FEDERAL RESERVE BANK OF ST. LOUIS

REVIEW

SECOND QUARTER 2016
VOLUME 98 | NUMBER 2

REVIEW

Volume 98 • Number 2

Review

Review is published four times per year by the Research Division of the Federal Reserve Bank of St. Louis. Complimentary print subscriptions are available to U.S. addresses only. Full online access is available to all, free of charge.

Online Access to Current and Past Issues

The current issue and past issues dating back to 1967 may be accessed through our Research Division website: **http://research.stlouisfed.org/publications/review**. All nonproprietary and nonconfidential data and programs for the articles written by Federal Reserve Bank of St. Louis staff and published in *Review* also are available to our readers on this website.

Review articles published before 1967 may be accessed through our digital archive, FRASER: **http://fraser.stlouisfed.org/publication/?pid=820**.

Review is indexed in Fed in Print, the catalog of Federal Reserve publications (**http://www.fedinprint.org/**), and in IDEAS/RePEc, the free online bibliography hosted by the Research Division (**http://ideas.repec.org/**).

Authorship and Disclaimer

The majority of research published in *Review* is authored by economists on staff at the Federal Reserve Bank of St. Louis. Visiting scholars and others affiliated with the St. Louis Fed or the Federal Reserve System occasionally provide content as well. *Review* does not accept unsolicited manuscripts for publication.

The views expressed in *Review* are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

Subscriptions and Alerts

Single-copy subscriptions (U.S. addresses only) are available free of charge. Subscribe here: **https://research.stlouisfed.org/publications/review/subscribe/**.

Our monthly email newsletter keeps you informed when new issues of *Review*, *Economic Synopses*, *The Regional Economist*, and other publications become available; it also alerts you to new or enhanced data and information services provided by the St. Louis Fed. Subscribe to the newsletter here: **http://research.stlouisfed.org/newsletter-subscribe.html**.

Economic Data

General economic data can be obtained through FRED® (Federal Reserve Economic Data), our free database with about 300,000 national, international, and regional data series, including data for our own Eighth Federal Reserve District. You may access FRED through our website: **http://research.stlouisfed.org/fred2**.

ISSN 0014-9187

In Memoriam: Robert H. Rasche

Robert Harold Rasche
1941-2016
Former Director of Research
Federal Reserve Bank of St. Louis

Bob Rasche was a champion of the St. Louis Fed. He devoted his energy and imagination to cultivating an environment of public service and rigorous economic research. Bob died June 2, 2016. He is survived by his wife of 52 years, Dorothy Anita Bensen; his children, Jeanette [Bart] Little and Karl [Tuhina] Rasche; and his brother, Richard A. [Jamie] Rasche.

Bob was born in New Haven, Connecticut. He received his BA in economics and mathematics from Yale University and his PhD in economics from the University of Michigan. Bob had a distinguished career in academia. He published many research papers in top scholarly journals and was well respected by his peers in the economics profession. The IDEAS/RePEc bibliographic database provides some details: Bob's publishing rank is in the top 3% of economists. His most cited paper comes from the St. Louis Fed *Review*, "Market Anticipations of Monetary Policy Actions," with co-authors William Poole and Daniel Thornton. And his most downloaded paper comes from the Carnegie-Rochester Conference Series on Public Policy: "Energy Price Shocks, Aggregate Supply and Monetary Policy: The Theory and the International Evidence."

In January 1999, Bob became the St. Louis Fed's director of research under Bank president William Poole. Bob was already well connected to the St. Louis Fed before he accepted the role: In August 1971, Bob took a leave of absence from the University of Pennsylvania and worked as a visiting scholar at the St. Louis Fed. In August 1976, Bob took a leave of absence from Michigan State University and returned to the St. Louis Fed, again, as a visiting scholar. Bob paid close attention to the research being conducted there and was an avid reader and user of the Bank's research publications and data services.

His work within the Federal Reserve and the context of central banking in general prepared him well for his tenure at the St. Louis Fed. He was an MIT research associate on the Federal Reserve Board—MIT Monetary Research Project; a visiting scholar at the San Francisco Fed; a visiting scholar at the Bank of Japan's Institute for Monetary and Economic Studies; and a prolific author of Shadow Open Market Committee position papers. Bob amassed a decades-deep understanding of the research environment and the overall mission of the Federal Reserve.

At the St. Louis Fed, Bob established objectives for conducting high-quality economic research, disseminating that research in scholarly journals and Bank publications, and bolstering the reach and influence of the Bank throughout the Federal Reserve System and academia.

Bob was also one of the most vocal and persistent advocates of delivering accurate data and information to the public. He had closely followed the St. Louis Fed's historical trajectory of collecting, analyzing, and sharing economic data. He understood and respected that tradition and viewed it as a vital public service. One of his legacies is that he expanded and enhanced FRED, the Bank's prime data service, at a key moment in its history: He made a compelling case to the Bank's president and its senior leaders that the St. Louis Fed's mission of public service should continue and thrive. And it did thrive. During Bob's tenure, the data and information services grew in size, scope, and recognition.

Many of the current projects in the St. Louis Fed's Research Division began with Bob Rasche, and existing lines of business were expanded and strengthened. Bob led and inspired his staff by his own example, with authentic passion and engagement both inside and outside the Bank:

Bob was a family archivist, compiling information on his own family history, which he organized into a multigenerational collection after he retired. He also envisioned and nurtured a realm of St. Louis Fed data services, including a database of historical "vintage" economic data (ALFRED) and a digital research archive of Federal Reserve and national economic history (FRASER).

Bob and his wife, Dottie, had many opportunities for domestic and international travel, including a number of cruises to Norway, the Baltic, and Alaska. With that expansive vision, he conceived of and brought about GeoFRED, an innovative extension of the FRED database that provides geographic maps of data.

As a teenager, Bob delivered the *New Haven Register* for six years. His work ethic and vigor persisted: He would arrive at the office early in the morning, passionately describing the ideas that had occurred to him overnight, and a new project was born. He also produced economic research his entire career. Current St. Louis Fed president, James Bullard, remarked on the occasion of Bob's retirement that Bob had published five academic papers in one year alone.

Bob was also an avid carpenter. He was curious, creative, and precise in his expectations and measurements of success, and he expected the finished products to be polished and true. Bob crafted valuable and lasting service to the public, the economics profession, and the Federal Reserve System. His contributions stand out and inspire those who continue this work. ■

Permazero

James Bullard

The financial crisis of 2007-09 and its aftermath turned monetary economics and policymaking on its head and called into question many of the conventional views held before the crisis. One of the most popular and enduring views in all of monetary economics since the 1970s, and indeed since the 1940s, has been that a nominal interest rate peg is poor monetary policy and that attempts to pursue such a policy would lead to ruin. Yet, post-crisis U.S. monetary policy could be interpreted as exactly that—an interest rate peg—and an extreme one at that, since the policy rate has remained near zero for nearly seven years. The author summarizes some recent academic work on the idea of a stable interest rate peg and what its implications may be for current monetary policy choices. He argues that a stable interest rate peg is a realistic theoretical possibility; that it has some mild empirical support based on a cursory look at the data; and that, should we find ourselves in a persistent state of low nominal interest rates and low inflation, some of our fundamental assumptions about how U.S. monetary policy works may have to be altered. (JEL E31, E52, E58)

Federal Reserve Bank of St. Louis *Review*, Second Quarter 2016, *98*(2), pp. 81-92.
http://dx.doi.org/10.20955/r.2016.81-92

MY CURRENT POLICY RECOMMENDATIONS

Let me begin by describing briefly my current monetary policy recommendations. Those of you who have followed my commentary during 2015 know that I have been an advocate of ending the Federal Open Market Committee's (FOMC's) near-zero nominal interest rate policy. My case has been straightforward. Essentially, I have argued that while the Committee's goals have been met, the Committee's policy settings remain as extreme as they have been at any time since the recession ended in 2009.

With respect to these goals, the current unemployment rate of 5 percent is statistically indistinguishable from the Committee's view of the equilibrium long-run rate of unemployment. In addition, the current year-over-year inflation rate, while low, reflects an outsized oil

James Bullard is president and CEO of the Federal Reserve Bank of St. Louis. This article is based on his keynote address at the Cato Institute's 33rd Annual Monetary Conference, "Rethinking Monetary Policy," Washington, D.C., November 12, 2015. This article originally appeared in the *Cato Journal*, Vol. 36, No. 2 (2016). The author thanks his staff for helpful comments.

price shock that occurred during 2014. A measure that tries to control for this effect—the Dallas Fed's trimmed mean inflation rate, measured year over year—is currently running at 1.7 percent, just 30 basis points below the FOMC's inflation target of 2 percent. By these measures, the Committee's goals have been met.

On the other hand, the Committee's policy settings remain far from normal. The policy rate remains near zero, and the balance sheet is very large relative to its pre-crisis levels. In the past, the Committee has acted to normalize policy well before goals have been completely met.

A simple and prudent approach to current policy is to move the policy settings closer to normal levels now that the goals of policy have been attained. There is no reason to continue to experiment with extreme policy settings.

Implicit in my argument is a desire to return to the 1984-2007 U.S. macroeconomic equilibrium, which involved relatively good monetary policy, relatively long economic expansions, and a higher nominal interest rate than we have today. Part of the nature of that equilibrium was a monetary policy that was relatively well understood by both financial market participants and monetary policymakers. We gained much experience with the equilibrium over this time period, and we think we know how it works, in part because it has been studied extensively from both a theoretical and empirical perspective.

RETHINKING MONETARY POLICY

Nevertheless, as the topic of this conference is "Rethinking Monetary Policy," I plan to devote the bulk of my remarks not to the return to the standard macroeconomic equilibrium that I recommend, but to the possibility that such a return is not achieved, despite the Committee's best efforts to engineer such an outcome for the U.S. economy.

We have, after all, been at the zero lower bound in the United States for seven years. In addition, the FOMC has repeatedly stressed that any policy rate increase in coming quarters and years will likely be more gradual than either the 1994 cycle or the 2004-06 cycle. In short, the FOMC is already committed to a very low nominal interest rate environment over the forecast horizon of two to three years. Perhaps short-term nominal rates will simply be low during this period, or perhaps the economy will encounter a negative shock that will propel policy back toward the zero lower bound.

Our experience is not unique. In Japan, the policy rate has not been higher than 50 basis points for two decades, and in the euro zone, the policy rate looks set to remain near zero at least through September 2016. The thrust of this talk is to suppose, for the sake of argument, that the zero interest rate policy (ZIRP) or near-ZIRP remains a persistent feature of the U.S. economy. How should we think about monetary stabilization policy in such an environment? What sorts of considerations should be paramount? Should we expect slow growth? Will we continue to have low inflation, or will inflation rise? Would we be at more risk of financial asset price volatility? What types of concrete policy decisions could be made to cope with such an environment? Would it require a rethinking of U.S. monetary policy?

I will provide tentative answers to all of these questions. But first, I want to argue that it may indeed be possible to converge to an equilibrium at the zero lower bound, and that this situation has some surprising consequences. Chief among these consequences is that the policy itself may put downward pressure on inflation in the medium and long term, rather than upward pressure as conventionally thought. This is a simple consequence of the Fisher equation having to hold in concert with monetary neutrality. I will now turn to developing this point.

PERMAZERO

Most analyses of U.S. monetary policy since the crisis of 2007-09 have suggested that the ZIRP in the United States is a temporary affair, one that was part of an important set of policy actions designed to mitigate a particularly large shock to the U.S. economy. But how temporary is it?

We have been at the zero lower bound for nearly seven years. This is well beyond an ordinary business cycle time. Normally, we would think of a shock hitting the economy, with the effects of that shock largely wearing off well within a seven-year time span. What are the consequences of spending such a long time with the policy rate at one value? Arguably, it is an interest rate peg.

In the 1970s and 1980s, the typical reply to this question was that an interest rate peg was poor policy. Trying to keep the nominal policy rate unnaturally low for too long a period would ultimately be inflationary, and indeed, this was widely viewed as a large part of the problem leading to global inflation during this era.[1] Indeed, during the past six years I have warned, along with many others, that the Committee's ZIRP has put the U.S. economy at considerable risk of future inflation. In fact, my monetarist background urges me to continue to make this warning right now!

In any case, after seven years, one might want to consider other models. One important possibility is that the 1970s were an era when U.S. monetary policy was not very credible with respect to fighting inflation, whereas the 2000s were an era when U.S. monetary policy had already earned a lot of credibility for keeping inflation low and stable. One way to interpret this is to say that market expectations of future inflation today move to stay in line with the FOMC's desired policy rate instead of becoming "unanchored" as they did in the 1970s. In particular, this would mean that a low nominal interest rate peg, far from being a harbinger of runaway inflation, would instead dictate medium- and longer-run low inflation outcomes.

This theme is sometimes labeled "neo-Fisherian" because it emphasizes that the Fisher equation holds in virtually all modern macroeconomic models. The Fisher equation states that the nominal interest rate can be decomposed into a real interest rate component and an expected inflation component. If we view the real interest rate as determined by supply and demand conditions in the private sector, then a permanent nominal interest rate peg would also pin down the long-run rate of inflation. The Fisher equation implies, among other things, that the monetary authority cannot choose the long-run value of the nominal policy interest rate separately from the long-run value of inflation.

This Fisher effect is well known and is not likely to be disputed in macroeconomic circles. However, how long before this Fisher effect sets in? Over what time period can the monetary authority maintain an interest rate peg before the peg itself begins to pull inflation expectations in a direction consistent with the peg? Is seven years a sufficient length of time? How about 20 years, as in Japan?

COCHRANE (2016)

A recent paper by Cochrane (2016) provides an interesting analysis of this issue in the context of the most canonical of modern macroeconomic models, the linearized three-equation New Keynesian model.[2] I will not provide any details of the model here, but for those who are unfamiliar with it, I will briefly describe its essential ingredients. The key friction in the model is that prices are sticky, meaning that they do not adjust immediately in response to supply and demand conditions. Households and firms solve optimization problems taking the friction as given. The policymaker controls a one-period nominal interest rate and through this channel can have temporary effects on real output and inflation. The Fisher equation holds at all times. The model can be described by three simple equations that depend on expectations of future real output, future inflation, and future monetary policy. The spirit of Cochrane's analysis is to suggest that neo-Fisherian effects are part of even the most ordinary of macroeconomic models used to inform current monetary policy.

Cochrane (2016) uses a solution technique for the model due to Werning (2012). We can think of the economy as continuing from the distant past to the distant future. The policymaker chooses the short-term nominal interest rate sequence, and, given this sequence, the model traces out what would happen to the real output gap (x) and inflation (π).

I use Cochrane's model to trace out the effects on the economy of the following thought experiment. Suppose the economy begins with the nominal interest rate equal to 2 percent, a real interest rate equal to 0 percent (for convenience), and an inflation rate equal to 2 percent. The Fisher equation holds, as it must, so that in the long run the policy rate will equal the inflation rate in this example. The policymaker then lowers the policy rate by 200 basis points to zero and leaves it there for a considerable time.

Figure 1 illustrates the effect of such a policy experiment in Cochrane's (2016) model. The green triangles show the policy rate, which begins at 200 basis points, but is lowered to zero at date 0. If the policy move is anticipated, as many actual policy moves are, then the effects on inflation are described by the red squares, and the effects on the real output gap are given by the black circles. If the policy change is completely unanticipated, then the effects on inflation are given by the orange squares, and the effects on the real output gap are given by the blue circles. In the case of a "surprise" policy move, nothing happens until the date of the move, whereupon the inflation and real output gap variables jump to the path they would have been on had the policy change been known in advance. For our purposes here, it does not matter that much if we focus on an anticipated or an unanticipated policy change.

Instead, I want to focus on the right-hand side of this picture, after the policy move has occurred. The policymaker has lowered the policy rate to zero, and in response, the real output gap has increased.[3] This is one way to gauge the real effects of monetary policy according to the model: A pure change in the policy rate, with no other shocks occurring, would temporar-

Figure 1

The Policy Rate Falls 200 Basis Points

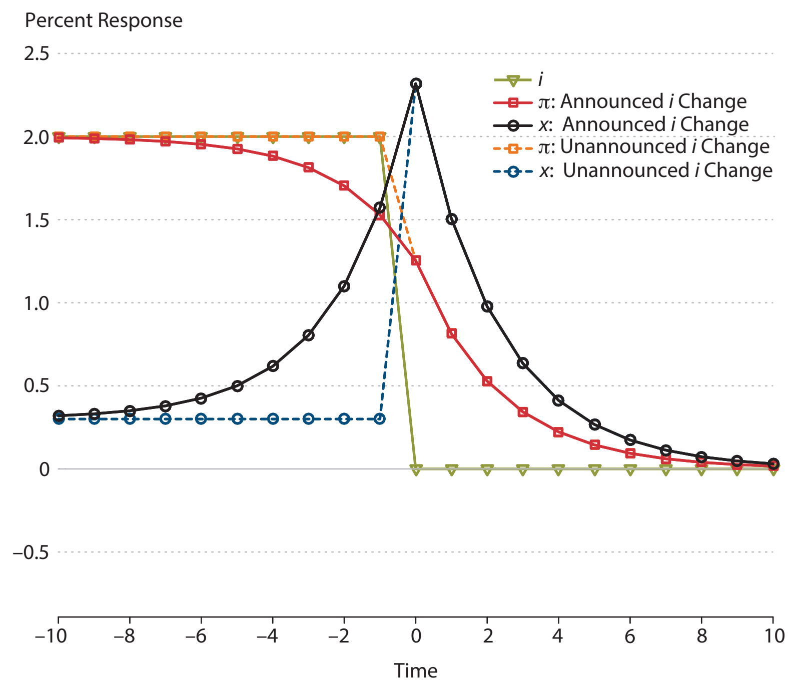

NOTE: i = policy rate; π = inflation, x = output gap.
SOURCE: Adapted from Cochrane (2016).

ily increase output. This is what the model is designed to do, and if we added more shocks to the model, the policymaker could use this power appropriately to smooth real output over time. Smoother output would be preferred to more volatile output by the households in the model, and thus the model provides a theory of monetary stabilization policy.

But now let us look at inflation in response to the policy change. It falls in response to the policy change, very little at first, but more substantially as the ZIRP continues. After about 2.5 years (10 quarters), at the far right of Figure 1, the transitory effects of the policy change have nearly completely died out. The real output gap is zero, the policy rate remains at zero, and the inflation rate has fallen to zero. This can be interpreted as a neo-Fisherian result: The policy rate is lowered, and after some transitory dynamics, the inflation rate falls to be consistent with the new interest rate peg.

It is clear from Figure 1 that, should the policymaker simply elect to keep the nominal interest rate at zero for a much longer time, nothing further would happen in this economy. The black, red, and green lines would simply remain at zero.

Cochrane's (2016) analysis, as I have translated it into Figure 1, yields a very different interpretation of current events compared with conventional wisdom. Conventional descrip-

Figure 2

A Gradual Policy Rate Increase

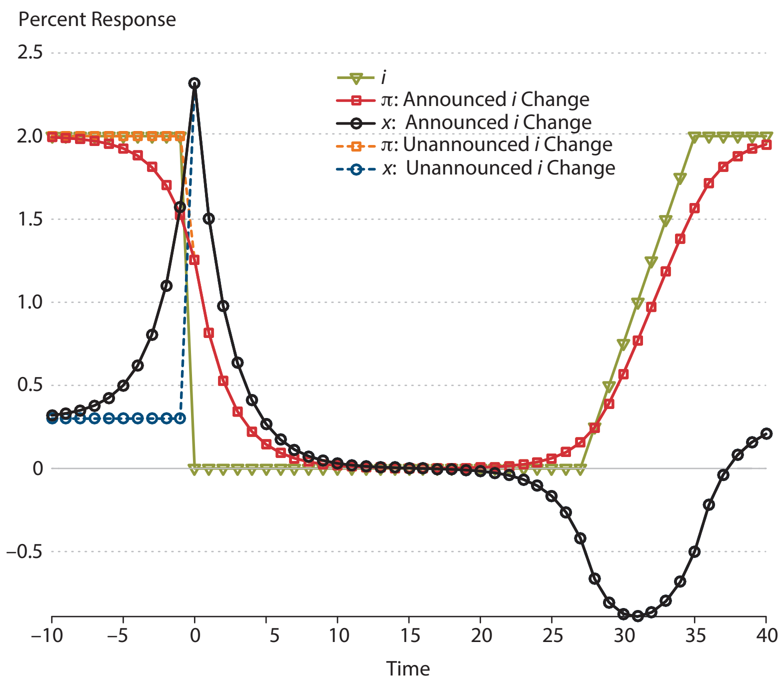

NOTE: i = policy rate; π = inflation, x = output gap.
SOURCE: Adapted from Cochrane (2016).

tions of current monetary policy, including my own description earlier in this very speech, suggest that the Committee's ZIRP is putting upward pressure on inflation, perhaps dangerously so. Figure 1 suggests otherwise.

