of default on. While this is motivated by empirical bankruptcy procedures (and can be relaxed by admitting only temporary exclusion or saving after default, as in Krueger and Perri, 1999), it remains true that the consequence of default is specified essentially exogenous to the model. In Krueger and Uhlig (2005) we endogenize the outside option via competition. The outside option of the agent after default is determined by the best consumption insurance contract a household can obtain from a competing financial intermediary, subject to the constraint that the intermediary has to at least break even with the contract. What we also show in that paper is that, even though the extent of consumption insurance depends on the outside option, the consumption *dynamics* is essentially the same as in the DCM model.

Bringing the theory to the data

The workhorse for our empirical analysis is the Consumer Expenditure (CE) Survey which reports data on earnings, hours and detailed consumption expenditures for a fairly large (5000-8000) repeated cross section of US households from 1980 to 2004. One important object we can compute from the data set is within-group income inequality, that is, inequality after controlling for fixed characteristics of the households such as sex, race and education: a statistic of this residual inequality (the variance of logs, say) is the best, if still imperfect, measure of earnings risk we can obtain from the cross sections of the CE.

One striking fact that emerges from the CE is that, over the last 25 years, within-group earnings inequality has increased substantially while within-group consumption inequality increased only very modestly. This fact suggests that US households were able to insulate fairly well their consumption profiles from idiosyncratic earnings risk. In Krueger and Perri (2002) we ask whether the two models discussed above are able to explain this fact, in a quantitatively satisfactory way. We first estimate a time-varying process for earnings risk. Following a large previous literature we model earning risk as the sum of two components: a very persistent autoregressive process and a purely transitory shock. We estimate this process on CE data (we are able to identify the two components due to a short panel dimension of the CE) and find that about half of the increase in earnings risk is driven by the persistent component and half by the transitory component. We then feed this process into both models and find that both predict an increase in consumption inequality substantially smaller than the increase in earnings inequality. Comparing models to consumption data we find that the DCM only slightly understates the increase in within-group consumption inequality while the SIM overstates it.

The DCM predicts very little increase in consumption inequality for two reasons: first households have a full set of state-contingent securities available so they can insure well against shocks even if they are persistent. Second the increased earnings risk makes defaulting and living in financial autarky more costly and thus borrowing constraints (which are the only limits to risk sharing) expand as a response. In other words, the increase in earnings risk makes risk sharing more valuable and the DCM predicts that credit/insurance markets will develop in order to provide more of it.

The reason why also the SIM predicts a more modest increase in consumption inequality, compared to the increase in income inequality, is that even with an uncontingent bond agents can effectively self-insure against the temporary earnings shocks so that the increase in earnings risk due to the increase of the variance of temporary shocks does not translate into consumption. This point was also made by Heathcote, Storesletten and Violante (2004).

Another important difference between the two models is the implication for consumer credit. The development of financial markets generated by the DCM implies a sizable increase in consumer credit that matches up well with what we observe in US data. In the SIM model, on the other hand, the increase in risk implies that households want to accumulate more assets for self insurance and make less use of credit lines. Thus along this dimension that model is less consistent with data as it generates a (small) decline in consumer credit.

In Krueger and Perri (2005) we evaluate the two models along a different dimension. We ask directly how household

consumption responds to earnings shocks, a feature empirically examined by Dynarski and Gruber (1997) with CE data. Our results mirror the ones derived in Krueger and Perri (2002): relative to the data the DCM understates the consumption response to income shocks while the SIM overpredicts it.

From our work we would draw as final assessment of the two models that the DCM model has the appealing feature that risk sharing is endogenous and responds to changes in fundamentals. It may, however, overstate the true insurance possibilities of households, due to the presence of a full set of state contingent securities. On the other hand, the SIM model probably understates the ability households have to insulate their consumption from income shocks and does not capture the fact that credit and insurance markets may evolve in response to change in fundamentals, such as the stochastic income process households face. We conjecture that a model that combines aspect of both models has the most chances of perfectly capturing the empirical facts we have focused on (note that Blundell, Preston and Pistaferri, 2004 come to similar conclusion by following a different methodology).

Welfare and policy implications

After having explored the positive consequences of an increase in earnings risk for consumption we were ultimately interested in the welfare implications of this increase. And if the welfare costs of this increase are large, is there something economic policy can do to reduce them? The two models discussed above provide very different answers to these questions.