What's going on? The model does have a Phillips curve in that today's inflation rate does depend in part on today's real output gap. When the policy rate is lowered, the output gap is higher than it otherwise would have been, and this does put upward pressure on inflation in the model.[4] However, the model also has a Fisher relation, which means that as the real output gap returns to normal (that is, monetary neutrality asserts itself), the inflation rate will have to fall to be consistent with the new level of the nominal interest rate. Another aspect is that the policymaker is viewed as choosing the interest rate sequence, and inflation follows as dictated by the Fisher equation. The policymaker cannot set the nominal interest rate and the inflation target in an inconsistent way.

A few of you may be aware of a closely related analysis by Benhabib, Schmitt-Grohé, and Uribe (2001) that I have championed in discussing dimensions of monetary policy since 2007-09.[5] In that analysis the Fisher relation also plays a prominent role, but the analysis is nonlinear and global. Benhabib, Schmitt-Grohé, and Uribe (2001) find two steady states, one of which is associated with a low nominal interest rate and inflation below target. Arguments in this context then center around which of the two steady states is the stable one in a reasonable expectation dynamic ("learning"). Often the argument is that the traditional steady state is the stable one and therefore the one worthy of the most attention from policymakers.[6] The Cochrane (2016) analysis is of a linear system, and, consequently, ideas about "getting stuck at the wrong steady state" are not nearly as clear. Rational expectations prevail at all times.[7]

To illustrate that policymakers can reverse their actions in the Cochrane (2016) model, Figure 2 illustrates an alternative policy experiment. This experiment is almost the same as the one described in Figure 1, except that the policymaker chooses the nominal policy rate sequence to remain at zero for seven years before gradually raising the policy rate back to 2 percent.

The left-hand side of Figure 2 simply repeats what is in Figure 1. The middle portion of Figure 2 shows how the case where the policy rate remains near zero simply keeps the inflation rate low and the output gap steady as the effects of the first policy move wear off. The gradual policy rate increase is shown in the right-hand portion of Figure 2 by the green triangles. This policy move is portrayed as being anticipated here, so inflation and the output gap begin to react before the actual date of liftoff. The rising rate environment puts downward pressure on the output gap, reversing the effects of the previous policy rate move. As before, inflation moves in tandem with the policy rate as the Fisher equation asserts itself.

Is this what will actually happen in the U.S. economy? Definitely not, since we are looking here at pure policy effects with no other shocks added to the model. At best, Figures 1 and 2 can illustrate the directions that monetary policy can be expected to push in this particular model; but a more realistic analysis would include additional shocks, and monetary policy would have to react appropriately to those changes in macroeconomic conditions. Still, the key point is that this canonical model has a clear interpretation in neo-Fisherian terms, and that this interpretation is hardly surprising, since the Fisher equation is built into the model.

I have spent a lot of my time with these particular figures because I think they are interesting and can communicate to a wide audience in the monetary policymaking community. But I do want to stress that the New Keynesian model is just one model in a sea of possibilities. In addition, it is a model that was designed to describe the relatively good monetary policy in the United States from 1984-2007, without features that turned out to be quite important during the 2007-09 crisis and its aftermath. While I do not have time to emphasize other more novel work here, let me just say that there is important recent work in monetary theory and policy that has tried to explain very low real rates of return on safe assets along with the implications for monetary policy. Andolfatto and Williamson (2015), for instance, think of all consolidated government debt as having value in conducting transactions. Their model has a liquidity premium on government debt under some circumstances and offers novel interpretations of current policy dilemmas. Caballero and Farhi (2015) similarly study safe asset shortages and suggest important ways that our understanding of the effectiveness of various policies at the zero lower bound would be affected. These are just some examples of interesting work going on outside the relatively narrow New Keynesian framework to try to come to terms with the reality of the post-crisis macroeconomy.

Figure 3

G-7 Countries' Aggregated Inflation and Policy Rates

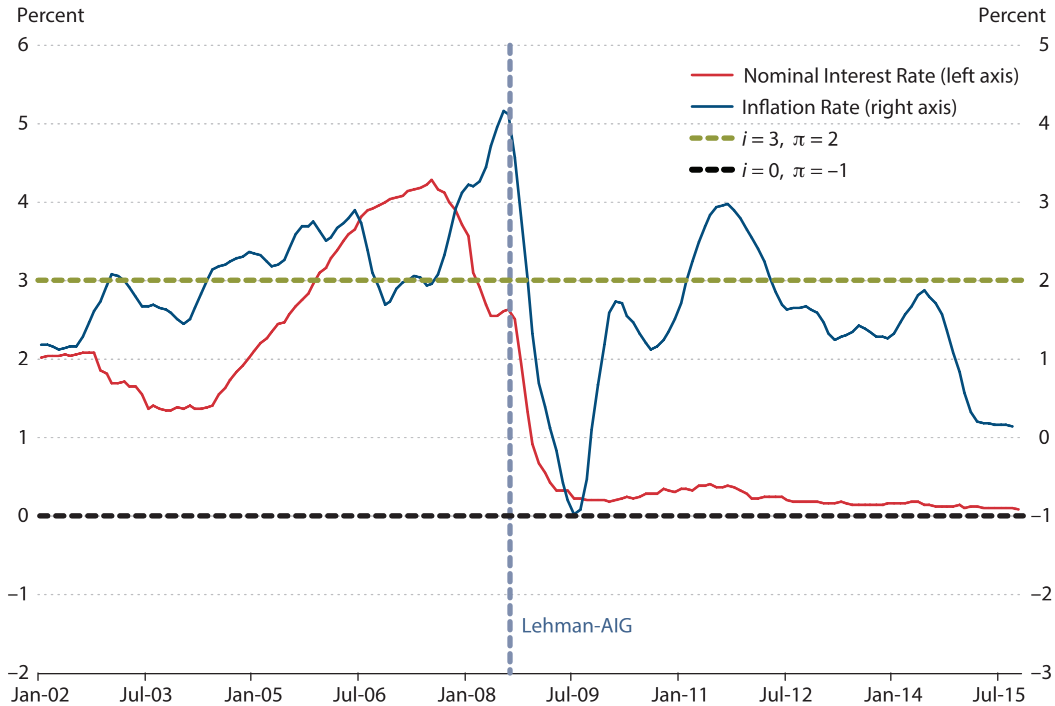

SOURCE: Organization for Economic Co-operation and Development's *Main Economic Indicators* and author's calculations. Last observation: September 2015.

EMPIRICAL EVIDENCE

Figures 1 and 2 suggest that low nominal interest rates and low inflation may go hand in hand, at least over relatively long horizons in which the policy rate is kept at a constant level. Over shorter horizons with more policy moves and more shocks, the correlation may not be very high. Policy rates have generally been very low, near zero, continuously in the G-7 economies since the 2007-09 period. Consequently, we may be able to look at the data since 2009 to see to what extent neo-Fisherian effects are exerting themselves in the G-7.[8]

To get at this issue in just one picture, Figure 3 shows the centered five-quarter moving average of the G-7 headline inflation rate and the average, GDP-weighted, G-7 nominal policy rate since 2002. In Figure 3, the inflation rate is the blue line on the right-hand scale, and the GDP-weighted nominal policy rate is the red line on the left-hand scale. The horizontal green line is an inflation rate of 2 percent, and the horizontal black line is an inflation rate of –1 percent. The vertical line in the middle of Figure 3 marks the Lehman-AIG event. On the left side of Figure 3, interest rates and inflation arguably behaved according to traditional interpretations of New Keynesian theory. On the right half of Figure 3, the nominal policy rate falls to near zero and remains there. Inflation initially falls across the G-7, but then impressively returns close to target. In fact, inflation was above target as of the beginning of 2012, about 2.5 years after the end of the recession in the United States. Since then, however, policy rates have remained near zero and inflation has drifted down, to the point where G-7 inflation is around zero today.

Conventional wisdom would have suggested that the zero policy rates in the G-7 were putting upward pressure on inflation during the nearly four years since January 2012, but instead, inflation fell. This could be viewed as consistent with neo-Fisherian effects asserting themselves. Of course, we have to be cautious about carrying such an explanation too far. There have been many other shocks during the past four years, notably a very large oil price shock beginning in the summer of 2014.

CONSEQUENCES

Let us suppose for the sake of argument that the G-7 economies will spend still more time at or near the zero lower bound. This would occur because either liftoff does not materialize in most or all countries or because additional negative shocks drive those countries that do raise their policy rates back to the zero lower bound. Prudent policymaking suggests that we should at least entertain this as a realistic possibility for the path of G-7 monetary policy in the coming years. What are the consequences of remaining in such a state for a long period of time?[9]

I can think of six consequences, based on the discussion in the earlier part of this speech:

First, consider the near-zero policy rate path illustrated on the right-hand side of Figure 1. In this situation, promising to keep the nominal interest rate sequence at the zero lower bound simply reinforces the equilibrium and does not provide accommodation as in the traditional New Keynesian equilibrium. Nothing happens in response to such promises. Policymakers would have to come to grips with such a situation.

Second, in such a situation, inflation remains persistently below the stated inflation target. The near-zero policy rate is not putting upward pressure on inflation, but is instead through the Fisher equation dictating a rate of inflation lower than the original target. It could be that policymakers do not intend to return to the original equilibrium—that is, they may intend to remain with the near-zero policy rate. In that case, policymakers may wish to lower the inflation target to remain more consistent with the actual inflation outcomes.

Third, longer-run economic growth would still be driven by human capital accumulation and technological progress, as always, but without the accompanying stabilization policy as conventionally practiced from 1984-2007. In principle, the economy would still be expected to grow at a pace dictated by fundamentals.[10]

Fourth, the celebrated Friedman rule would arguably be achieved so that household and business cash needs are satiated. In many monetary models this is a desirable state of affairs.

Fifth, the risk of asset price fluctuations may be high. In the New Keynesian model, the near-ZIRP with little or no response to incoming shocks is associated with equilibrium indeterminacy. This means there are many possible equilibria, all of which are consistent with rational expectations and market clearing. In a nutshell, a lot of things can happen. Many of the possible equilibria are exceptionally volatile. One could interpret this theoretical situation as consistent with the idea that excessive asset price volatility is a risk.

Sixth, and finally, the limits on operating monetary policy through ordinary short-term nominal interest rate adjustment in this situation would surely continue to fire a search for alternative ways to conduct monetary stabilization policy. The favored approach during the past five years within the G-7 economies has been quantitative easing, and there would surely be pressure to use this or related tools.[11]

CONCLUSION

During 2015, I was an advocate of beginning to normalize the policy rate in the United States. My arguments have focused on the idea that the U.S. economy is quite close to normal today based on an unemployment rate of 5 percent, which is essentially at the Committee's estimate of the long-run rate, and inflation net of the 2014 oil price shock only slightly below the Committee's target. The Committee's policy settings, in contrast, remain as extreme as they have ever been since the 2007-09 crisis. The policy rate remains near zero, and the Fed's balance sheet is more than $3.5 trillion larger than it was before the crisis. Prudence alone suggests that, since the goals of policy have been met, we should be edging the policy rate and the balance sheet back toward more normal settings.

Implicit in my argument has been a yearning to return to the monetary equilibrium of 1984-2007, which is one around which a great deal of theory and empirical work has been done. We would be returning to a world in which monetary policy is better understood, the effects of policies are more closely calibrated, and private sector expectations can move and adapt to ordinary adjustments of the policy rate.

My current policy views have not changed. But in the spirit of the conference, I have tried to contribute to the topic of "Rethinking Monetary Policy" by focusing on a situation where the nominal policy rate and the inflation rate remain low, either because liftoff does not materialize or because future negative shocks to the economy force a return to the zero interest rate policy. I have illustrated by reference to relatively new research how such a situation could become permanent. In addition, I have suggested several consequences of remaining at such an equilibrium over the long term. It is my hope that my characterizations here will spur further thinking and research on these important topics. ■

NOTES

[1] See, for instance, Sargent and Wallace (1975) for an argument that an interest rate peg is associated with price level indeterminacy.

[2] See Woodford (2003) and Galí (2015).

[3] The long-run real output gap in this model is not zero unless the long-run inflation rate is zero, so the initial real output gap on the left-hand side of this picture is somewhat positive. This is not material to the argument here, but has been discussed extensively in the New Keynesian literature.

[4] The inflation decline is mitigated by the increase in real activity.

[5] See Bullard (2010) and Bullard (2015).

[6] See Eusepi (2007), Evans (2013), and Benhabib, Evans, and Honkapohja (2014).

[7] García-Schmidt and Woodford (2015) delve into this question and, in particular, consider departures from rational expectations.

[8] For state-of-the-art empirical analysis of the issues discussed here, see Aruoba and Schorfheide (2015) and Aruoba, Cuba-Borda, and Schorfheide (2014).

[9] Cochrane (2014) addresses how U.S. monetary policy might operate in a zero policy rate and large balance sheet environment.

[10] Endogenous growth theories that mix long-run growth prospects with monetary policy practice are rare and of dubious empirical validity.

[11] For some recent arguments concerning the future of monetary policy in a low interest rate environment, see Haldane (2015). For a theoretical analysis of quantitative easing at the zero lower bound, see Boel and Waller (2015).

REFERENCES

Andolfatto, David and Williamson, Stephen. "Scarcity of Safe Assets, Inflation, and the Policy Trap." Federal Reserve Bank of St. Louis Working Paper No. 2015-002A, January 2015; https://research.stlouisfed.org/wp/2015/2015-002.pdf.

Aruoba, S. Borağan; Cuba-Borda, Pablo and Schorfheide, Frank. "Macroeconomic Dynamics Near the ZLB: A Tale of Two Countries." Penn Institute for Economic Research Working Paper No. 14-035, June 19, 2014.

Aruoba, S. Borağan and Schorfheide, Frank. "Inflation During and After the Zero Lower Bound." Unpublished manuscript, University of Pennsylvania, 2015.

Benhabib, Jess; Evans, George W. and Honkapohja, Seppo. "Liquidity Traps and Expectation Dynamics: Fiscal Stimulus or Fiscal Austerity?" *Journal of Economic Dynamics and Control*, August 2014, *45*, pp. 220-38; http://dx.doi.org/10.1016/j.jedc.2014.05.021.

Benhabib, Jess; Schmitt-Grohé, Stephanie and Uribe, Martin. "The Perils of Taylor Rules." *Journal of Economic Theory*, 2001, *96*(1-2), pp. 40–69; http://dx.doi.org/10.1006/jeth.1999.2585.

Boel, Paolo and Waller, Christopher J. "On the Theoretical Efficacy of Quantitative Easing at the Zero Lower Bound." Federal Reserve Bank of St. Louis Working Paper No. 2015–027A, September 2015; https://research.stlouisfed.org/wp/2015/2015-027.pdf.

Bullard, James. "Seven Faces of 'The Peril.'" Federal Reserve Bank of St. Louis *Review*, September/October 2010, *92*(5), pp. 339-52; https://research.stlouisfed.org/publications/review/10/09/Bullard.pdf.

Bullard, James. "Neo-Fisherianism." Remarks delivered at the "Expectations in Dynamic Macroeconomic Models" conference, Eugene, Oregon, August 13, 2015; https://www.stlouisfed.org/~/media/Files/PDFs/Bullard/remarks/Bullard-Expectations-in%20Dynamic-Macroeconomic-Models-08-13-2015.pdf.

Caballero, Ricardo J. and Farhi, Emmanuel. "The Safety Trap." Unpublished manuscript, 2015, Harvard University.

Cochrane, John H. (2014) "Monetary Policy with Interest on Reserves." *Journal of Economic Dynamics and Control*, December 2014, *49*, pp. 74-108; http://dx.doi.org/10.1016/j.jedc.2014.09.003.

Cochrane, John H. "Do Higher Interest Rates Raise or Lower Inflation?" Unpublished manuscript, Chicago Booth School of Business, 2015, updated February 10, 2016; http://faculty.chicagobooth.edu/john.cochrane/research/papers/fisher.pdf.

Eusepi, Stefano. "Learnability and Monetary Policy: A Global Perspective." *Journal of Monetary Economics*, May 2007, *54*(4), pp. 1115-31; http://dx.doi.org/10.1016/j.jmoneco.2006.02.003.

Evans, George W. "The Stagnation Regime of the New Keynesian Model and Recent U.S. Policy," in Thomas J. Sargent and Jouko Vilmunen (eds.) *Macroeconomics at the Service of Public Policy*. Chap. 3. Oxford, UK: Oxford University Press, 2013; http://dx.doi.org/10.1093/acprof:oso/9780199666126.003.0004.

Galí, Jordi. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Second Edition. Princeton, NJ: Princeton University Press, 2015.

García-Schmidt, Mariana and Woodford, Michael. "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis." NBER Working Paper No. 21614, National Bureau of Economic Research, October 2015.

Haldane, Andrew G. "Stuck." Remarks delivered at the Open University, London, UK, June 30, 2015.

Sargent, Thomas J. and Wallace, Neil. "'Rational' Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule." *Journal of Political Economy*, April 1975, *83*(2), pp. 241-54.

Werning, Iván. "'Managing' a Liquidity Trap: Monetary and Fiscal Policy." Unpublished manuscript, MIT, March 2012.

Woodford, Michael. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press, 2003.

Secular Stagnation and Monetary Policy

Lawrence H. Summers

This article is based on the author's Homer Jones Memorial Lecture delivered at the Federal Reserve Bank of St. Louis, April 6, 2016. (JEL E52, E61, H50)

Federal Reserve Bank of St. Louis *Review*, Second Quarter 2016, *98*(2), pp. 93-110.
http://dx.doi.org/10.20955/r.2016.93-110

I have been engaged in thinking, writing, provoking, and analyzing around the issue of secular stagnation: the issue of protracted sluggish growth, why it seems to be our experience, and what should be done about it.[1] This paper summarizes my current thinking on those topics and reflects on the important limits monetary policy experiences in dealing with secular stagnation.

The recent sluggish performance of the economy, both in the United States and particularly abroad, motivates this discussion of secular stagnation. Figure 1 depicts successive estimates of potential GDP and the behavior of actual GDP for the United States. While the Congressional Budget Office (CBO) finds a much lower output gap than it did in 2009, that is wholly the result of downward revisions to its estimate of potential GDP (see Figure 2).

Relative to what the CBO thought potential was in 2007, the gap has never been greater. The situation is more negative in Europe, shown in Figure 3, where actual GDP has performed worst and the magnitude of the downward revision has been substantially greater.

The contrast between what happened after the fall of 2008 and what happened after the fall of 1929 was a notable macro-economic achievement. Collapse and catastrophe happened after the fall of 1929 with the Great Depression. There was nothing like the Great Depression after 2008, in part because of aggressive actions of the Federal Reserve and the Treasury. That was a bipartisan achievement. Key steps such as the Troubled Asset Relief Program (TARP) were initiated during the Bush administration. Further steps such as the American Recovery and Reinvestment Act took place during the Obama administration.

But although the recent policy response to the financial crisis far surpassed that which followed the 1929 crash, twelve years after the acute shock the economy is likely to be in a

Lawrence H. Summers is president emeritus and a distinguished professor at Harvard University, and he directs the university's Mossavar-Rahmani Center for Business and Government.

Figure 1

Downward Revision in Potential GDP, United States

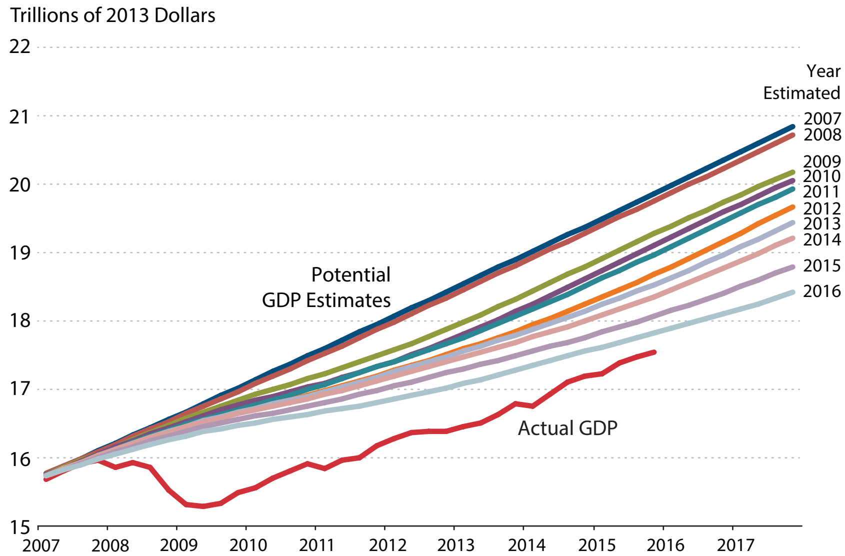

SOURCE: CBO Budget and Economic Outlooks 2007-15; Bureau of Economic Analysis.

Figure 2

Growth at a Permanently Lower Plateau?

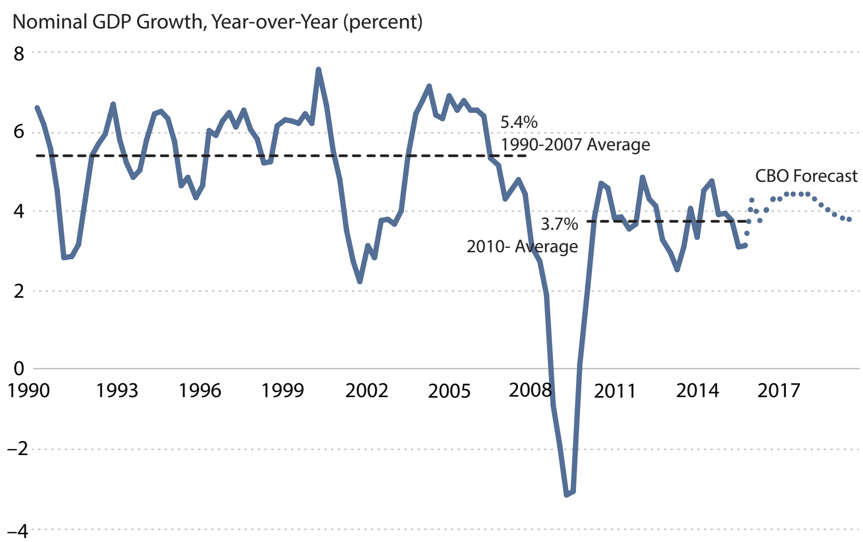

SOURCE: Bureau of Economic Analysis; CBO.

Figure 3

Downward Revision in Potential GDP, Eurozone

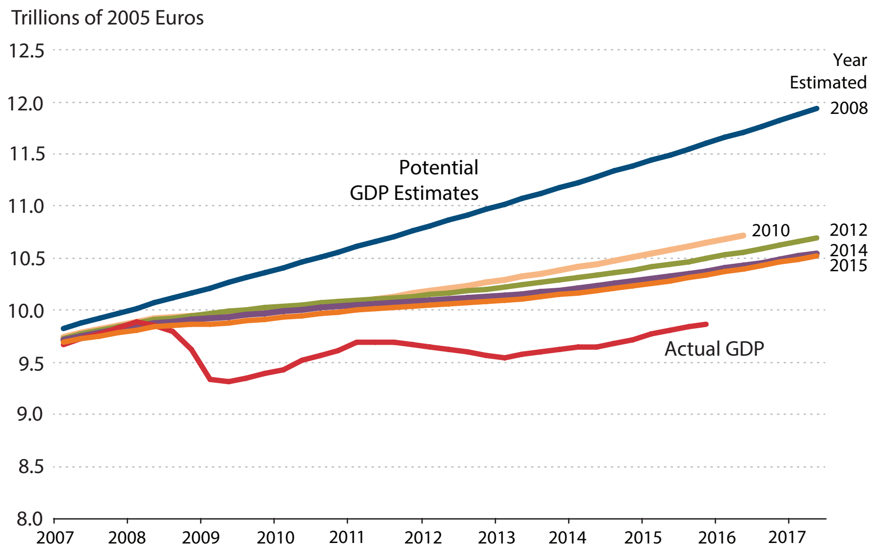

SOURCE: IMF World Economic Outlook Database; Bloomberg.

Figure 4

Great Recession Very Damaging

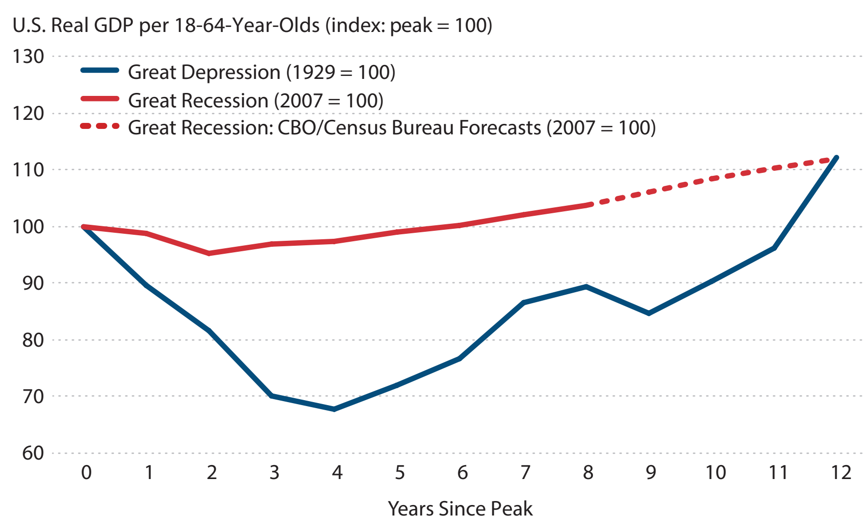

SOURCE: Bureau of Economic Analysis; NBER; CBO; Census.

similar place. As shown in Figure 4, the current forecast is that GDP per working age American will grow by the same amount from 2007-2019 as it did from 1929-1941. The 1930s are thought of as a lost decade, and the 2010s are an equivalently lost decade. While the economy imploded much more dramatically in the 1930s, it also recovered more dramatically in the 1930s.

At the end of the 1930s, the Harvard economist Alvin Hansen, who was a leading disciple of Keynes, put forward the theory of secular stagnation. Essentially his idea was that, for a variety of reasons including importantly demography and the exhaustion of investment opportunity with the end of the American frontier, there was going to be a tendency in the United States for saving to chronically exceed investment—and therefore for the economy to be short of its full potential a substantial fraction of the time and for there to be at least some deflationary bias in the economy.[2] Hansen turned out to be completely wrong but completely wrong in a way that suggests that at some future point he could turn out to be right. What actually happened was that the Second World War came along, which brought about a vast increase in demand, and the government became a mass absorber of savings by running large deficits. That propelled the economy very rapidly forward to the point where the unemployment rate went from 11 percent in 1940 to close to 1 percent by 1942.[3]

It was generally believed that, after the Second World War, the economy would return to depression. That turned out to be wrong for several reason. The postwar economy was aided by pent-up demand released by the end of wartime credit controls and rationing; the economy was also aided by a massive government project to build out suburbia. The economy was also importantly aided by an unexpected (and still partially unaccounted for) rise in fertility to nearly four children per woman, which created the Baby Boom. And so the prediction of a chronic excess of saving over investment that Hansen made did not turn out to be hugely relevant in his period. My thesis is that Hansen was a couple of generations too early but that the issues he identified can be a chronic problem for capitalist economies. Moreover, they quite likely are a chronic problem for the economies of the industrialized world today.

Table 1 offers some data on interest rates as inferred from swaps. The numbers are quite small, whether measured by swaps or government bonds. Nominal interest rates are low in the United States and extraordinarily low in Japan and Europe. Inflation rates are expected to be well below 2 percent for the *next ten years*. Nowhere do markets think that central banks are going to get even very close to their 2 percent inflation targets.

Perhaps most striking is what the expectation is for real interest rates, which can be thought of as the terms for which one can exchange purchasing power this year for purchasing power in the future. The real interest rate in the industrial world is essentially zero. Now, one can

Table 1

Zero Long-Term Real Rate: Swaps

	10-year interest rates and expected inflation		
	USA*	Japan	Euro
Nominal swap rate	1.26	0.14	0.24
– Inflation swaps	1.51	0.21	1.08
Real swap rate	–0.25	–0.07	–0.83

NOTE: *Adjusted for the 0.35-percentage-point average difference between CPI and the Federal Reserve's preferred PCE inflation rate.

SOURCE: Bloomberg.

argue technical matters involving risk premiums and liquidity premiums and the differences between swaps and government bonds. And while those arguments are interesting, irrespective of any technical adjustments, markets are signaling that for the next 10 years industrial world central banks will uniformly miss their inflation targets despite interest rates far below historical norms.

To preview my argument: If one thought ex ante that there would a big increase in saving and a big reduction in investment propensity, then one would expect that interest rates would fall very dramatically. And so this fall in rates is consistent with a market judgment that the secular stagnation hypothesis is true.

Markets have come to this judgment slowly and painfully. Figure 5 depicts the overnight-indexed swap market forwards for short-term interest rates. There is an old joke about the price of shale oil. This joke no longer applies, but the joke is that the price at which it would be profitable to extract shale oil has been constant for 40 years: at *the current price of oil plus $10*. And in the same way, the view of the market has been that normalization is 6 to 9 months away and will take place at a reasonably rapid rate. And that has been the consistent view of the market since 2009. That view has turned out to be very overly optimistic when compared with what happened. And as overly optimistic as the market has been, the Federal Reserve has been even more so. Figure 6 depicts how the Federal Reserve has consistently been above the market, particularly in the past couple of years, with the consistent overoptimism leading to persistent downward revisions in interest rate projections.

In sum, the evidence of unusually stagnant growth is overwhelming, as is the evidence that there is a market expectation of extraordinarily low inflation and extraordinarily low real interest rates going forward. For example, the swap curve in Figure 7 implies that rates are never going much above 2 percent.

This is not all about the financial crisis. This is not all about the current business cycle. There has been a trend—for as long as there have been indexed bonds to use to measure it—downward of real interest rates over the long term.

Figures 8 and 9 depict that one could have used simple extrapolation in 2007 to more or less predict the general area in which real interest rates are today even without the impending crisis. This is not an artifact of near-term developments or of particular things that the Federal Reserve is doing as seen by the so-called five-year five-year rate, which is the market expectation of what the five-year real interest rate will be five years from now. So it is not about anything ephemeral or current; it is essentially the same picture of a very substantial and continuing decline over the longer term.

Capital is mobile around the world, so economic theory predicts the industrial world as a whole would experiencing similar dynamics. And sure enough, the foreign 10-year yield adjusting for expected inflation exhibits exactly this picture of a sharply and continuing declining real interest rate (Figure 10).

What is the natural way to explain these phenomena? The natural way to explain the phenomena is to imagine that there has been a set of forces that have pushed up the saving propensity and pushed down the investment propensity, and therefore the interest rate that has been necessary to equilibrate them has been under substantial downward pressure and

Figure 5

Market Has Been Too Optimistic…

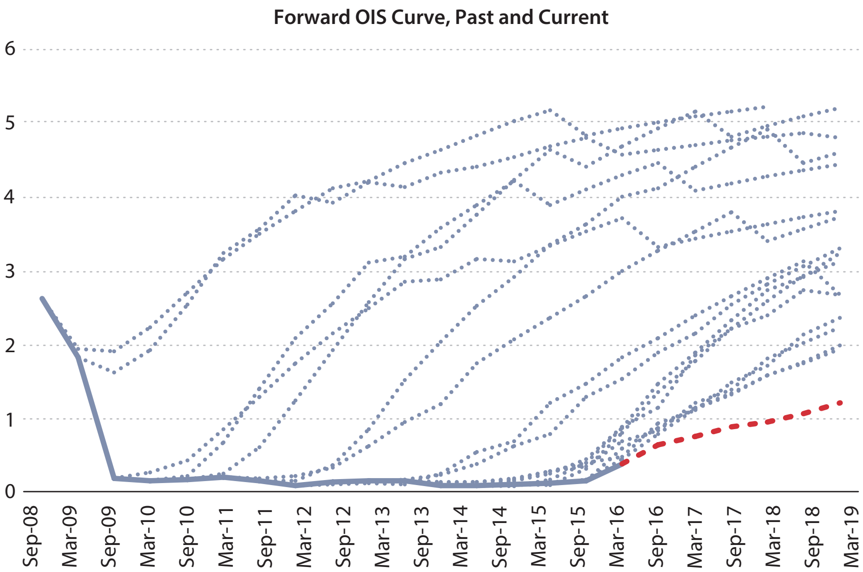

SOURCE: Bloomberg.

Figure 6

…But Consistently More Pessimistic Than the Fed…

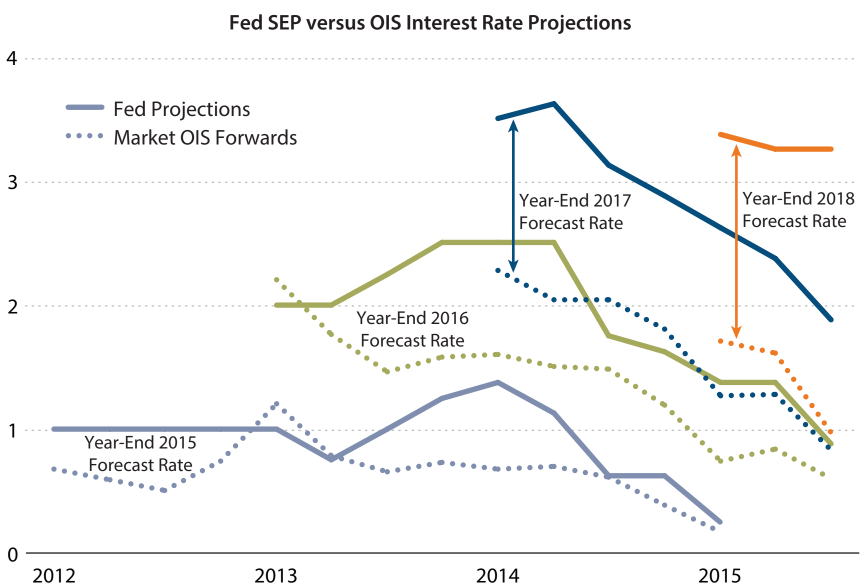

SOURCE: Federal Reserve Summary Economic Projections Fed Funds Median Projection; Bloomberg.

Figure 7

…And Is Now Pricing Very Low Rates

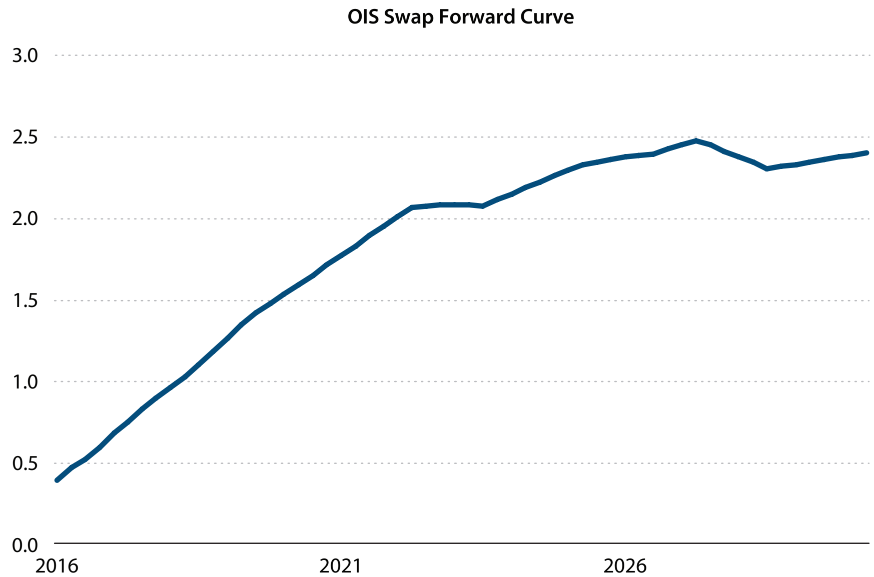

SOURCE: Bloomberg.

Real Rates Have Fallen Steadily…

Figure 8

U.S. TIPS 10-Year Real Yield

SOURCE: Bloomberg.

Figure 9

U.S. TIPS 5- to 10-Year Real Yield

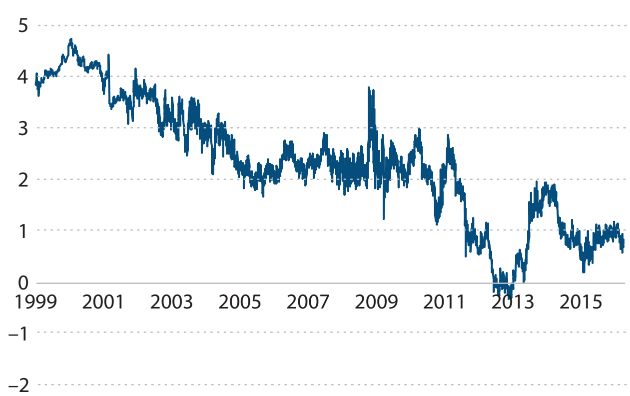

SOURCE: Bloomberg.

Figure 10

…Along With Rates in the Rest of the World

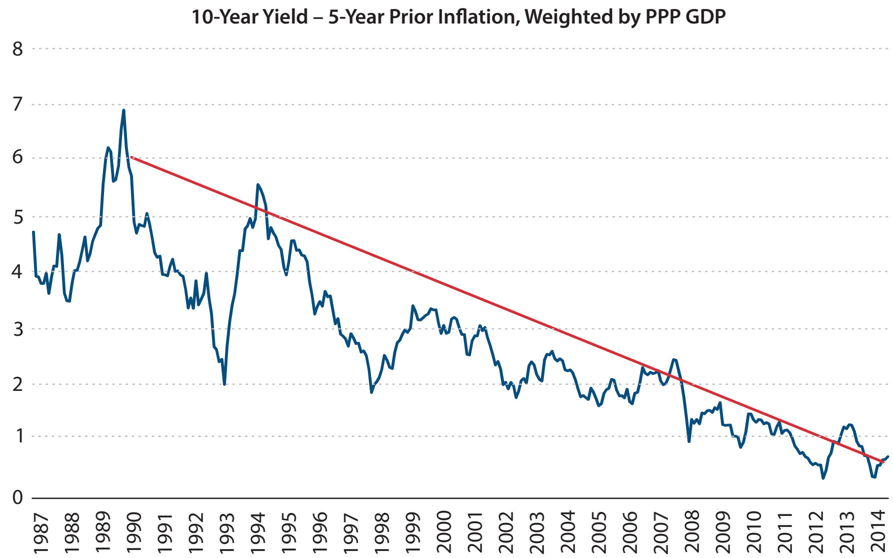

SOURCE: Bloomberg; IMF.

that at some points they have not been equilibrated. And so the mechanism that equilibrated them with excess saving and insufficient demand has been a lower level of income than would otherwise take place. Is that a plausible view? Well, one way to answer the question (although I do not think it is ultimately the most persuasive) is to try to measure what the neutral real interest rate is: more technically, to try to measure what the intercept is in the Taylor rule and thereby answer the question "when the economy is normal, what real interest rate do we observe and has that changed?"

If the economy is strong, the real interest rate would be expected to be above the then-current rate. If the economy is slow, we figure the neutral rate is below the then-current real rate. A variety of econometricians have attempted to do this, shown in Figure 11. Most conclude that the current real neutral rate is substantially negative and has been trending somewhat downward over time (Figure 12).

The Fed has noticed and over the past few years has revised downward its estimate of the neutral real rate. But the Fed's adjustments (Figure 13) are much smaller than the trends that the market is seeing as reflected in long-term real rates. And they are pretty small compared with the adjustments that are suggested by the econometric estimates to measure neutral real rates.

So what I have not yet discussed—but which is maybe the most important question for thinking about these issues—is, "What is it that caused saving to rise and investment to decline and therefore to create this downward pressure, this tendency toward stagnation?"

There is actually a plenitude of factors that could bring about this outcome (Figure 14). Increases in the savings propensity might be expected to come from rising inequality and the

Figure 11

Empirical Estimates of the Real Natural Rate

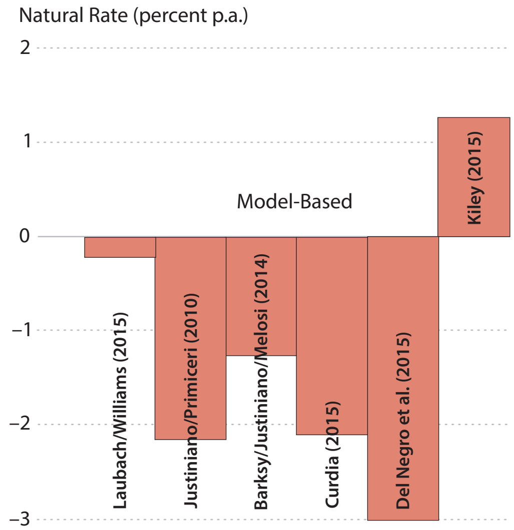

SOURCE: Anna Cieslak discussion of "The Equilibrium Real Funds Rate: Past, Present and Future"; http://www.brookings.edu/~/media/Events/2015/10/interest-rates/Disc_HHHW_02.pdf?la=en.

Figure 12

Decline in Real Equilibrium Rate

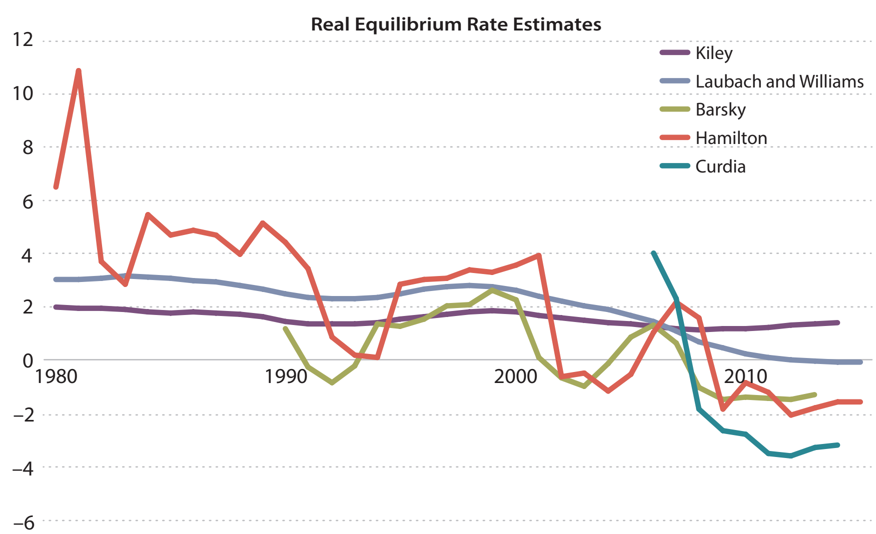

SOURCE: Michael Kiley "What Can the Data Tell Us About The Equilibrium Real Interest Rate," Laubach & Williams "Measuring The Natural Rate Of Interest Redux," Hamilton et al. "The Equilibrium Real Funds Rate," Vasco Curdia "Why So Slow? A Gradual Return For Interest Rates," Barsky et al. "The Natural Rate & Its Usefulness for Monetary Policy Making."