In the DCM, in principle the endogenous increase in risk sharing can mitigate the adverse welfare consequences of increased earnings risk. Actually in Krueger and Perri (1999, 2002) we show that there can be situations in which the increase in risk sharing opportunities triggered by the increase in earnings risk is so large that overall consumption risk falls, and welfare rises. In those situations economic policies intended to reduce income volatility (such as unemployment insurance) may have the perverse effect of increasing consumption inequality, because those policies may crowd out the private provision of consumption insurance more than one-to-one. This crowding-out mechanism is similar to the effect at work in Attanasio and Rios-Rull (2001).

On the other hand in the context of the SIM an increase in earnings risk is always welfare reducing as self-insurance can only partially offset it. As a consequence policies that reduce earnings risk are welfare improving. For example, in Conesa and Krueger (2005) we find that in the SIM the optimal income tax code is likely to be progressive because it provides a partial substitute for missing private insurance markets.

Because theory does not give an unambiguous answer to the welfare question, in Krueger and Perri (2004) we use a more empirically guided approach. More concretely, we ask how much would a household in 1973 have been willing to pay to avoid the increase in earnings dispersion that has taken place from 1973 to 2002. In order to do so we first estimate stochastic processes for household consumption and hours worked that are consistent with the evolution of the empirical cross-sectional distributions and with one year consumption mobility matrices from the CE. For consumption we also estimate separate stochastic processes for the between- and within-group component of dispersion, thus capturing both the change in consumption risk and the change in permanent consumption dispersion. Consistently with our previous work we find that the increase in consumption risk has been very modest and thus it has had a very mild welfare impact. On the other hand the increase in between-group dispersion, although not extremely large either, has a much more persistent nature and thus more important welfare consequences. To quantify these we employ a standard lifetime utility framework, together with our estimates of the stochastic processes for the relevant variables. We find that the welfare losses for a substantial fraction of the US population amount to 2 to 3 percent of lifetime consumption and that for some groups (in particular households with low education) the cost can be as large as 6% of lifetime consumption. Heathcote, Storesletten and Violante (2004, 2005) use incomplete markets models to assess the welfare consequences of the recent increase in wage inequality and find numbers comparable to ours. Their approach also captures the interesting effect that, in a model with endogenous labor supply, an increase in wage dispersion raises earnings risk but also raises average earnings, so that the negative welfare impact of higher risk is further mitigated.

Risk sharing across generations

If wages and returns to capital are imperfectly correlated, then there is scope to share aggregate wage and capital income risk across generations. Young households derive most of their income from labor, whereas old households finance old-age consumption mostly via income generated from their assets. If financial markets are incomplete in that households cannot trade a full set of contingent claims on aggregate uncertainty, then a policy such as social security that provides old, asset rich households a claim to labor income, may endow households with welcome risk diversification. In Krueger and Kubler (2002, 2004) we show that even if an economy is dynamically efficient in the

sense of Samuelson's seminal work on the Overlapping Generations model, the introduction of social security may constitute a Pareto-improving reform because it helps to achieve a better allocation of wage/return risk across households. But we also show that for this argument to work quantitatively, shocks to private asset returns have to be as big as return risk to the US stock market, fairly uncorrelated with wage risk and households have to be very risk averse and fairly willing to intertemporally substitute consumption. High risk aversion (a coefficient of relative risk aversion of at least 15) is needed for households to value better risk allocation, while high intertemporal elasticity of substitution is required to keep in check the capital-crowding out effect of social security in general equilibrium. We conclude that, for a realistically calibrated OLG economy the *intergenerational* risk-sharing effects alone are unlikely to provide a normative argument for the introduction of social security. However, Conesa and Krueger (1999) argue that the positive *intragenerational* insurance and redistribution effects from the current US social security system may be sufficient to make a transition from the current system to no social security undesirable for a majority of households currently alive.

Risk Sharing across countries

One type of risk that has received a lot of attention in the macroeconomic literature is country specific aggregate risk. Booms and recessions are not perfectly synchronized across nations. Thus international risk sharing could greatly reduce the costs of business cycles.

However, some early research (Backus, Kehoe and Kydland, 1992) has shown that, in the context of a standard one-good complete markets international business cycles model (IRBC), complete cross-country risk sharing is not consistent with basic business cycles facts, suggesting that international risk sharing might be limited. In Kehoe and Perri (2002) we analyze whether limited enforcement of international contracts could be responsible for limited risk sharing. We characterize and solve the IRBC model with limited enforcement and find that this imperfection can greatly reduce the amount of international risk sharing in the model. We also find that, although the IRBC model with limited enforcement can account for business cycle facts much better than the complete markets model, discrepancies remain between theory and data.