Figure 13

Fed Estimates of Neutral Rate Have Fallen

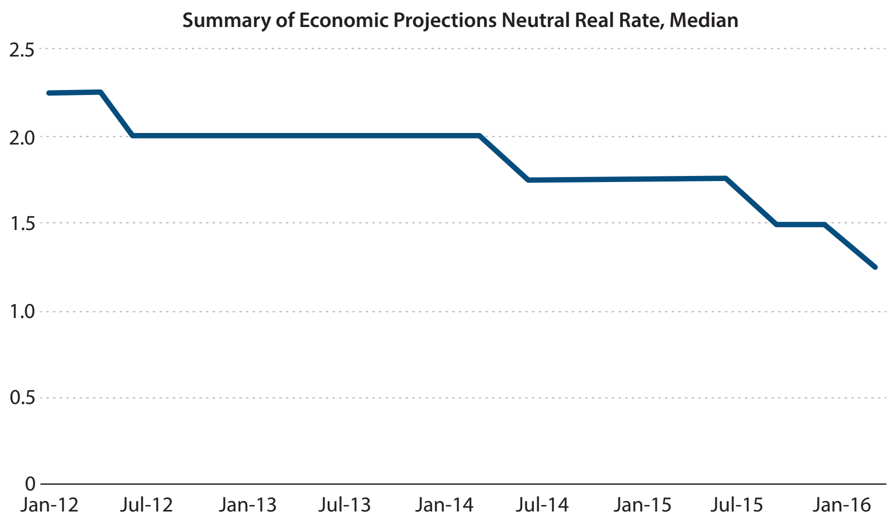

SOURCE: Federal Reserve Quarterly Summary of Economic Projections, 2012 to the present.

observation that the rich save more of their income than the poor. It might be expected to come from what we know we have seen: very substantial reserve accumulation in developing countries. And from some developing countries, notably China in the past nine months, a very substantial capital flight. Indeed, the capital flight from China over the past nine months has been, depending on how one measures it, about three times as large as all the capital flight from all the emerging market crises of the 1990s. Longer life expectancy, more resistance to debt, more uncertainty, household deleveraging, and paying back debt are forms of saving.

What about the investment propensity? The labor force of the industrialized world is not going to grow over the next 20 years. That has implications for the demand for new equipment to equip workers, the demand for new housing to house them, and the need for new business structures to give them a roof over their heads as they work. Think about the de-massification of the economy. I was speaking with somebody in real estate not long ago who said that 15 years ago a good law firm would require 1300 square feet per lawyer. But today at least some law firms are trying to allocate less than 600 square feet per lawyer. Technology has lessened the demand for space-consuming filing cabinets, paralegals, and assistants. Offices can also be smaller since people are likely to spend more time with their laptops working at home. All of that means downward pressure on physical investment. Other examples include e-commerce and the demand for malls, Air B&B and the demand for hotels. Not to mention Uber and the driverless car that will come and alter the demand for and number of automobiles.

People have tried to measure this at least since Alan Greenspan in the 1990s and have found that the total number of pounds of stuff per dollar of GDP is in a secular downward trend. And since a lot of that stuff is capital that means less investment. Capital goods are

Figure 14

Why Might Equilibrium Real Rates Have Fallen?

Increased Savings

1. Changes in distribution of income and profits share
2. Reserve accumulation or capital flight
3. Increasing deleveraging and retirement preparation

Decreases in Investment Propensity

1. Declining growth rate of population and/or technology
2. Demassification of the economy
3. Fall in price of capital goods

Other Factors

1. Interaction between inflation and after-tax real rates
2. Increased frictions in intermediation
3. Increased global safe asset demand

getting cheaper, relative to other goods—most obviously anything to do with information technology. And that means one dollar of saving buys a lot more capital than it used to. That reduces the demand for capital as well. There are other factors, including the fact that financial intermediation works less efficiently because people are more nervous because of what happened and because of regulation. These forces push interest rates down. All of this taken together makes it highly plausible that we have seen a very substantial and structural increase in saving and decrease in investment resulting in low rates, resulting in a tendency toward economic sluggishness.

In important respects this is the deep cause of the financial crisis. After all, during the 2003-07 period the economy did fine overall. It did not overheat, it did not perform spectacularly. It did fine. And yet what did it take to propel it to "fine"? It took what we now see as a vast erosion of credit standards and the mother of all bubbles in the housing market. Without those things, growth would have been inadequate in 2003-07. Or, to put the point differently, if the economy had been overheating, the Fed would have provided the interest rate increases that would have cut off the housing bubble much sooner. And so this structural tendency toward too much saving relative to investment is both the reason why we have sluggishness now and the reason why the periods of robust growth increasingly prove not to be sustainable. Indeed, here is a potted version of recent American economic history. We have had a very slow recovery. Before that we had the financial crisis. Before that we had growth on what we now know to have been an unsustainable financial foundation. Before that we had the recession and slow recovery from 2001 to 2003. And before that we had rapid growth propelled by the Internet bubble. And so if one asks when was the last time that the American economy grew at a rapid rate with clearly strong financial underpinnings, one would have to go back

Figure 15

No Obvious Reason for Equilibrium Rate to Rise

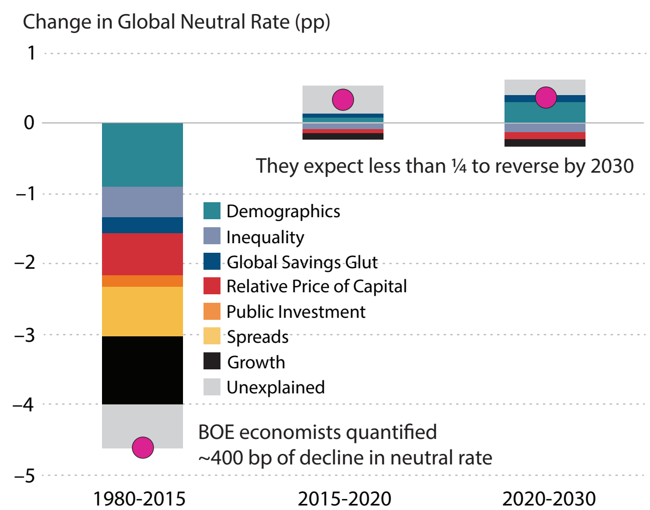

SOURCE: Bank of England Staff Working Paper #571, "Secular drivers of the global real interest rate" by Lukasz Rachel and Thomas Smith; http://www.bankofengland.co.uk/research/Pages/workingpapers/2015/swp571.aspx.

close to a generation ago to find that time. And the underlying reason is this excess of saving over investment.

Through most of this period the Federal Reserve has been inclined to say, "Well, the economy is held back by headwinds but the headwinds will recede as a recovery continues and then we will be in a position to normalize rates." That is the view that was behind all the overoptimism displayed in Figure 6 showing expected rate increases that never materialized. Belief in temporary headwinds was a very reasonable thing to think and say in 2009 or 2010. There were all kinds of huge aftershocks and aftereffects of the financial crisis. It may well have been a reasonable thing to say in 2011 or even in 2012. But six or seven years after the financial crisis, it is hard to understand what the headwind is that one could reasonably expect to recede in the next two or three years. And I believe the answer is that it was not a temporary headwind; it was a permanent headwind driven by the factors that I have just described.

Lukasz Rachel and Thomas Smith at the Bank of England have attempted to quantify and forecast each one of these factors. Their judgment is that, on a global basis, real rates declined 400 basis points between 1980 and 2015 and that there is no particular reason to think that any significant part of that is coming back over the next 15 years (Figure 15).

I believe there is a good chance that judgment is correct. More saving, less investment, very low interest rates, problematic financial sustainability, and a tendency toward economic weakness. That is the secular stagnation hypothesis. Now, there are some alternatives to it (Figure 16). Some of them are kind of the same, and some of them I think are incomplete or inadequate as an explanation.

Figure 16

Differential Diagnosis

Is weakness the result of deleveraging (Rogoff)?

- Does not account for long-term downtrend in rates
- What temporary headwinds still exist seven years into recovery?

Is stagnation mostly on the supply side (Gordon)?

- Lack of supply should lead to inflation, not deflation
- Likely correct that potential output and productivity growth is slow

Is this a global savings glut (Bernanke)?

- Emerging market surpluses have relatively declined
- A glut of savings and excess of savings over investment are very similar diagnosis

Is this a classic liquidity trap (Krugman)?

- Liquidity traps usually thought of as temporary

My Harvard colleague Ken Rogoff[4] talks about a debt supercycle coming to an end and the need for deleveraging. A good story for 2011; not a very good story eight years after the crisis when financial conditions seem very much to have normalized. Also, not a story that explains why there has been a secular downward trend in real interest rates for between 20 and 30 years and that the pre-crisis trend tracks the current level of real interest rates. Bob Gordon suggests, importantly, that there is a great deal of structural weakness on the supply side[5]—that whatever happens with demand, the economy's potential is rising less rapidly than it once did. And there is powerful evidence for that. The fact that GDP growth has been very weak but that unemployment is really quite low suggests that something very bad has happened to productivity. And that may well have to do with the various factors that Gordon described.

However, economists have a classic test for telling whether something is a supply-side phenomenon or a demand-side phenomenon. If the price of apples goes up and one wants to know whether it is because there is a shortage of apples or because apples have come into fashion, they should look at what happens to the quantity of apples. If the price of apples has gone up and there are a lot more apples being sold, then it is because there was more demand. If the price of apples goes up and the quantity of apples goes down, there was likely a problem at the orchards. What has happened recently? There have been steadily declining rates of underlying inflation in the industrial world. And long-term inflation expectations, despite the recovery, are very low by historical standards. So, yes, there are supply issues; but I think it would be a mistake to explain everything in terms of supply issues. The last time there were major supply issues was in the 1970s, and inflation today does not resemble inflation then.

Former Fed Chairman Ben Bernanke has suggested that low interest rates have been caused by a savings glut.[6] I do not really want to argue with that. A savings glut is, after all, a lot like an excess of saving over investment, which is what I have posited. Bernanke sees the

Table 2

Current Forecasts Imply 35 Percent Chance of Recession Within Two Years

	2016	2017
Fed forecast	2.20	2.15
– Recessionary growth	0.40	0.40
Recessionary forecast error	1.80	1.75
% RMSE of GDP forecasts	1.80	2.10
	1.00	0.83
Percent chance recession	16	20

SOURCE: Federal Reserve Summary of Economic Projections; Federal Reserve Updated Forecast Errors.

Table 3

Always a Decent Chance of Recession Intermediate Term

	Three+ year-old expansions Percent of time recession within		
	2 Years	3 Years	5 Years
Japan	30%	40%	54%
Germany	53%	74%	98%
U.K.	28%	40%	63%
U.S.	43%	63%	88%

SOURCE: NBER; Economic Cycle Research Institute.

problem very much in terms of excess saving in developing countries. I look at the length of the trend, the variety of other factors that we have seen, and conclude that the problem is likely global and more permanent in nature.

Finally, Paul Krugman and others talk about the liquidity trap. Again, the liquidity trap is closely related to the difficulty in lowering interest rates enough to balance saving and investment that I have been talking about. The liquidity trap is usually, and in all of Krugman's formal modeling, treated as a temporary phenomenon. And it is the essence of the argument that this excess of saving over investment may be a phenomenon for the next era, not simply a temporary and cyclical phenomenon.

Here is why this is a critical problem that should be a focus of concern for policymakers in the United States. I do not know when the next recession is going to come. If one simply takes the prevailing Fed forecast and assumes the historical level of forecast accuracy, one would conclude recession odds of about one in six this year and about one in three over the next two years. (See Table 2.)

But another way to ask the question is to just ask what are the odds, if an expansion has been going for three years, that it will end sometime in the next two years? And here is the answer to that question in Table 3. One can interpret the data many ways, but I think it is hard to escape the conclusion that the odds are better than 50 percent that within three years the U.S. economy will have gone into recession.

Figure 17 shows why the possibility of recession is so concerning. It lists the decline in the real federal funds rate in the last nine business cycles. Both real and nominal short rate declines have averaged around 5 percentage points. Now, given what I have been arguing and the market is saying about future interest rates, it is highly unlikely that when the next recession comes there will be nearly enough room for the standard monetary policy response. The strategy for the past 40 years has been to respond to recessions with monetary easing. That monetary easing has typically been about 5 percentage points. And a best guess of the market of where interest rates will be when the next recession comes is that they will be at 2 percent.

Figure 17

Large Rate Cuts Are Often Necessary

	Real funds rate easings		
	Start	Final	Easing
May 1960	1.9	–0.1	2.0
August 1966	3.1	0.7	2.3
November 1970	4.5	–0.9	5.4
November 1974	6.4	–1.6	8.0
May 1981	8.7	–0.1	8.8
September 1984	7.6	3.4	4.2
November 1990	5.5	0.1	5.4
December 2000	4.8	–0.4	5.2
August 2007	3.3	–1.1	4.4
			5.1

SOURCE: Bloomberg; Core PCE Deflator from Bureau of Economic Analysis.

Figure 18

Monetary Policy Options

- Raise inflation target
- Forward guidance
- Broader QE
- Negative rates

And so even if the next recession is typical and no worse, policymakers will either have to rely very substantially on unconventional monetary policy or something else.

Now, what are the monetary policy options that one can talk about? (See Figure 18.) The Fed has expressed confidence in its ability to use so-called unconventional policy. Incidentally, when you have responsibility for something and you are in government, it is a good idea to look like you have the problem under control even if you are not altogether certain that you do. So I do not fault the rhetoric of the Fed. But I do not think it is remotely realistic to think that there is anything like the equivalent of 300 basis points in unconventional monetary policy that is likely to be available when the next recession comes.

What are the tools that people talk about? Principally there are four. There is an argument that the Fed should raise its inflation target. That will be an interesting and important issue at some point, but it is like discussing whether I should lose 40 pounds or lose 60 pounds—a much more interesting issue after I have lost the first 20. And when inflation is consistently below 2 percent, discussing whether optimal is 3 or 4 or 5 percent seems a tad jejune.

There is the idea of forward guidance. This is the suggestion that the Fed can promise that it is going to keep rates low in the future and that will make longer-term rates low and will make everybody optimistic and will make them spend. Here is the problem: People already do not think the Fed's going to raise rates very much. There is not much more to promise than what has already been priced-in. People are already anticipating that rates are not going much above 2 percent.

There is broader quantitative easing. Now, I think quantitative easing was very effective at the beginning when there was substantial illiquidity in many fixed-income markets and when there were large credit spreads and steep term premia. When there are not steep term premia and when there are not large credit spreads, what rate is it that you are actually going to reduce substantially? What is the quantitative easing program that, starting with the 10-year Treasury rate of 1.73 percent, is going to do anything like what a normal 3- or 5-percentage-point reduction in interest rates will accomplish?

And finally, there is the idea of a venture into negative rates, which it seems to me is a quite uncertain experiment. Andrew Haldane at the Bank of England has made a fairly thorough study of this[7]: There is interest rate data going back to Babylonian times, and there is no historical experience with negative rates. But in a world where you can put $20 bills in a suitcase or $100 bills in a suitcase, maybe interest rates can be –30 basis points or –60 basis points. Whatever the true minimum lower bound is, it is not all that big relative to what is necessary the next time a recession comes.

So there is a real question as to how effective unconventional monetary policy can be. But even if it could be effective, there are two other major problems with unconventional monetary policy. The first is what I would call the pull forward problem. The pull forward problem is the way monetary policy works by pulling spending forward. When the interest rate is lower, you buy a new car sooner than you otherwise would. When the interest rate is lower, you invest a lot in order to adjust your capital stock to a higher level because the interest rate is lower.

So if you have a temporary demand shortfall, a reduction in interest rates is terrific because it pulls forward demand to the moment where you have the shortfall. But if you have a permanent demand shortfall, at some point that pull forward catches up with you. At some point, yes, we increased demand in 2015; but the way we did it was by pulling demand from 2016 into 2015. And so unless you are going to have ever easier monetary policy at some point, the treadmill catches up with you. And then there is this: What is the consequence of the spending being generated with these unconventional hyper-liquidity measures? They are surely increasing the risk of making the mistakes of 2005-07. There are surely questions about the quality of the investment that people are not willing to make with a 1.75 percent 10-year rate that they would be willing to make with a 1.25 percent rate.

My conclusion is not that the Federal Reserve has been wrong to try to respond to slack deflationary pressure in excess of saving over investment by letting interest rates and letting financial conditions adopt. My conclusion, though, is that the process has surely run into very severe diminishing returns and cannot be relied on the next time a recession comes. What is the implication of that? I would suggest two. A first implication is that surely it is very important to avoid the next recession coming anytime soon. That one should be prepared to take risks that one would not take in normal times to avoid recession. But the possible risks of recession and entrenched inflation expectations below 2 percent vastly exceed, in terms of their consequences, a hypothetical situation in which inflation rose to 2.3 or 2.7 percent. And that awareness needs to be present in the setting of monetary policy.

Second, normalization is the wrong objective in abnormal times. Normalization relative to historical experience is not appropriate in conditions where the natural, neutral, normal interest rate is very different from what it has been historically. More profoundly, traditional economics has assumed that changes in demand can naturally be offset by the workings of monetary policy. So in the conventional analysis of a policy area, changes in demand are abstracted away. But in a world where monetary policy is much more limited in its impact, the demand aspects of all other policies should be considered because they will not naturally be sterilized by the actions of the Federal Reserve. In that regard, I believe there should be a radical rethinking of the role of fiscal policy in economic stabilization. That with monetary policy able to do much less, fiscal policy will need to do much more. This rethinking of fiscal policy should reflect the fact that the sustainability of debt depends crucially on a comparison of interest rates and growth rates; and when interest rates are chronically and systemically low, debt burdens that would otherwise become imprudent become much more prudent, particularly in the context of getting an economy to grow.

And I will conclude by observing that at a moment when money has never been cheaper, a moment when materials cost are near record lows, a moment when those with the capacity to do construction work are still out of work in disproportionate numbers, it is difficult to imagine a better moment to stop paint chipping off school buildings or to fix LaGuardia Airport. And the United States right now has the lowest infrastructure investment rate that it has had since the Second World War. If you adjust for depreciation and round to the nearest integer, that investment rate is zero. And that seems to me manifestly inappropriate. And the only thing that is more wrong than the current low infrastructure investment rate is the absence of any systematic planning to rapidly build up and launch infrastructure when the next recession comes.

Policymakers need to adapt to a world that is very different. A world where the problems are as much of deflationary risks as of inflationary risks, of insufficient demand as of inadequate supply. A world of limits on the power of monetary policy. A world where fiscal policy matters, whether one has my kind of views of fiscal policy directed at public investment or other views of fiscal policy directed at stimulating private investment. This is not a matter of political ideology, but it is a matter of economic reality. A moment of very different economic conditions will require very different economic policy responses. And that is something both those charged with monetary policy responsibility and those who create the political frameworks in which monetary policy operates need to keep very much in mind in this era of secular stagnation. ■

NOTES

[1] This paper is adapted from a speech that included the following introduction: "Before I get into the substance of what I want to say, let me salute the St. Louis Fed. I was not a fan of hard-core monetarism. I think neo-Fisherian economics has an 80 to 90 percent chance of being mostly a confusion. But I think it is hugely important to our processes that there be experienced practical people exploring and pushing doctrines that are counterintuitive to the conventional wisdom. And so I salute the St. Louis Fed, which I think actually has stood out within the Federal Reserve System for nerve and verve in attempting to take on a role of intellectual leadership. And I think it is very, very important. And I have paid as much or more attention—even while disagreeing with some of it—to the research that has come out of the St. Louis Fed as I have to the research that has come out of any of the other 11 regional Reserve Banks. And so I salute Jim Bullard on his leadership, and I salute his predecessor, Bill Poole, and predecessors going farther back—as well as those who succeeded Homer Jones in the research department. Believe me, you are making a very important contribution to our national life."

[2] "Secular" refers to the permanence of the stagnation.

[3] https://research.stlouisfed.org/fred2/data/M0892AUSM156SNBR.txt.

[4] http://voxeu.org/article/debt-supercycle-not-secular-stagnation.

[5] http://piketty.pse.ens.fr/files/Gordon2015.pdf.

[6] http://www.brookings.edu/blogs/ben-bernanke/posts/2015/03/31-why-interest-rates-low-secular-stagnation.

[7] http://www.bankofengland.co.uk/publications/Documents/speeches/2015/speech828.pdf.

Market Power and Asset Contractibility in Dynamic Insurance Contracts

Alexander K. Karaivanov and Fernando M. Martin

The authors study the roles of asset contractibility, market power, and rate of return differentials in dynamic insurance when the contracting parties have limited commitment. They define, characterize, and compute Markov-perfect risk-sharing contracts with bargaining. These contracts significantly improve consumption smoothing and welfare relative to self-insurance through savings. Incorporating savings decisions into the contract (asset contractibility) implies sizable gains for both the insurers and the insured. The size and distribution of these gains depend critically on the insurers' market power. Finally, a rate of return advantage for insurers destroys surplus and is thus harmful to both contracting parties. (JEL D11, E21)

Federal Reserve Bank of St. Louis *Review*, Second Quarter 2016, *98*(2), pp. 111-27.
http://dx.doi.org/10.20955/r.2016.111-127

1 INTRODUCTION

Households face fluctuations in their incomes but desire stable consumption. Prime examples of shocks to income are variations in labor status and changes in health. Maintaining savings in liquid and low-risk assets—for instance, in the form of government bonds or savings accounts—allows households to mitigate the impact of negative income shocks on their standard of living. Similarly, positive income shocks provide the opportunity to accumulate savings to use in bad times. However, savings are an imperfect way to insure against idiosyncratic shocks: For instance, the return on a deposit does not increase because the depositor is laid off or sick. Hence, a natural way to complement self-insurance through savings is to contract with an insurer (private or government-run) willing to absorb an agent's individual risk. In a perfect world, the parties would sign a long-term contract that maximizes the surplus generated by the relationship and fully specifies the time paths of consumption and savings of the insured for all possible combinations of future income states.

In practice, however, economic actors often cannot commit or are legally barred from committing to a long-term contract. For example, consider typical labor, housing, and personal or property insurance contracts: Costless renegotiation or switching providers is always

Alexander K. Karaivanov is a professor of economics at Simon Fraser University. Fernando M. Martin is a senior economist at the Federal Reserve Bank of St. Louis. Alexander Karaivanov acknowledges the financial support of the Social Sciences and Humanities Research Council of Canada (grant No. 435-2013-0698).

possible, although sometimes only at fixed time intervals. In addition, while insurers are frequently aware of an agent's net worth or assets, they may or may not have the ability to control private asset accumulation. The latter ability, however, can be key to the interplay between self-insurance and market- or government-provided third-party insurance (e.g., Arnott and Stiglitz, 1991). As an example, government social security schemes (old-age insurance) usually have both voluntary and controlled/forced savings components. Various mixtures of components exist around the world.

We study the above issues and trade-offs in a multiperiod risk-sharing setting that features a risk-neutral insurer and a risk-averse agent endowed with a stochastic income technology and the ability to save at a fixed rate of return. We assume that the parties cannot commit to a long-term contract: Both the agent and the insurer can commit only to *one-period* risk-sharing contracts. In this setting, we show that there are still large gains from third-party insurance and the ability to incorporate the agent's savings decisions into the insurance contract.

Specifically, we model the interaction between the agent and insurer by assuming that they periodically bargain over the terms of the contract. Formally, we do so by adopting the solution concept of a Markov-perfect equilibrium (MPE), as in Maskin and Tirole (2001). This solution captures our notion of limited commitment, since contract terms are a function of only payoff-relevant variables (in our setting, the agent's assets and the income realization) and the idea that bygones are bygones. That is, the past does not matter beyond its effect on the current state.

We find that the agent's asset holdings are a key feature of Markov-perfect insurance contracts, as the assets determine the agent's endogenous outside option. Given that feature, we analyze the role of asset/savings contractibility by comparing the case of "contractible assets" (when the insurer can fully control the agent's savings decisions) with the case of "non-contractible assets" (when the agent can privately decide on the amount of his savings, even though the asset holdings are observed by the insurer). In many situations, governments, insurance companies, banks, and so on may have information about agents' assets but, for legal or other reasons, are unable to directly control agents' savings choices. In other situations—for example, social security—the opposite is true.

We show that asset contractibility affects the insurance contract terms and the degree of achievable risk-sharing compared with self-insurance, except in the limit when insurance markets are perfectly competitive (free entry). Intuitively, whenever the insurer has market power (not necessarily monopoly power) and thus can generate positive profits from insuring the agent, private asset accumulation provides the agent with an instrument to "counter" the insurer by controlling his future outside option. Essentially, larger savings by the agent today imply a larger outside option tomorrow since the agent would be better able to self-insure. On the insurer's side, however, a larger outside option for the insured implies lower profits. We show that this misalignment of incentives between the contracting parties, which originates in the commitment problem, causes a welfare loss to *both* sides when the agent's assets are non-contractible.

Numerically, we assess the degree to which the presence of third-party insurance improves agents' welfare beyond that achievable on their own through savings. We show that the welfare gains for the poorest agents (zero assets) can be as high as 4.5 percent of their autarky consumption per period. This number is significantly larger than the cost of business cycle fluctuations (about 0.1 percent), a common benchmark for welfare calculations in macroeconomic applications. In terms of the role of asset contractibility, the largest welfare loss if agents' savings are non-contractible is about 0.4 percent of autarky consumption per period.

We also find that the market power of insurance providers significantly affects the welfare gains that agents derive in Markov-perfect insurance contracts and, to a lesser extent, the welfare losses when agents' assets are non-contractible. The welfare gains from third-party insurance are strictly decreasing in the insurers' market power, whereas the welfare costs of asset non-contractibility peak at an intermediate value of market power, somewhere between the monopolistic insurer case and perfectly competitive insurance markets.

Finally, our numerical results suggest that both the insured and the insurer are better off if there is no return on assets differential between them. A higher intertemporal return—or, equivalently, discount rate—for the insurer relative to the insured reduces the total surplus that can be generated in the risk-sharing relation. Furthermore, differences in the parties' rates of return on assets amplify the distortions in the time profiles of consumption and savings (relative to the equal return benchmark) that arise from the limited commitment friction.

Our article builds on and extends in several dimensions our previous analysis (Karaivanov and Martin, 2015). In that article, we introduced the idea of Markov-perfect insurance contracts and showed that limited commitment on the insurer's side is restrictive only when he has a rate of return advantage over agents with sufficiently large asset holdings. The limited commitment friction makes assets carried by agents essential in an MPE, as they cannot be replaced with promises of future transfers. In contrast, if the insurer and the insured have equal rates of return on carrying assets over time, we showed that Markov-perfect insurance contracts result in an equivalent consumption time path as a long-term contract to which only the insurer can commit because assets and promised utility are then interchangeable. While we retain the basic idea of Markov-perfect insurance, our analysis here differs in two important aspects. First, unlike in Karaivanov and Martin (2015), we allow agents' assets to be non-contractible. Second, instead of assuming an arbitrary asset-dependent but otherwise exogenous outside option for the agent, we endogenize the division of the gains from risk-sharing by defining and analyzing a bargaining problem between the parties.

This article also differs from the literature on optimal contracts with hidden savings (see Allen, 1985, and Cole and Kocherlakota, 2001, among others) that assumes that the principal has no ability to monitor the agent's assets. The main result in these articles is that no additional insurance over self-insurance may be possible, unlike in this article. On the technical side, our assumption of observable assets (even if non-contractible) helps us avoid dynamic adverse selection and the possible failure of the revelation principle with lack of commitment (Bester and Strausz, 2001), while still preserving the empirically relevant intertemporal implications of savings non-contractibility.

More generally, in the dynamic mechanism design literature, allowing agents to accumulate assets in a principal agent relationship typically yields one of the following three results, depending on the specific assumptions about the information or commitment structure:

(i) An agent's assets play no role (when the insurer can control the agent's consumption). (ii) Assets eliminate the insurer's ability to smooth the agent's consumption beyond self-insurance (Allen, 1985, and Cole and Kocherlakota, 2001). Or (iii) the environment becomes highly intractable (Fernandes and Phelan, 2000, and Doepke and Townsend, 2006). In contrast, we show that Markov-perfect insurance contracts result in simple dynamic programs with a single scalar state variable and avoid the curse of dimensionality, including the case with non-contractible savings.

2 THE ENVIRONMENT

Consider an infinitely lived, risk-averse agent who maximizes discounted expected utility from consumption c. The agent's flow utility is $u(c)$, with $u_{cc} < 0 < u_c(c)$ and u satisfying Inada conditions.[1] The agent discounts the future by factor $\beta \in (0,1)$. Each period the agent receives an output/income endowment, which he can consume or save. Output is stochastic and takes the values $y^i > 0$ with probabilities $\pi^i \in (0,1)$ for all $i = 1,\ldots,n$, with $n \geq 2$ and where $\sum_{i=1}^{n} \pi^i = 1$. Without loss of generality, let $y^1 < \ldots < y^n$.

The risk-averse agent would like to smooth consumption over output states and over time. We assume that the agent can carry assets a over time by means of a savings (storage) technology with fixed gross return $r \in (0,\beta^{-1})$. Let $\mathbb{A} = [0,\bar{a}]$ denote the set of feasible asset holdings, where $\bar{a} \in (0,\infty)$ is chosen to be sufficiently large that it is not restrictive. In contrast, the lower bound on $\mathbb{A}$ is restrictive and represents a borrowing constraint. Assuming that assets cannot be negative means that the agent cannot borrow—that is, he can only save.

Suppose that the agent has no access to insurance markets and therefore can rely only on self-insurance through savings—running up and down a buffer stock of assets as in Bewley (1977). In this situation, which we label "autarky," the agent's optimal consumption and savings decisions depend on his accumulated assets and are contingent on the output realization. That is, given realized output y^i, the agent carries into the next period assets $a^i \geq 0$ and consumes $c^i \equiv ra + y^i - a^i$.

Formally, the agent's problem in autarky can be written recursively as

$$\Omega(a) = \max_{\{a^i \geq 0\}_{i=1}^{n}} \sum_{i=1}^{n} \pi^i \left[u\left(ra + y^i - a^i\right) + \beta \Omega(a^i) \right]. \tag{1}$$

By standard arguments (e.g., Stokey, Lucas, and Prescott, 1989), our assumptions on u ensure that the autarky value function $\Omega(a)$ is continuously differentiable, strictly increasing, and strictly concave for all a $\in \mathbb{A}$. The autarky (self-insurance) problem is a standard "income fluctuation" problem, versions of which have been studied, for instance, by Schechtman and Escudero (1977) and Aiyagari (1994), among many others. The properties of the solution are well known: imperfect consumption smoothing (c^i differs across states with different y^i); consumption c^i and next-period assets a^i in each income state increasing in current assets a; asset contraction (negative savings) in the lowest income state(s); and asset accumulation (positive savings) for some range of asset holdings in the highest income state(s).

Since the rate of return on assets is assumed to be smaller than the agent's discount rate, $r < \beta^{-1}$, the agent saves only to insure against consumption volatility.[2] In particular, there is a *precautionary* motive for saving because the agent wants to mitigate the chance of ending up with zero assets, a situation in which he would be unable to self-insure against negative income shocks. Note that since assets provide the same return in all output states, the agent is unable to insure perfectly against income fluctuations. Thus, there is a demand for additional insurance as addressed in the next section.

3 INSURANCE

Suppose there exists a risk-neutral, profit-maximizing insurer. Throughout the article, we assume that the insurer can costlessly observe output realizations y^i and the agent's assets a. The insurer can borrow and lend, without restrictions, at gross rate $R > 1$. The insurer's future profits are also discounted at the rate R. The parameter R can have either a technological or preference interpretation. The special case $r = R$ can be thought of as the insurer having the ability to carry resources intertemporally using the same savings technology as the agent. If, instead, $R = \beta^{-1}$, we can think of the agent and insurer as having the same discount factor—a standard assumption in the literature. In general, we allow R to take any value between these bounds, as stated in Assumption 1 below.

Assumption 1 $0 < r \leq R \leq \beta^{-1}$, *with* $\mathrm{r} < \beta^{-1}$, *and* $R > 1$.

3.1 The Agent's Savings Decision

Suppose the insurer, while observing the agent's assets a, cannot directly control the agent's savings decision—namely, the choice of a'. We can think of the insurance arrangement between agent and insurer in any time period as the exchange of output y^i for gross transfer τ^i (this includes the insurance premium or payoff in the different states of the world). Transfers are allowed to depend on the agent's accumulated assets a, since assets affect how much insurance the agent demands.

Suppose the agent is offered insurance for the current period. What is his savings decision given transfers τ^i? Call period consumption $c^i \equiv ra + \tau^i - a'^i$, as implied by the insurance transfer τ^i, the gross return on the agent's current assets ra, and the agent's savings decision a'^i. Let $v(a'^i)$ denote the continuation value for the agent carrying assets a'^i into the next period. The function v is an equilibrium object that depends on all future agent-insurer interactions, which in turn depend on the level of assets carried into the future. The consumption/savings problem of the agent can then be written as follows:

$$\max_{\{a'^i\}} \sum_i \pi^i \left[u\left(ra + \tau^i - a'^i\right) + \beta v\left(a'^i\right) \right].$$

With the Lagrange multiplier $\xi^i \pi^i \geq 0$ associated with the non-borrowing constraint $a'^i \geq 0$, the first-order conditions are

$$-u_c\left(ra+\tau^i-a'^i\right)+\beta v_a\left(a'^i\right)+\xi^i=0$$

for all $i = 1,\ldots,n$. In other words, given an insurance contract for the current period and anticipating future interactions (contracts) between the agent and the insurer, which yield the continuation value v, the agent's savings decision is characterized by

$$u_c\left(c^i\right)-\beta v_a\left(a'^i\right)\geq 0, \text{ with equality if } a'^i>0. \tag{2}$$

When the agent's savings are non-contractible, the insurer must take into account the agent's savings decision given by (2) when deciding on the insurance transfers τ^i. We call this the agent's incentive-compatibility constraint, as any insurance contract that allows the agent to make his own savings decisions must respect condition (2).

Below we also consider the alternative case in which the agent's savings can be specified (enforced) as part of the insurance contract. In this case, inequality (2) does not restrict the design of the insurance contract offered to the agent.

3.2 Markov-Perfect Insurance

We assume that the agent and the insurer can bargain over the insurance terms each period. The insurance contract is negotiated every period since we assume a limited commitment friction—neither the agent nor the insurer can commit to honor any agreement beyond the current period. This limited commitment friction could be motivated by legal, regulatory, or market reasons. For example, in many real-life situations (labor contracts, housing rental, home and car insurance, and so on) the parties are allowed to (costlessly) modify or renegotiate the contract terms at fixed points of time (e.g., yearly).

If the parties do not reach an agreement, they revert to their respective outside option from then on. Of course, given the limited commitment friction, both parties know that any agreement spanning more than one period is subject to renegotiation and cannot be committed to. The outside option for the agent is autarky, with value $\Omega(a)$ as derived previously. The outside option for the insurer is zero profits.

To model the bargaining game between the agent and the insurer, we adopt the Kalai (1977) solution, which picks a point on the utility possibility frontier depending on a single parameter θ. This parameter can be interpreted as the agent's "bargaining power." Specifically, in Kalai's bargaining solution, a larger value of θ implies that the agent obtains surplus closer to his maximum feasible surplus, while the insurer obtains surplus closer to his outside option. The converse is true for lower values of θ. The limiting case $\theta \to 1$ corresponds to the agent receiving his maximum possible surplus and the insurer receiving his outside option of zero profits. This situation can be interpreted as a market setting with perfect competition and free entry by insurers. In contrast, in the opposite limiting case, $\theta \to 0$, the agent receives his outside option, while the insurer receives maximum (monopoly) profits. Formally, the Kalai bargaining solution postulates a *proportional surplus-splitting rule*, which takes the form $(1 - \theta)S^A = \theta S^I$, where S^A is the agent's surplus, defined as the difference between the agent's value in the contract and his outside option, and S^I is the insurer's surplus, defined analogously.

Let $\overline{y} \equiv \sum_{i=1}^{n} \pi^i y^i > 0$ denote expected output. The insurer's expected period profit is therefore $\overline{y} - \sum_{i=1}^{n} \pi^i \tau^i$. Equivalently, using $c^i = ra + \tau^i - a'^i$, we can rewrite the insurer's profit in terms of the agent's consumption and next-period assets as $\overline{y} + ra - \sum_{i=1}^{n} \pi^i \left(c^i + a'^i\right)$. The participation constraints of the contracting parties are therefore

$$\sum_{i=1}^{n} \pi^i \left[u\left(c^i\right) + \beta v\left(a'^i\right)\right] \geq \Omega(a)$$

$$\overline{y} + ra - \sum_{i=1}^{n} \pi^i \left[c^i + a'^i - R^{-1}\Pi\left(a'^i\right)\right] \geq 0,$$

where v and Π denote the (endogenous) agent and insurer continuation payoffs, respectively, both as functions of the agent's asset holdings.

Assuming $\theta \in (0,1)$, we can write the insurance contract with Kalai bargaining as

$$\max_{\{c^i, a'^i \geq 0\}} \overline{y} + ra - \sum_{i=1}^{n} \pi^i \left[c^i + a'^i - R^{-1}\Pi\left(a'^i\right)\right],$$

subject to (2) for all $i = 1,\ldots,n$ and

(3) $$(1-\theta)\left[\sum_{i=1}^{n} \pi^i \left(u\left(c^i\right) + \beta v\left(a'^i\right)\right) - \Omega(a)\right] - \theta\left[\overline{y} + ra - \sum_{i=1}^{n} \pi^i \left(c^i + a'^i - R^{-1}\Pi\left(a'^i\right)\right)\right] = 0.$$

The insurer's profits are maximized subject to the agent's incentive-compatibility constraint and the proportional surplus-splitting rule.

Since the insurer observes the output realization y^i and there are no private information issues or intratemporal commitment problems, it is optimal that the agent receives full insurance—that is, $c^i = c$ for all i. Formally, this can be shown by taking the first-order conditions with respect to c^i in the constrained maximization problem above and noticing that they are fully symmetric with respect to i. Intuitively, the risk-averse agent is fully insured against his idiosyncratic income fluctuations and all income risk is absorbed by the risk-neutral insurer. Unlike in alternative settings (e.g., with moral hazard or adverse selection), here there are no gains from making the agent's consumption state-contingent since output realizations are exogenous and not affected by any agent actions or type. Assuming a symmetric solution, we also obtain $a'^i = a'$ for all i. In this case (which is assumed hereafter), the insurance contract can be written as

$$\max_{c, a' \geq 0} \overline{y} + ra - c - a' + R^{-1}\Pi(a'),$$

subject to

(4) $$u_c(c) - \beta v_a(a') \geq 0, \text{ with equality if } a' > 0$$

and

$$(1-\theta)\left[u(c)+\beta v(a')-\Omega(a)\right]-\theta\left[\overline{y}+ra-c-a'+R^{-1}\Pi(a')\right]=0. \tag{5}$$

We formally define an *MPE* and a *Markov-perfect insurance contract* in our setting as follows.

Definition 1 *Consider a risk-averse agent with autarky value* $\Omega(a)$, *as defined in* (1), *and bargaining power* $\theta \in (0,1)$ *contracting with a risk-neutral insurer.*
(i) An MPE is a set of functions $\{\mathcal{C}, \mathcal{A}, v, \Pi\} : \mathbb{A} \to \mathbb{R} \times \mathbb{A} \times \mathbb{R} \times \mathbb{R}_+$ *defined such that, for all* $a \in \mathbb{A}$:

$$\{\mathcal{C}(a), \mathcal{A}(a)\} = \underset{c,a' \geq 0}{\operatorname{argmax}}\, \overline{y}+ra-c-a'+R^{-1}\Pi(a'),$$

subject to (4) and (5), and where

$$v(a)=u(\mathcal{C}(a))+\beta v(\mathcal{A}(a))$$
$$\Pi(a)=\overline{y}+ra-\mathcal{C}(a)-\mathcal{A}(a)+R^{-1}\Pi(\mathcal{A}(a)).$$

(ii) For any $a \in \mathbb{A}$, *the Markov-perfect contract implied by an MPE is the transfer schedule*:

$$\mathcal{T}(a)=\mathcal{C}(a)+\mathcal{A}(a)-ra.$$

Solving for an MPE involves finding a fixed point in the agent's value function v and the insurer's profit function Π. We briefly characterize the properties of the MPE with bargaining using the first-order conditions of the insurance problem. With Lagrange multipliers μ, λ, and ζ associated with the constraints (4), (5), and $a' \geq 0$, respectively, the first-order conditions are

$$-1+\mu u_{cc}(c)+\lambda\left\{(1-\theta)u_c(c)+\theta\right\}=0 \tag{6}$$

$$-1+R^{-1}\Pi_a(a')-\mu\beta v_{aa}(a')+\lambda\left\{(1-\theta)\beta v_a(a')-\theta\left[-1+R^{-1}\Pi_a(a')\right]\right\}+\zeta=0. \tag{7}$$

The values of the Lagrange multipliers—specifically, whether or not they are zero—are critical to understanding the equilibrium properties.

Lemma 1 *In an MPE, the Lagrange multiplier on the surplus-splitting rule (5) is positive—that is,* $\lambda > 0$.