In some recent work (Heathcote and Perri, 2005) we are exploring this issue in the context of a richer model, namely the IRBC model with two goods and with taste shocks. In that context we find that a high degree of international risk-sharing is consistent with several observations for developed economies, especially in the last 10-15 years. In particular for this period, it is consistent with most international business cycle facts (including the relatively low cross-country correlation of consumption), with the proportion of foreign asset in country portfolios (the international diversification puzzle) and with the low observed correlation of the real exchange rate with relative consumption. This suggests that one of the roles of financial globalization (which has happened in the last 15-20 years) has been to improve international risk sharing among developed countries.

What Next

In our empirical work on inequality a crucial component for the evolution of consumption inequality are service flows from consumer durables. Our empirical results also suggest that these services make up a growing share of consumption of households. This motivates us to explore an extension of the limited commitment model that explicitly incorporates consumer durables and collateralized debt, in the same spirit as Fernandez-Villaverde and Krueger (2002). The asset pricing implications of such a model have already successfully been explored by Lustig and van Nieuwerburgh (2004). We intend to use this model to assess to what extent relaxed collateral constraints and improved risk sharing can affect the dynamics of aggregate expenditures on durables over the business cycle, and more concretely, whether these factors have had role in the decline of US Business cycles volatility that many researchers have documented

References

- Alvarez, Fernando, and Urban Jermann (2000), [Efficiency, Equilibrium, and Asset Pricing with Risk of Default](#), *Econometrica*, 68, 775-798.
- Atkeson, Anthony, and Robert E. Lucas, Jr. (1992), [On Efficient Distribution with Private Information](#), *Review of Economic Studies*, 59, 427-453.
- Attanasio, Orazio, and Stephen Davis (1996), [Relative Wage Movements and the Distribution of Consumption](#), *Journal of Political Economy*, 104, 1227-1262.

- Attanasio, Orazio, and José-Víctor Ríos-Rull (2000), [Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?](#), *European Economic Review*, 44, 1225-1258.
- Backus, David, Patrick Kehoe and Finn Kydland (1992), [International Real Business Cycles](#), *Journal of Political Economy*, 101, 745-775.
- Blundell, Richard, Luigi Pistaferri and Ian Preston (2005), [Consumption Inequality and Partial Insurance](#), mimeo, Stanford University
- Conesa, Juan Carlos, and Dirk Krueger (1999), [Social Security Reform with Heterogeneous Agents](#), *Review of Economic Dynamics*, 2, 757-795.
- Conesa, Juan Carlos, and Dirk Krueger (2005), [On the Optimal Progressivity of the Income Tax Code](#), *NBER Working Paper* 11044.
- Dynarski, Susan, and Jonathan Gruber (1997), Can Families Smooth Variable Earnings?, *Brookings Papers on Economic Activity*, 229-284
- Fernández-Villaverde, Jesús, and Dirk Krueger (2002), [Consumption and Saving over the Life Cycle: How Important are Consumer Durables?](#), *Proceedings of the 2002 North American Summer Meetings of the Econometric Society: Macroeconomic Theory*.
- Heathcote, Jonathan, and Fabrizio Perri (2004), [The International Diversification Puzzle is not as Bad as You Think](#), mimeo, New York University
- Heathcote, Jonathan, Kjetil Storesletten and Giovanni Violante (2003), [The Macroeconomic Implications of Rising Wage Inequality in the US](#), mimeo, Georgetown University.
- Heathcote, Jonathan, Kjetil Storesletten and Giovanni Violante (2005), Insurance and Opportunities: The Welfare Implications of Rising Wage Dispersion, mimeo, Georgetown University.
- Kehoe, Timothy, and David Levine (1993), [Debt Constrained Asset Markets](#), *Review of Economic Studies*, 60, 865-888.
- Kehoe, Patrick, and Fabrizio Perri (2002), [International Business Cycles with Endogenous Incomplete Markets](#), *Econometrica*, 70, 907-928.
- Kocherlakota, Narayana (1996), [Implications of Efficient Risk Sharing without Commitment](#), *Review of Economic Studies*, 63, 595-609.
- Krueger, Dirk, and Felix Kubler (2002), [Intergenerational Risk Sharing via Social Security when Financial Markets are Incomplete](#), *American Economic Review*, 92, 407-410.
- Krueger, Dirk, and Felix Kubler (2002), [Pareto Improving Social Security Reform when Financial Markets are Incomplete!?](#), *NBER Working Paper* 9410.
- Krueger, Dirk, and Fabrizio Perri (1999), [Risk Sharing: Private Insurance Markets or Redistributive Taxes](#), *Federal Reserve Bank of Minneapolis Staff Report* 262.
- Krueger, Dirk, and Fabrizio Perri (2002), [Does Income Inequality Lead to Consumption Inequality? Evidence and Theory](#), *NBER Working Paper* 9292
- Krueger, Dirk, and Fabrizio Perri (2004), [On the Welfare Consequences of the Increase in Inequality in the United States](#), in Mark Gertler and Kenneth Rogoff (eds.) *NBER Macroeconomics Annual 2003*, 83-121, The MIT Press, Cambridge, MA.
- Krueger, Dirk, and Fabrizio Perri (2005), Understanding Consumption Smoothing: Evidence from the U.S. Consumer Expenditure Survey Data, forthcoming, *Journal of the European Economic Association*.
- Krueger, Dirk, and Harald Uhlig (2005), Competitive Risk Sharing Contracts with One-Sided Commitment, mimeo,