Proof. Rearrange (6) as $\lambda\{(1 - \theta)u_c(c) + \theta\} = 1 - \mu u_{cc}$. Given that $u_c > 0$ and $\theta > 0$, the sign of λ is the same as the sign of the right-hand side. Since (4) is an inequality constraint, $\mu \geq 0$. Thus, given $u_{cc} < 0$, the right-hand side of the previous expression is strictly positive, which implies $\lambda > 0$. ■

If, in addition, $\mu > 0$, then (4) implies an interior solution for future assets, and so $\zeta = 0$.[3] Conditions (6) and (7) can then be solved to obtain the values of μ and λ. The optimal consumption and savings (c,a') implied by the Markov-perfect contract with an interior solution for assets are characterized by

(8) $$u_c(c) = \beta v_a(a')$$

and

(9) $$(1-\theta)\left[u(c)+\beta v(a')-\Omega(a)\right]=\theta\left[\bar{y}+ra-c-a'+R^{-1}\Pi(a')\right].$$

We further describe the properties of the MPE insurance contracts numerically in Section 4.1.

3.2.1 Discussion. If assets are contractible and there is a strictly positive rate of return differential between the parties (the case $R > r$), it would be optimal for assets to be carried over time at the higher rate R. However, since in our setting the insurer cannot commit to future transfers, the only way it could take over all the agent's assets would be to appropriately compensate him today. Doing so would imply inducing disproportionately high consumption today, which is not optimal for intertemporal smoothing reasons. This implies that the agent carries assets over time at the lower rate r. Note that the key problem is that the insurer is unable to commit to a long-term disbursement of the returns from assets through future transfers. In contrast, if the insurer could commit to an infinitely long contract, one can show that it is optimal to extract all of the agent's assets at the initial date (see Karaivanov and Martin, 2015, for details).

When assets are non-contractible, the agent can use savings to influence his future outside option $\Omega(a')$. Hence, a conflict between the parties arises whenever the insurer has market power. The insurer would prefer the agent to hold less assets, which implies higher demand for market insurance by the agent because of his lower ability to self-insure and, thus, higher profits for the insurer. In contrast, the agent would prefer larger future assets, a', which would raise his outside option, $\Omega(a')$, by providing a better ability to self-insure. The interplay of these incentives is illustrated in the numerical analysis below.

3.3 Special Cases: Monopoly and Perfect Competition

We previously wrote the Markov-perfect insurance problem for any $\theta \in (0,1)$. To gain more intuition about the properties of its solution, we describe what happens in two limiting cases—as θ goes to 0 or 1. The limiting case $\theta \to 0$ implies that the agent has no bargaining power and corresponds to the case of a monopolist insurer. Note that as $\theta \to 0$, the surplus-splitting rule (5) converges to $u(c) + \beta v(a') = \Omega(a)$. Since the agent's value in an MPE is $v(a) = u(\mathcal{C}(a)) + \beta v(\mathcal{A}(a))$, it follows that $v(a) = \Omega(a)$; that is, the agent always receives present value equal to his outside option. In other words, when the agent faces a monopolist insurer, the insurer receives all gains from the contract and the agent receives the same value as in autarky. Note that this applies regardless of whether constraint (4) binds. However, as we show in the numerical analysis, the savings decision of the agent affects, in general, the profits that the insurer can extract.

The other limiting case, $\theta \to 1$, can be interpreted as the agent having maximum bargaining power (and the insurer having zero bargaining power) and corresponds to the setting of perfect competition (free entry by insurers). Note that as $\theta \to 1$, the surplus-splitting rule (5) converges to $\bar{y} + ra - c - a' + R^{-1}\Pi(a') = 0$. Since in an MPE $\Pi(a) = \bar{y} + ra - \mathcal{C}(a) -$

$\mathcal{A}(a) + R^{-1}\Pi(\mathcal{A}(a))$, this implies $\Pi(a) = 0$; that is, the insurer receives zero expected present value profits. This holds for all asset levels $a \in \mathbb{A}$ and all periods. In turn, this implies that $\Pi(a) = \bar{y} + ra - \mathcal{C}(a) - \mathcal{A}(a) = 0$, or equivalently, $\Pi(a) = \bar{y} - \mathcal{T}(a) = 0$. In other words, if $\theta \to 1$, the insurer makes zero expected profits *per period*. Since this also implies that $\Pi_a(a) = 0$ for all $a \in \mathbb{A}$, as $\theta \to 1$ the first-order conditions (6) and (7) simplify to

$$-1+\mu u_{cc}(c)+\lambda=0$$
$$-1-\mu\beta v_{aa}(a')+\lambda=0.$$

As shown in Proposition 5 in Karaivanov and Martin (2015), with free entry by insurers the agent's value function $v(a)$ is strictly concave. Thus, $v_{aa} < 0$, which, together with $u_{cc} < 0$, implies that the above conditions are satisfied if and only if $\mu = 0$. Intuitively, when the agent receives all the surplus from the risk-sharing contract, there is no misalignment between the insurer and the agent in the values of assets to be held in savings and, thus, how much assets to save and, thus, the incentive-compatibility constraint (4) does not bind.

4 THE ROLE OF ASSET CONTRACTIBILITY

4.1 Theoretical Analysis

Does asset contractibility matter for the degree of insurance and the time profiles of consumption and savings? In other words, how important is it for risk-sharing whether the insurer can or cannot bind the agent to a specific savings level? To answer these questions, we investigate whether, and under what conditions, the incentive-compatibility constraint (4) binds in an MPE. If the constraint does not bind, then whether saving decisions can or cannot be contracted on would not matter for risk-sharing. If the constraint does bind, however, then clearly the agent and the insurer have conflicting views of what savings should be. In the proposition below, we show that asset contractibility generally does matter for the contract terms.

Proposition 1 *In an MPE, if $\Pi_a(a') < 0$ for some $a \in \mathbb{A}$ such that $a' = \mathcal{A}(a) > 0$, then the incentive-compatibility constraint (4) binds—that is, the Lagrange multiplier μ is positive.*

Proof. Suppose $\mu = 0$. Then (6) and Lemma 1 imply $1 - \lambda\theta = \lambda(1 - \theta)u_c(c) > 0$. Since $a' > 0$, we have $\zeta = 0$ and so we can rearrange (7) as

$$R^{-1}\Pi_a(a')(1-\lambda\theta)=\lambda(1-\theta)\left[u_c(c)-\beta v_a(a')\right].$$

The left-hand side is negative since, by assumption, $\Pi_a(a') < 0$ and since, as shown above, $1 - \lambda\theta > 0$. The right-hand side, however, is nonnegative by (4), $\lambda > 0$, and $\theta \in (0,1)$—a contradiction. ■

Proposition 1 shows that as long as the insurer's profits are strictly decreasing in the agent's assets for some $a' > 0$ in $\mathbb{A}$ at which the agent is not borrowing constrained, then Markov-perfect insurance contracts in which the insurer is able to specify and control agent savings (equivalently consumption) differ from Markov-perfect contracts in which the insurer is unable to do so. That is, asset contractibility matters for any asset level a satisfying the proposition conditions. Insurer's profits that monotonically decrease in the agent's assets (holding bargaining power θ constant) naturally arise—for example, if the agent's preferences exhibit decreasing absolute risk aversion. In that case, richer agents have lower demand for market insurance (they can do more smoothing with their own assets) compared with poorer agents. The borrowing constraint $a' \geq 0$ is also less likely to bind for richer agents. See Section 4.2 for an illustration.

We can gain more intuition by looking at the special cases when θ approaches its bounds. As shown in Karaivanov and Martin (2013), in the monopolistic insurer case (when $\theta \to 0$), if u is unbounded below and satisfies a mild technical condition, MPE contracts with and without asset contractibility differ and asset contractibility affects the insurer's profits. The reason is that the commitment friction creates a misalignment in the asset accumulation incentives of the contracting parties. Intuitively, the agent can use his ability to save privately to increase his outside option, since Ω is strictly increasing in a, thereby ensuring higher future transfers. This strategy counters the principal's desire, coming from profit maximization, to drive the agent toward the lower utility bound $\Omega(0)$.

As $\theta \to 1$—the case of free entry by insurers—we showed that (i) the insurer makes zero expected profits *per period* for all assets levels a and (ii) $\mu = 0$, the savings incentive-compatibility constraint (4) does not bind. In this case, since all of the surplus goes to the agent, the objectives of the two sides are perfectly aligned. And because the insurer makes zero expected profits per period, asset contractibility is irrelevant: The insurance contract is the same, regardless of whether the insurer can control the agent's savings. The result that the insurer makes zero profits per period with free entry is critical, as it does not allow the insurer to exploit his rate of return advantage when $r < R$ if assets are contractible.

4.2 Numerical Analysis

We illustrate and quantify the effects of asset contractibility in Markov-perfect insurance contracts using a numerical simulation. We adopt the parameterization we used previously (Karaivanov and Martin, 2015). Specifically, suppose $u(c) = \ln c$ and pick the following parameter values: $\beta = 0.93$, $r = 1.06$, $R = 1.07$, $y^1 = 0.1$, $y^2 = 0.3$, and $\pi^1 = \pi^2 = 0.5$. These parameters imply expected output $\bar{y} = 0.2$. For market power, we choose $\theta = 0.5$ as the benchmark and analyze the effects of varying it below.

We use the following method to compute the various cases. We begin by computing the autarky problem. We use a discrete grid of 100 points for the asset space but allow all choice variables to take any admissible value. Cubic splines are used to interpolate between grid points. The upper bound for assets $\bar{a}$ is set to 5, which ensures that the asset accumulation functions always cross the 45-degree line (i.e., the upper bound is never restrictive). Next we compute the MPE assuming $\theta = 1$ (perfect competition), since in this case asset contractibility does not matter. We use the first-order conditions of the autarky and MPE problems to compute the numerical solutions for each case. Having solved the MPE with $\theta = 1$, we use it as the starting point to compute an MPE for other assumptions on market power and asset contractibility. These problems are solved using standard value function iteration methods.

Figure 1

Consumption and Savings

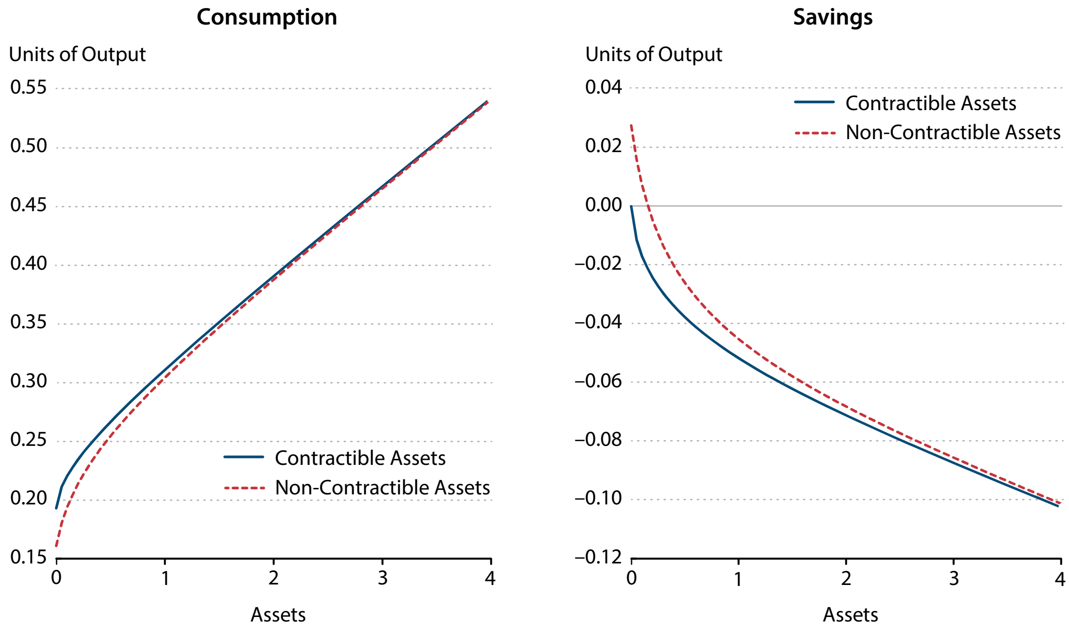

Figure 1 displays the agent's consumption c and *net savings* $a' - a$ as a function of the agent's current asset level a. The solid line corresponds to the case with contractible assets (i.e., when constraint (4) is not imposed). The dashed line corresponds to the case when the agent's choice of a' is not contractible (i.e., when constraint (4) is imposed). As shown, the agent's consumption is strictly increasing in his assets, while net savings are decreasing in assets. Allowing the savings decision to be part of the insurance contract results in higher consumption and lower savings for the agent. Intuitively, when assets are contractible, the insurer wants to push the agent's assets toward zero as this generates a lower outside option for the agent and more profits for the insurer. In addition, less assets are carried over time at the agent's rate of return r instead of the higher return R.

The long-run implications of asset contractibility are also significantly different. When the agent's assets are not contractible, if we start with an agent with some initial assets a_0 and use the computed MPE to simulate the insurance contract for infinitely many periods, then the agent's assets converge in the limit to a positive value. This is shown by the dashed line in the right panel of Figure 1, which shows that savings $a' - a$ is above zero for sufficiently low asset values and below zero for sufficiently high asset levels. In contrast, when savings are contractible, the agent's assets converge to zero in finite time, as proven in Karaivanov and Martin (2015).

Figure 2 shows the implications of asset (non-)contractibility for the agent's welfare and the insurer's profits. Agent welfare is measured as the per-period consumption equivalent

Figure 2

Welfare and Profits

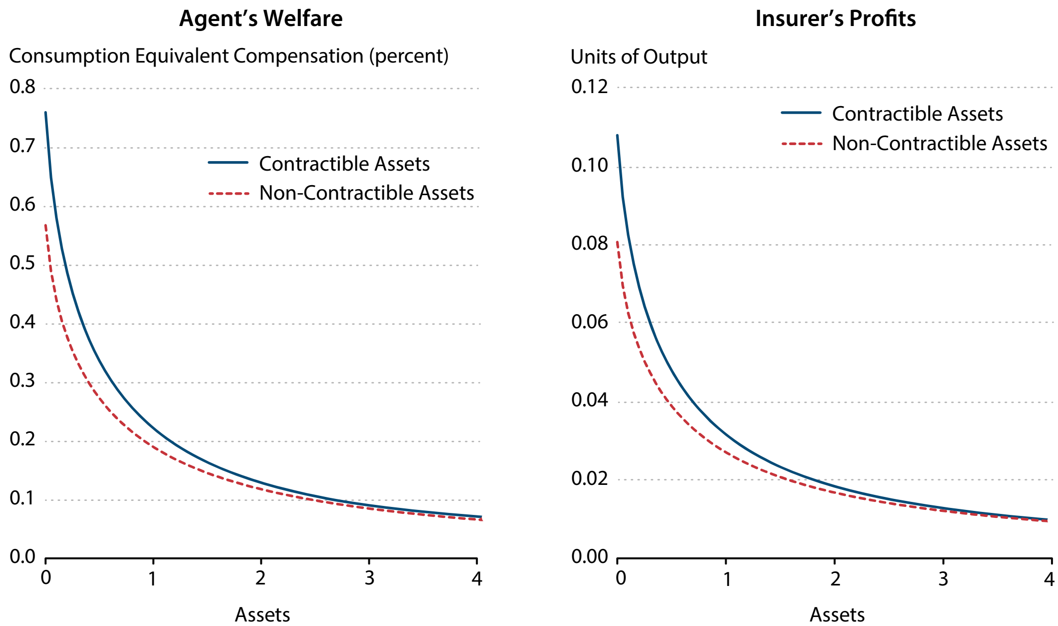

compensation the agent would require in autarky to be indifferent between remaining in autarky and accepting the insurance contract. Formally, for any $a \in \mathbb{A}$, we define the welfare gains as

$$\Delta(a) \equiv \exp\{(1-\beta)[v(a)-\Omega(a)]\}-1.$$

The insurer's profits are measured as the expected net present value $\Pi(a)$, which is expressed in output units. As shown, both the agent's welfare and the insurer's profits are strictly decreasing in the agent's assets a. This is intuitive: At lower asset levels the agent is less able to self-insure and therefore benefits more from additional insurance. That is, the surplus generated in an insurance contract, which is proportionally split between the parties, is larger when the agent's wealth is lower.

Note that the welfare gains for the agent in an MPE relative to self-insurance can be substantial: At the extreme, at zero assets (no ability to self-insure), they amount to almost 0.8 percent of consumption per period. The welfare gains are still significant at higher asset levels, converging toward 0.1 percent of autarky consumption per period, which is about the same as the estimated cost of business cycle fluctuations for the average agent (see Lucas, 1987). The welfare loss that arises if the agent's assets are non-contractible (the difference between the solid and dashed lines in Figure 2) can be large too: At zero assets, it is about 0.19 percent. This difference, however, becomes negligible at high asset levels.

Turning to the insurer's profits, we see that they are the largest when the insurer contracts with an agent with zero assets (given our log utility, this corresponds to the highest demand for insurance and the least ability to self-insure). In this case, the net present value of profits equals 54 percent and 40 percent of the expected per-period output ($\bar{y} = 0.2$) for the cases with and without contractible assets, respectively. As we can see, the ability to contract on the savings decision can also significantly boost the insurer's profits, in addition to the agent's welfare.

5 EXTENSIONS

5.1 Market Power

We now analyze how the degree of the insurer's market power affects the results. That is, how do Markov-perfect insurance contracts change when we vary the bargaining power parameter θ? The proportional surplus-splitting rule (5) directly implies that raising the agent's bargaining power θ strictly increases the agent's net surplus from market insurance, $v(a) - \Omega(a)$, relative to the insurer's present value profits $\Pi(a)$.

Using the parameterization from the previous section, we quantify the effects of market power on the agent's welfare and the insurer's profits. Figure 3 shows the consumption equivalent compensation $\Delta(a)$ and the insurer's profits at zero assets, plotted as a function of the parameter θ. Recall that higher θ can be interpreted as *lower* market power for the insurer. As the figure shows, unsurprisingly, the agent's welfare increases with his bargaining power, while the insurer's profits decrease. As we converge to a more competitive environment (higher θ), the agent's welfare increases considerably. In the extreme, at $\theta \to 1$ (perfectly competitive insurance market), the consumption equivalent compensation value of insurance in an MPE for an agent with zero wealth is about 4.5 percent of his autarky consumption per period. At the other extreme, when $\theta \to 0$ (monopolistic insurer), the profits of an insurer facing an agent with zero wealth are the largest, with a net present value about 65 percent of expected per period output.

Figure 3 also shows that both the agent and the insurer lose (in terms of welfare or profits) when the agent's assets are not contractible over the whole range $\theta \in (0,1)$. Interestingly, the agent's largest welfare loss from savings non-contractibility, equal to about 0.4 percent of autarky consumption, occurs at an interior value for the bargaining power parameter, at around $\theta = 0.8$. Remember that the agent cannot benefit from asset contractibility in the monopoly case ($\theta \to 0$) since in that case all gains from controlling the agent's assets go to the insurer. Also, as argued previously, the agent does not benefit from asset contractibility in the case of perfect competition ($\theta \to 1$) since in that case the MPEs with and without asset contractibility coincide (see Sections 3.3 and 4.1). The insurer's largest loss from asset non-contractibility occurs as $\theta \to 0$ (the monopoly case), with a magnitude slightly higher than 14 percent of expected per-period output.

5.2 The Rate of Return R

We next analyze the effects of varying the insurer's intertemporal rate of return R. Increasing R is equivalent to decreasing the factor by which the insurer discounts future

Figure 3

Market Power

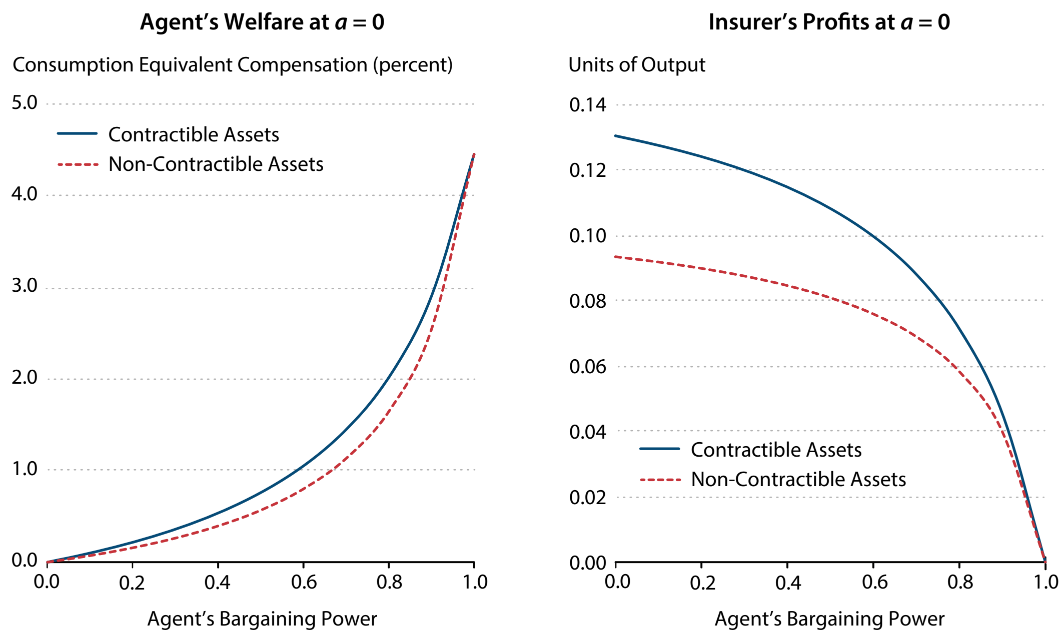

profits—that is, making the insurer more impatient. Note that there is no direct productivity effect of varying R as the agent's output technology—and hence total resources—are independent of R. In addition, the agent's autarky problem (1) remains the same.

Figure 4 plots the agent's welfare gains in an MPE relative to self-insurance, as measured by $\Delta(a)$, and the insurer's present value profits $\Pi(a)$ as R varies over its full range, from $R = r = 1.06$ to $R = 1/\beta \approx 1.075$. All other parameters, including the bargaining power θ, are held fixed at their respective benchmark values. In the interest of providing the clearest intuition for the results, we focus on the case of zero assets, $a = 0$. All other asset levels provide a similar qualitative picture (details are available upon request).

Two main results are evident from Figure 4. First, both the agent's welfare gains relative to autarky and the present value of the insurer's profits are strictly decreasing in R. The intuition for this result is found by examining the direct effect of varying R on the agent's and insurer's surplus in the contract. If the decision variables c and a' were held fixed, the agent's surplus, $u(c) + \beta v(a') - \Omega(a)$, would be constant in R, while the insurer's surplus, $\bar{y} + ra - c - a' + R^{-1}\Pi(a')$, would be strictly decreasing in R. At $a = 0$, when assets are contractible $a' = 0$ and thus, when R increases, the only way to satisfy the proportional surplus-splitting constraint (5) is to decrease the agent's consumption. When $a > 0$, savings decisions do vary with R and, hence, there are further effects on welfare and profits.[4] Our numerical simulations show that, for the chosen parameters, the overall effect still moves in the same direction as when the agent has zero assets. The difference in welfare gains as R varies can be substantial. For example,

Figure 4

Rates of Return

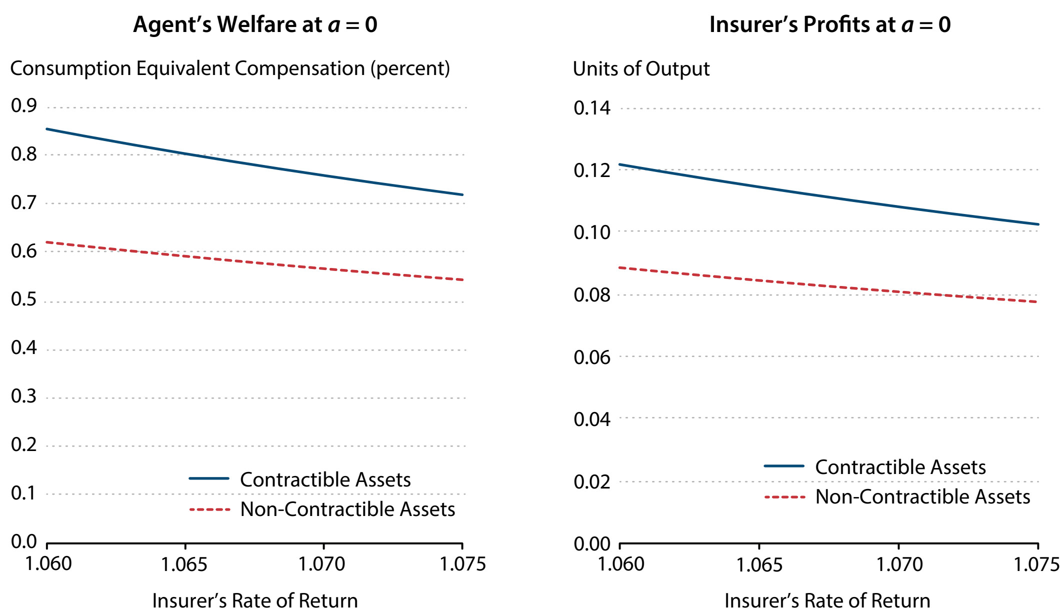

at zero assets, moving from $R \approx 1/\beta$ to $R = r$ results in a welfare increase for the agent equivalent to 0.14 percent of his autarky consumption per period.

Second, Figure 4 shows that our results on the effects of asset (non-)contractibility continue to hold for all admissible values of R. Making assets contractible increases both the agent's welfare and the insurer's profits (compare the dashed lines with the solid lines). Quantitatively, at zero assets, the welfare gains from making the agent's assets contractible are the highest at $R = r = 1.06$ and are equivalent to 0.23 percent of autarky consumption per period, compared with 0.19 percent at the benchmark value of $R = 1.07$ or 0.18 percent at $R \approx 1/\beta$.

6 CONCLUSION

We study the role of assets contractibility, market power, and the rate of return differential between insured and insurers in a dynamic risk-sharing setting with a limited commitment friction. We find significant welfare effects along all three dimensions. Potential lessons from our analysis with relevance for actual insurance markets with commitment frictions similar to those we model indicate the desirability of increased competition, extending the ability to condition insurance terms on both the current assets and the savings of the insured, as well as mitigating the possibility of a large return on assets differentials between insurance providers and households or firms. ■

NOTES

[1] Throughout the article, we use subscripts to denote partial derivatives and primes for next-period values.

[2] That is, the agent would not save if output were constant over time.

[3] Generically, an interior solution for asset choice implies $a' > 0$. However, it is possible to have an interior solution, where $a' = 0$ and where the nonnegativity constraint, although satisfied with equality, does not bind. In either case, $\zeta = 0$.

[4] In particular, the agent would prefer to contract with an insurer whose intertemporal rate of return R is closer to the agent's rate of return r as this mitigates the distortion in the time profiles of consumption and savings arising from the commitment friction (see Karaivanov and Martin, 2015, Section 3.2 for additional details).

REFERENCES

Aiyagari, S. Rao. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics*, August 1994, *109*(3), pp. 659-84; http://dx.doi.org/10.2307/2118417.

Allen, Franklin. "Repeated Principal-Agent Relationships with Lending and Borrowing." *Economics Letters*, 1985, *17*(1-2), pp. 27-31; http://dx.doi.org/10.1016/0165-1765(85)90121-1.

Arnott, Richard and Stiglitz, Joseph E. "Moral Hazard and Nonmarket Institutions: Dysfunctional Crowding Out of Peer Monitoring?" *American Economic Review*, March 1991, *81*(1), pp. 179-90.

Bester, Helmut and Strausz, Roland. "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case." *Econometrica*, July 2001, *69*(4), pp. 1077-98; http://dx.doi.org/10.1111/1468-0262.00231.

Bewley, Truman F. "The Permanent Income Hypothesis: A Theoretical Formulation." *Journal of Economic Theory*, December 1977, *16*(2), pp. 252-92; http://dx.doi.org/10.1016/0022-0531(77)90009-6.

Cole, Harold L. and Kocherlakota, Narayana R. "Efficient Allocations with Hidden Income and Hidden Storage." *Review of Economic Studies*, July 2001, *68*(3), pp. 523-42; http://dx.doi.org/10.1111/1467-937X.00179.

Doepke, Matthias and Townsend, Robert M. "Dynamic Mechanism Design with Hidden Income and Hidden Actions." *Journal of Economic Theory*, January 2006, *126*(1), pp. 235-85; http://dx.doi.org/10.1016/j.jet.2004.07.008.

Fernandes, Ana and Phelan, Christopher. "A Recursive Formulation for Repeated Agency with History Dependence." *Journal of Economic Theory*, April 2000, *91*(2), pp. 223-47; http://dx.doi.org/10.1006/jeth.1999.2619.

Kalai, Ehud. "Proportional Solutions to Bargaining Situations: Interpersonal Utility Comparisons." *Econometrica*, October 1977, *45*(7), pp. 1623-30; http://dx.doi.org/10.2307/1913954.

Karaivanov, Alexander K. and Martin, Fernando M. "Dynamic Optimal Insurance and Lack of Commitment." Working Paper No. 2011-029, Federal Reserve Bank of St. Louis, October 2011, revised July 2013; https://research.stlouisfed.org/wp/more/2011-029.

Karaivanov, Alexander K. and Martin, Fernando M. "Dynamic Optimal Insurance and Lack of Commitment." *Review of Economic Dynamics*, April 2015, *18*(2), pp. 287-305; http://dx.doi.org/10.1016/j.red.2014.05.001.

Lucas, Robert E. Jr. *Models of Business Cycles*. Oxford, England: Basil Blackwell, 1987.

Maskin, Eric and Tirole, Jean. "Markov Perfect Equilibrium: I. Observable Actions." *Journal of Economic Theory*, October 2001, *100*(2), pp. 191-219; http://dx.doi.org/10.1006/jeth.2000.2785.

Schechtman, Jack and Escudero, Vera L.S. "Some Results on an 'Infinite Fluctuation Problem.'" *Journal of Economic Theory*, December 1977, *16*(2), pp. 151-66; http://dx.doi.org/10.1016/0022-0531(77)90003-5.

Stokey, Nancy L.; Lucas, Robert E. Jr. and Prescott, Edward C. *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press, 1989.

Student Loans Under the Risk of Youth Unemployment

Alexander Monge-Naranjo

While most college graduates eventually find jobs that match their qualifications, the possibility of long spells of unemployment and/or underemployment—combined with ensuing difficulties in repaying student loans—may limit and even dissuade productive investments in human capital. The author explores the optimal design of student loans when young college graduates can be unemployed and reaches three main conclusions. First, the optimal student loan program must incorporate an unemployment compensation mechanism as a key element, even if unemployment probabilities are endogenous and subject to moral hazard. Second, despite the presence of moral hazard, a well-designed student loan program can deliver efficient levels of investments. Dispersion in consumption should be introduced so the labor market potential of any individual, regardless of the family's financial background, is not impaired as long as the individual is willing to put forth the effort, both during school and afterward, when seeking a job. Third, the amounts of unemployment benefits and the debt repayment schedule should be adjusted with the length of the unemployment spell. As unemployment persists, benefits should decline and repayments should increase to provide the right incentives for young college graduates to seek employment. (JEL D82, D86, I22, I26, I28, J65)

Federal Reserve Bank of St. Louis *Review*, Second Quarter 2016, *98*(2), pp. 129-58.
http://dx.doi.org/10.20955/r.2016.129-158

Many college graduates may face spells of unemployment and/or underemployment before they find jobs that match their qualifications. These spells may be long, especially for some college majors, and can lead to serious financial difficulties, including obtaining credit and repaying student loans and other forms of debt. Aside from their direct welfare costs, the hardship and volatility during the early stages of labor market participation can impair—and even dissuade altogether—productive investments in human capital, especially for those from more modest family backgrounds.

In this article, I explore the optimal design of student loan programs in an environment in which younger individuals, fresh out of college, may face a substantial risk of unemploy-

Alexander Monge-Naranjo is a research officer and economist at the Federal Reserve Bank of St. Louis. The author thanks Carlos Garriga, Lance Lochner, Guillaume Vandenbroucke, and Stephen Williamson for useful comments and suggestions. Faisal Sohail provided excellent research assistance.

ment. In my model, the risk of unemployment is endogenous and subject to incentive problems. In particular, I assume that a problem of "moral hazard" (hidden action) distorts the implementation of credit contracts. More specifically, I examine an environment in which costly and unverifiable effort determines the probability of younger workers finding a job. The costs of unemployment for a young worker are in terms of zero (or very low) earnings and missing opportunities to gain experience that would enhance his or her labor earnings for subsequent periods. Moral hazard and other incentive problems have been studied extensively by economists in a wide array of areas ranging from banking and insurance to labor markets. Yet, only recently has the explicit consideration of incentive problems been introduced in the study of optimal student loan programs.[1] Despite the extensive literature on unemployment insurance since the 1990s (e.g., Wang and Williamson, 1996, and Hopenhayn and Nicolini, 1997), the integration of an unemployment insurance scheme within the repayment structure of student loans and the optimal design of such a scheme is an aspect that remains unexplored.

In this article, I first consider a simple three-period environment. In the first period, a young person decides on his or her level of schooling investment. In the second period, a hidden effort governs the probability of unemployment. In the third period, all workers find employment, but their earnings are affected by their schooling level and their previous employment. I contrast the resulting allocations from two contractual arrangements: the first-best (i.e., unrestricted efficient) allocations and the optimal student loan programs when effort is a hidden action (moral hazard). I then extend the simple environment by dividing the potential postcollege unemployment spell into multiple subperiods. I use this extension to examine the optimal design of unemployment insurance and compare the human capital investments resulting from a suboptimal scheme without unemployment insurance. In all these cases, I restrict the credit arrangement so the creditor expects to break even in expectation (i.e., in average over all possible future outcomes). Therefore, my conclusions can apply not only to government-run programs, but also, under similar enforcement conditions, to privately run student loan programs.

I derive three main conclusions. First, the optimal student loan program must incorporate, as a key element, a transfer mechanism should college graduates face post-schooling unemployment. This conclusion holds even if unemployment probabilities are endogenous and job searching might be subject to moral hazard. This simple and perhaps not surprising result is worth highlighting given the limited scope for insurance in existing student loans. An unemployment insurance mechanism not only alleviates the welfare cost of potentially catastrophic low consumption for the unemployed, but can also help to enhance human capital formation as individuals and their families would not need to self-insure by means of lower-return assets and reduced schooling.

Second, and related to the last point, despite the presence of moral hazard, a well-designed student loan program can deliver efficient levels of investments for at least a segment of the population. Here, dispersion in consumption should be introduced so the labor market potential of any individual, regardless of family financial background, is not impaired. This result is conditional on the individual's willingness to exert effort, which might be subject to wealth effects. However, once the effort and abilities of a person are factored in, the investments in schooling should be completely independent of one's family's wealth.

A third important result concerns the dynamics of the unemployment benefits and the repayment of debts. Once I consider a model with possible multiperiod unemployment and repeated search effort, the unemployment benefits should decline with the length of the unemployment spell. Moreover, the debt balance and its repayment should also be adjusted upward the longer a person stays unemployed. While these two features are well understood in the literature on unemployment insurance, they are not incorporated in actual student loan programs. I believe this is an interesting margin to explore: By enhancing the ability to provide both insurance and incentives. it also can enhance the formation of human capital, especially for those individuals with high ability but low family resources.

In the next section, I examine data for recent cohorts of U.S. college graduates and show that unemployment and underemployment are significant risks for them right after college. In Section 2, I describe the basic environment for analysis; in Sections 3 and 4 I characterize the allocations under the first-best and under optimal loan programs under moral hazard. Section 5 solves the optimal repayment in the multiperiod environment and discusses the allocations. Section 6 concludes. The appendix discusses additional aspects of the optimal student loan and compares it with other contractual environments.

1 POSTCOLLEGE UNEMPLOYMENT AND UNDEREMPLOYMENT

Recent work on college education choices has called attention to the rather high risk involved in investments in education. For instance, Chatterjee and Ionescu (2012) highlight the fact that a sizable fraction of college students fail to graduate. Furthermore, as emphasized by Lee, Lee, and Shin (2014), even successful graduates face a large and widening dispersion in labor market outcomes, possibly including the option of working in jobs and occupations that do not require their college training. Thus, even if a college education might greatly enhance the set of labor market opportunities, such an education is a risky investment that comes at the cost of tuition and forgone earnings; also, graduates' ex post returns may even render repayment of student loans difficult.

To be sure, some of these risks and volatilities are more prevalent at the beginning of a person's labor market experience. A college education does not fully preclude a younger, unexperienced worker from facing more difficulties in finding a job than an older, more mature, experienced, and better-connected worker. To illustrate this point, I use 2011 cross-sectional earnings and unemployment data from the American Community Survey (ACS) to report the unemployment and earnings of college graduates (Figures 1 and 2).[2] In both figures, the blue columns correspond to the average recent college graduate (between 22 and 26 years of age), while the red columns represent the average more experienced graduate (between 30 and 54 years of age). In both figures, graduates from more than 170 majors are grouped into 13 broader areas.

Figure 1 shows that the unemployment rates are uniformly higher for recent graduates than for more experienced ones. The differences in the rates are very pronounced for some

Figure 1

2011 Unemployment Rates for Recent and Experienced College Graduates by Major

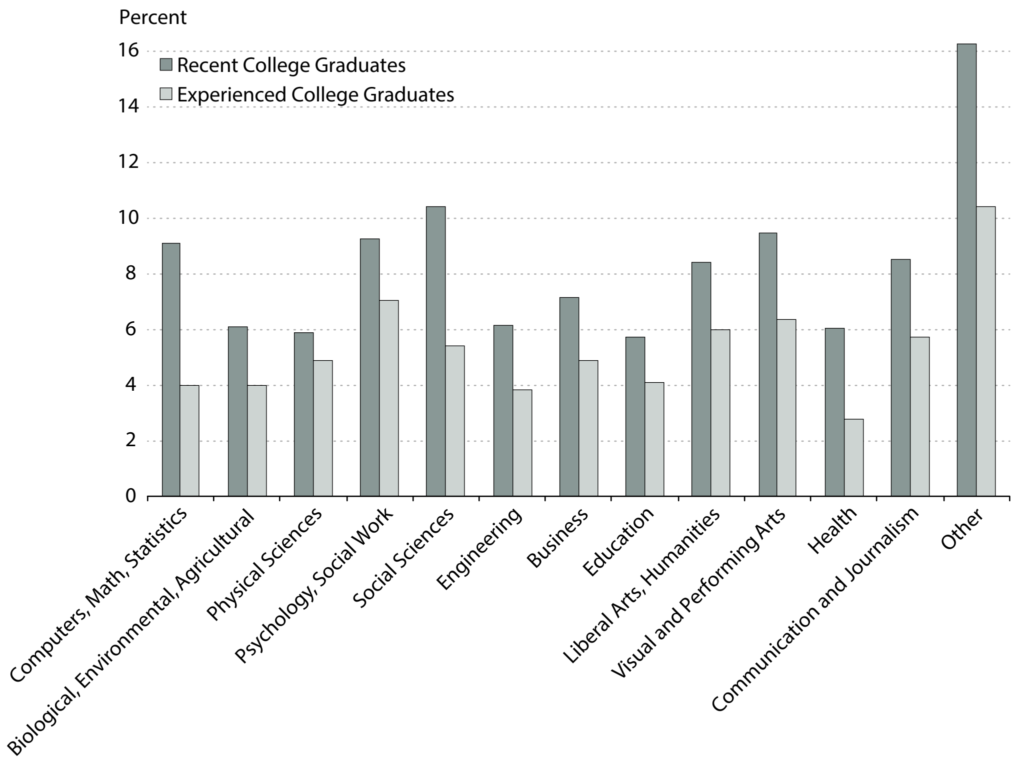

SOURCE: American Community Survey.

groups of majors: as much as 5 percentage points higher for fields such as computer science, math and statistics, social sciences, and others. The rates are much lower for other fields such as education, business, and, especially, physical sciences. However, the unemployment gaps are significant in all groups of majors, supporting the notion that graduates in all fields take time to find jobs. In the meantime, they experience higher rates of unemployment than their more established peers.

Figure 2 shows the 2011 labor earnings for the same groups of graduates who are employed by major. This figure also shows a very clear pattern: More recent graduates earn significantly less than more experienced graduates. In fact, for all but three majors (education, health, and other), recent graduates earn less than half that of their more experienced peers. Indeed, recent graduates earn as little as 41 percent as much as their older peers in biological sciences and liberal arts and humanities; in education that ratio is the highest among all majors, at 60 percent.