Goethe University Frankfurt.

Lustig, H, and S van Nieuwerburgh (2004), A Theory of Housing Collateral, Consumption Insurance and Risk Premia, mimeo, UCLA.

Storesletten, Kjetil, Chris Telmer and Amir Yaron (2004), [Consumption and Risk Sharing over the Life Cycle](#), *Journal of Monetary Economics*, 51, 609-633.

EconomicDynamics Interviews Ellen McGrattan on Business Cycle Accounting and Stock Market Valuation

Ellen McGrattan is Monetary Advisor at the Research Department of the Federal Reserve Bank of Minneapolis and Adjunct Professor of Economics at the University of Minneapolis. She has worked on a large number of topics, such as business cycles, equity premiums, optimal debt and solution methods. McGrattan's [RePEc/IDEAS entry](#).

EconomicDynamics: With V.V. Chari and Pat Kehoe, you show that the driving forces in business cycle model are well summarized by efficiency, labor and investment wedges. Using the Great Depression and the 1982 recession in the US, you argue that investment wedges are not relevant. In work with Prescott, you demonstrate that the recent stock market boom can be traced to changes in dividend taxation. Are these two result not contradictory?

Ellen McGrattan: No they are not contradictory.

But to explain that requires some background. In 'Business Cycle Accounting', we propose a methodology---one that is, in my opinion, much better than the SVAR methodology---to isolate promising classes of business cycle theories. There are two parts. The first is to show that a large class of models are equivalent to a prototype growth model with time-varying "wedges" resembling time-varying productivity, labor income tax rates, investment tax rates, and government consumption. The second is the accounting part: measure wedges using data and feed them into the prototype growth model to determine their contributions to aggregate fluctuations. We find that the investment "wedge" (which looks just like a time-varying tax rate on investment) does not contribute significantly to aggregate fluctuations. Therefore, models in which frictions manifest themselves as investment wedges are not promising for studying business cycles. These include those with credit market frictions such as Bernanke and Gertler (1989).

In the paper with Ed entitled 'Taxes, Regulations, and the Value of US and UK corporations', we consider the dramatic secular changes in the value of US and UK corporate equities that occurred between the 1960s and 1990s, when there was little change in corporate capital stocks, after-tax corporate earnings, or corporate net debt. In particular, we ask what growth theory predicts for equities given estimates of taxes and productive capital stocks. There were two innovations that we made that are worth noting. The first innovation was a method to estimate the value of corporate intangible capital, which is not included in measures of productive capital but adds to the value of corporations. Our estimate for intangible corporate capital is large, roughly 2/3 as big as tangible corporate capital. The second innovation was to bring public finance back into finance and relate the large movements in equity values to large movements in the effective tax rate on corporate distributions (e.g., dividends). A key proposition is that a decline in the tax rate on corporate distributions implies a rise in stock values and (if revenues are rebated back) no change in the reproducible cost of capital. This is what we see in the data.

Now let me go back to your question about possible inconsistencies between Chari-Kehoe-McGrattan and McGrattan-Prescott. One reason they are not inconsistent is the key tax rate for MP is the tax rate on corporate distributions. The level of the tax rate on distributions does not enter the dynamic Euler equation, only the growth rate if it is time-varying. If the variation in tax rates quarter by quarter is not large, then the implied investment wedge in big downturns is relatively small and not particularly relevant for cyclical behavior. MP focus on secular change over 40 or 50 years.

ED: Again with V.V. Chari and Pat Kehoe, you have recently worked on sudden stops and how financial crises alone cannot trigger drops in output. In fact, such a crisis would increase output. What critical ingredient is our basic intuition missing here?

EM: The basic intuition of the paper is simple. Using the idea in 'Business Cycle Accounting,' we show an equivalence