In sum, a simple look at the cross-sectional data from the ACS clearly indicates that within each field younger graduates (i) have more difficulty finding a job than more experienced ones and (ii) their earnings are lower when they are employed. But while the ACS makes it easy to compare different cohorts of college graduates, it does not follow them over time. The ACS

Figure 2

2011 Labor Earnings for Recent and Experienced College Graduates by Major

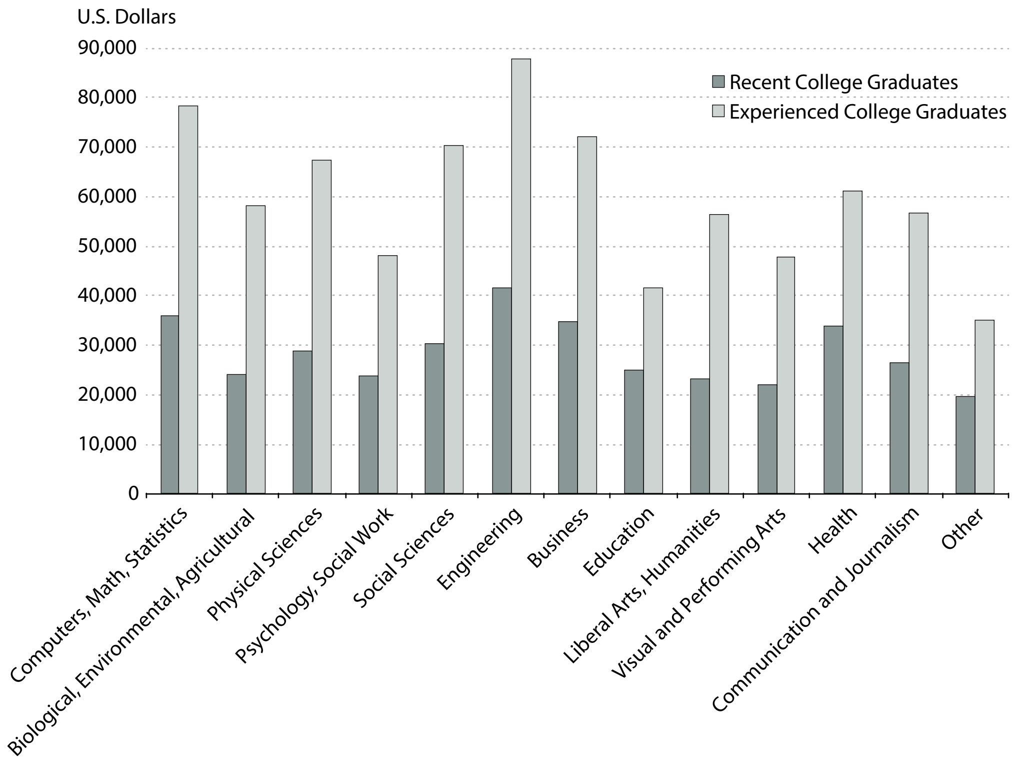

SOURCE: American Community Survey.

data cannot establish the transitions of college graduates from the early periods of their labor market experience to the more mature ones in terms of employment, earnings, and repayment of their student loans. To this end, I now consider the Baccalaureate and Beyond Longitudinal Survey 2008-12 (B&B:08/12) of college students who graduated in the 2007-08 academic year.[3] The survey collects the employment records of individuals in 2009 and in 2012, about one and four years after graduation, respectively. It follows just one cohort of graduates, so comparison across cohorts cannot be made with this dataset.

As in Lochner and Monge-Naranjo (2015a)—but for the B&B:93/03 survey—I aim to report unemployment and underemployment for a typical American college student. In what follows, I exclude noncitizens, the disabled, and individuals who received their baccalaureate degree at age 30 or older as their labor market experience involves a number of other issues. For the same reason, I also exclude those with more than 12 months of graduate work. Tables 1 and 2 document that unemployment and underemployment are very relevant risks for recent U.S. college graduates, especially in their first few years following graduation. Table 1 shows the average percentage of the months in which students remained unemployed since graduation (i.e., Number of months unemployed/Numbers of months since graduation × 100). The

Table 1

Mean Percent of Time Unemployed Since Graduation (July 2008)

College major	2009	2012
Business	10.0	6.1
Education	10.1	6.7
Engineering	6.9	4.2
Health professions	6.5	3.1
Public affairs	8.8	5.6
Biological sciences	9.3	6.0
Math/science/computer science	6.6	4.1
Social science	10.7	8.9
History	12.9	7.5
Humanities	12.0	9.1
Psychology	9.3	7.0
Other	10.9	7.1
All	9.8	6.6

SOURCE: Baccalaureate and Beyond, 93/03.

Table 2

Percent of Graduates with Primary Job Unrelated to College Education

College major	2009	2012
Business	19	14
Education	12	9
Engineering	9	11
Health professions	10	9
Public affairs	21	19
Biological sciences	20	23
Math/science/computer science	18	10
Social science	33	30
History	45	36
Humanities	43	29
Psychology	29	27
Other	29	23
All	24	19

SOURCE: Baccalaureate and Beyond, 93/03.

first column reports the percentage up until 2009 and the second column the average percentage three years later, in 2012. Each row reports the average by college major and for the whole sample.

The results show fairly high unemployment rates. On average, one of every 10 college graduates remains unemployed in the first year after graduation. Of course, 2009 is not a typical year since the United States was in the middle of the so-called Great Recession and the overall unemployment rate was high.[4] Moreover, the unemployment rate for this sample of college-educated individuals, on average, is much lower than for the rest of U.S. workers. Note also that there is significant dispersion. While health professionals, math/science, and computer science professionals all had unemployment rates lower than 7 percent, most others were closer to 10 percent. In the extreme, history majors found themselves unemployed 13 percent of the time—that is, almost one of every seven.

The other salient result is the rapid decline in this measure of unemployment three years later. The overall unemployment rate falls by one-third, from 9.83 percent to 6.55 percent. These employment gains occur across all majors, with remarkable gains in history, business, and education. With just four years of labor market experience, the unemployment rate for this young cohort compared favorably with the overall U.S. civilian unemployment rate: 8.2 percent in July 2012.

Table 2 reports the other form of labor market unemployment: the possibility of employment that does not use the person's main skills. From the B&B:08/12, I obtain the fraction of individuals reported as employed at the time of the survey but whose job or occupation is not directly related to the person's college education.

Table 3

Repayment Status of Borrowers Graduating in 1992-93

Status	1998	2003
Fully repaid (%)	26.9	63.9
Repaying (%)	65.1	27.8
Deferment/forbearance (%)	3.8	2.5
Default (%)	4.2	5.8

SOURCE: Lochner and Monge-Naranjo (2015a).

Table 2 also shows some remarkable results. For the first year after graduation, one of every four employed college graduates ends up working in a job unrelated to his or her education. For some majors such as social sciences, humanities, and especially history, the ratios are much higher. After four years, the ratios are lower but still high, around one of every five. With the exception of biological sciences and engineering, the ratios decline for all other majors; in some cases such as business, history, humanities, and especially math and computer science, the ratios decline substantially.

In sum, Tables 1 and 2 support the view hinted at by the ACS data that it not only may take time for a recent college graduate to find a job, but also may take an even longer time to find a job matching his or her acquired skills, abilities, and vocations.

The early postcollege stages are also associated with higher difficulties in repaying student loans. For a better perspective on the life cycle of payments for student loans, Lochner and Monge-Naranjo (2015a) examine the repayment patterns of an older cohort of borrowers: those in the B&B:93/03, the cohort of students who graduated in the 1992/93 academic year. Table 3 (from Lochner and Monge-Naranjo, 2015a) reports repayment status for borrowers as of 1998 and 2003—around 5 and 10 years after graduation. In both years, graduates repaying their loans plus those who had already fully repaid their loans account for 92 percent of the borrowers. Not surprisingly, the fraction of those with fully repaid loans is much higher 10 years after graduation.

More interestingly, the fraction of borrowers who applied for and received a deferment or a forbearance (postponement of repayment without default) was significantly higher in the early years after graduation. In 1998, this fraction accounted for 3.8 percent of borrowers. Five years later, in 2003, the percentage fell to 2.5 percent. These figures suggest that deferment and forbearance are important forms of non-repayment. The declining share of borrowers engaging in this form of non-repayment may be a reflection of lower volatility in the labor market of graduates as times passes, but it also may reflect the fact that fewer borrowers can quality for deferment and forbearance as they age. Indeed, the counterpart is that default rates rise from 4.2 percent to 5.8 percent between 1998 and 2003.[5]

Table 4

Repayment Status Transition Probabilities

	Repayment status in 2003		
Repayment status in 1998	**Repaying/fully paid**	**Deferment/forbearance**	**Default**
Repaying/fully paid (%)	93.9	2.0	4.0
Deferment/forbearance (%)	74.9	16.5	8.5
Default (%)	54.4	3.8	41.8

SOURCE: Lochner and Monge-Naranjo (2015a).

Table 4 (from Lochner and Monge-Naranjo, 2015a) shows the transition rates for different repayment states from 1998 to 2003. The rows in the table list the probabilities of (i) being in repayment (including those whose loans are fully repaid), (ii) receiving a deferment or forbearance, or (iii) being in default 10 years after graduation in 2003 conditional on each of these repayment states five years earlier (in 1998). Note that most (94 percent) of those in good standing five years after graduation are also in good standing 10 years after graduation. Also, most (75 percent) of those in deferment or forbearance in 1998 transition to good standing five years later. More than half (54 percent) of those in default in 1998 transition to good standing five years later. The general pattern indicates that repayment is more difficult early on, but many who face hardship repaying and even declare default eventually move to good standing. And once a college graduate is in good standing, there is a strong tendency for him or her to remain in that state.

A few words of caution are in order. These findings do not indicate that the risks are irrelevant because they are not necessarily permanent. Temporary and transition costs can be high for borrowers who may respond by underinvesting in their education. Moreover, the low persistence (and eventual reduction) in the fraction in deferment/forbearance may not be driven by younger workers finding a good job but instead because those mechanisms are designed only to temporarily help borrowers early on, and older borrowers cannot typically qualify for a deferment or forbearance. Supporting this view is the fact that the default rate is higher 10 years after graduation.

To summarize, this section reviews consistent evidence that new college graduates have a fairly higher incidence of unemployment and underemployment and lower earnings, relative to both contemporaneous older cohorts (from ACS data) and their own future (from B&B data). These findings are valid for all majors but to different degrees.[6] The section also provides evidence that new graduates seem to encounter more difficulties repaying their loans and that existing insurance devices such as deferments and forbearances are more widely used during the early postcollege years. These empirical patterns motivate the simple question of this article: What should be the optimal design of student loan programs given the risk of unemployment for recent graduates?

2 A SIMPLE MODEL OF UNEMPLOYMENT FOR RECENT GRADUATES

I analyze schooling investments in the presence of youth unemployment in a very stylized and tractable environment. In such an environment, I consider the best feasible arrangement in the presence of moral hazard and compare it with the best possible arrangement when incentive problems can be ruled out—that is, the so-called first-best allocations.

2.1 The Environment

Consider an individual whose life is divided into three periods or stages: early youth, $t = 0$; youth, $t = 1$; and maturity, $t = 2$. Schooling takes place in early youth, whereas labor market activity takes places during the youth and maturity periods. Preferences are standard. As of $t = 0$, each person's utility is driven by consumption in all periods $\{c_t\}_{t=0}^{2}$ and effort during youth e while forming human capital in school:

$$U = u(c_0) - v(e) + \beta E\left[u(c_1) + \beta u(c_2)\right],$$

where $0 < \beta < 1$ is a discount factor and $u(\cdot)$ is a standard increasing, twice continuously differentiable, strictly concave utility function for consumption streams c_t during $t = 0, 1, 2$. I assume that the utility has a constant relative risk aversion:

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma},$$

where $\sigma > 0$. For most of the analysis, I assume $\sigma < 1$. For brevity, in some formulas I retain the use of $u(\cdot)$ and its derivative $u'(c)$. The expectation $E[\cdot]$ is over the possible employment outcomes.

To remain focused solely on youth unemployment, I consider an environment in which the only risk is whether a college graduate finds a job right after school (i.e., at $t = 1$) or has to wait until maturity ($t = 2$). For the stylized model, the exertion of effort e is relevant only for the first period at a disutility cost $v(\cdot)$, an increasing function. The probability of finding a job during the youth period—and thus avoiding unemployment altogether—is an increasing function $p(\cdot)$ of the effort e exerted by the person at the end of period $t = 0$.[7] Naturally, the function $p(\cdot)$ is bounded between (0,1); moreover, when effort e is treated as a continuous variable, $p(\cdot)$ is a strictly concave function.

The exertion of effort is a central aspect of my analysis. I consider the two leading assumptions in the literature of optimal contracts under moral hazard. In the first case, effort is a continuous variable $e \in [0,\infty)$ and the higher the effort, the higher the probability of finding a job right after school. In the second case, effort is binary, $e \in \{0,1\}$; that is, a person either works hard or not at all. If the former, the probability of finding a job is $0 < p_H < 1$. If the latter, the probability is $0 < p_L < p_H$. Here, the subscripts L and H stand for low effort and high effort, respectively.

In the first period, $t = 0$, the key investment is schooling. It is denoted by $h \geq 0$ and measured in consumption goods. Aside from h, the effort e must also be considered investments during early youth. Investments h enhance the person's ability a to generate output once employed during both postcollege periods, $t = 1$ and $t = 2$, while e is an investment in increasing the probability of realizing the young person's earnings potential. The initial ability a is given as of $t = 0$ and summarizes innate talents and acquired skills during the person's earlier years of childhood. Given ability a and investments h, the labor market income profile of the person is described as follows: For $t = 1$, the person can be employed or unemployed:

$$y_1 = \begin{cases} ah^{\alpha} & \text{if employed at } t = 1 \\ b & \text{otherwise,} \end{cases}$$

and

$$y_2 = \begin{cases} Gah^{\alpha} & \text{if employed at } t = 1 \\ ah^{\alpha} & \text{otherwise,} \end{cases}$$

for $t = 2$. In this setting, I assume that mature workers are always employed, so the only possible incidence of unemployment is for the young. Here $G \geq 1$ allows for the possibility of on-the-job training/learning by doing during period $t = 1$, thereby increasing the earnings for the second period. Also, $0 < \alpha < 1$, indicating decreasing returns to investment in education. Finally, $b \geq 0$ is a possibly nonzero unemployment minimum consumption level. I allow $b > 0$ only to ensure that the problem is well defined for complete financial autarky when $\sigma > 1$.

All individuals are initially endowed with positive resources, $W > 0$. These resources can be used at $t = 0$ to consume, invest in human capital h, or save for future consumption. This initial wealth can be seen as a family transfer, but for concreteness, here I assume that it cannot be made contingent on the employment regardless of whether the person finds a job at $t = 1$ or only at $t = 2$.

Finally, for simplicity, I assume that the relevant interest rate (i.e., the cost for lenders of providing credit) and the rate of return to savings are both equal to the discount rate. That is, the implicit interest rate is given by $r = \beta^{-1} - 1$.

3 THE FIRST-BEST

I now consider the extreme when markets are complete and allocations are the unrestricted optima. In the case of complete markets, any intertemporal and interstate (insurance) exchange can be performed. Unrestricted allocations indicate that any possible incentive problems can be directly handled and do not distort the intertemporal and insurance exchanges. Such allocations are useful as the efficient benchmark for the allocations from any other market arrangements and government policies.

In my environment, it suffices that the asset structure is one in which an individual's loan repayment is fully contingent on the labor market outcomes at $t = 1$. Let d_e and d_u denote the amount of resources that an individual at $t = 0$ commits to pay to lenders, which are assumed to be risk neutral and competitive. The amount d_e is the repayment due if a graduate finds a job right after school (hence the subscript e), an event that happens with probability $p(e)$. The amount d_u is the repayment the borrower must deliver (or receive) when $d_u < 0$ if he or she is unemployed after school, an event that happens with probability $1 - p(e)$. The borrower's effort is fully known (and controlled) by lenders, as are the probabilities for the two outcomes. As borrowers are risk neutral, they are willing to give the borrower $\beta p(e)d_e$ at $t = 0$ in exchange for a repayment d_e at $t = 1$ if the borrower finds a job right away. Similarly, a borrower would receive (or pay) an amount $\beta(1-p(e))d_u$ at $t = 0$ in exchange for d_u units at $t = 1$ if unemployed.

The sequential budget constraints for the individual are as follows: For $t = 0$, the first-period consumption c_0 and the investments in human capital h are financed by either initial wealth W or borrowing:

$$c_0 + h = W + \beta\left[p(e)d_e + \left(1 - p(e)\right)d_u\right]. \tag{1}$$

For periods $t = 1$ and $t = 2$, the constraints for the consumption levels of young employed ($c_{e,1}$ and $c_{e,2}$) and young unemployed individuals ($c_{u,1}$ and $c_{u,2}$) are given by the present value constraints:

$$c_{e,1} + \beta c_{e,2} = ah^{\alpha} - d_e + \beta Gah^{\alpha} \tag{2}$$

and

$$c_{u,1} + \beta c_{u,2} = -d_u + \beta ah^{\alpha}. \tag{3}$$

Then, the optimization problem for the young person at $t = 0$ is choosing the investments in human capital h; the exertion of effort in successfully finding a job right after school e; the financial decisions d_e and d_u; the initial consumption c_0; and the contingent plan $c_{e,1}$, $c_{e,2}$, $c_{u,1}$, and $c_{u,2}$ to maximize

$$U(a,W) = \max_{\vec{c},d,h,e} u\left(c_0\right) - v(e) + \beta p(e)\left[u\left(c_{e,1}\right) + u\left(c_{e,2}\right)\right] + \left(1 - p(e)p(e)\left[u\left(c_{u,1}\right) + u\left(c_{u,2}\right)\right]\right),$$

subject to constraints(1), (2), and (3). Here $\vec{c}$ is the vector of all date-state consumption.

The solution to this problem is very familiar. First, it is straightforward from their first-order conditions (FOCs) that $c_{e,1}$ and $c_{e,2}$ must be equal to each other; the same applies for $c_{u,1}$ and $c_{u,2}$. Then, from (2) $c_{e,2} = c_{e,1} = (ah^{\alpha}(1 + \beta G) - d_e)/(1 + \beta)$ and from (3) $c_{u,1} = c_{u,2} = (ah^{\alpha}(1 + \beta G) - d_e)/(1 + \beta)$. Second, from the FOCs of d_e and d_u, it is easy to show that early consumption c_0 must be equal to all consumptions—that is, $c_0 = c_{e,2} = c_{e,1} = c_{u,1} = c_{u,2} = c$ for some consumption level c to be determined by the individual's wealth W, ability a, and actions h and e. As expected, complete markets lead to perfect consumption smoothing over time and across states of the world, constraints (1), (2), and (3) reduce, after some trivial simplification, to the present value constraint:

$$c\left(1 + \beta + \beta^2\right) = W - h + \beta\left[\beta + p(e)\left(1 + \beta(G - 1)\right)\right]ah^{\alpha}.$$

Solving this expression for the consumption level c, the optimal schooling investment problem for a young person is simplified to

$$U(a,W)=\max_{e,h}\left(1+\beta+\beta^2\right)u\left(\frac{W-h+\beta\left[\beta+p(e)\left(1+\beta(G-1)\right)\right]ah^\alpha}{\left(1+\beta+\beta^2\right)}\right)-v(e).$$

Given any e, schooling investments should maximize $W-h+\beta[\beta+p(e)(1+\beta(G-1))]ah^\alpha$, the person's lifetime earnings net of schooling costs. The FOC for this implies an optimal investment function:

$$h^*(a;e)=\left\{a\alpha\beta\left[\beta+p(e)\left(1+\beta(G-1)\right)\right]\right\}^{\frac{1}{1-\alpha}}. \tag{4}$$

Some simple but very important implications are evident for schooling. First, conditional on effort e, investments are independent of the individual's wealth W. This simple result is the basis for a large empirical literature on credit constraints and education.[8] Second, investments are always increasing in the person's ability a, exerted effort e, the probability of finding a job right after school $p(\cdot)$, and the gains from experience G. Efficient investments are driven by expected returns as risks are efficiently insured by lenders.

With the efficient investment (4), the maximized lifetime present value of resources for the young person is $W+(1/\alpha-1)h^*(a;e)$. In the continuous case, the optimal exertion of effort would be given by the condition[9]

$$\left(\frac{W+\left(1/\alpha-1\right)h^*(a;e)}{\left(1+\beta+\beta^2\right)}\right)^{-\sigma}h^*(a;e)\times\frac{p'(e)}{p(e)}=\alpha\left[\frac{\beta}{p(e)\left[1+\beta(G-1)\right]}+1\right]v'(e),$$

where I have already used $u'(c)=c^{-\sigma}$.

A wealth effect implies that the optimal exertion of effort decreases with initial wealth W. With respect to ability, the implied relationship is more complex. If the wealth effect is weaker than the substitution effect—that is, if $\sigma<1$—then the relationship between ability and effort is always positive. However, depending on the level of wealth W, if $\sigma>1$, a negative relationship would hold.[10]

The case of binary effort can be substantially simpler. For any ability a, there is always a threshold level $\bar{W}(a)$ for which wealth levels W above the threshold imply zero effort, $e=0$, since consumption is so plentiful that the marginal gains of exerting effort are low compared with the utility cost of $e=1$. For wealth levels below $\bar{W}(a)$ the optimal exertion of effort is $e=1$. The attained utility is given by

$$U^*(a,W)=\begin{cases}\left(1+\beta+\beta^2\right)u\left(\frac{W+\left(1/\alpha-1\right)h^*(a;1)}{\left(1+\beta+\beta^2\right)}\right)-v(1) & \text{if } W\leq\bar{W}(a),\\ \left(1+\beta+\beta^2\right)u\left(\frac{W+\left(1/\alpha-1\right)h^*(a;0)}{\left(1+\beta+\beta^2\right)}\right)-v(0) & \text{otherwise.}\end{cases}$$

That is, whether $\bar{W}(a)$ is increasing or decreasing is governed by whether or not $\sigma<1$.

3.1 Discussion: First-Best Allocations and Repayment

Except for the determination of effort, the first-best allocation is familiar to most readers. It is worthwhile, however, to reexamine its implications for consumption and investment, especially for borrowing and repayment of credit at $t = 0$, some of which is used to invest in college.

First, the optimal investment in schooling is determined entirely by the individual's ability a and effort e. While family wealth W may be correlated with a person's ability and may influence his or her exertion of effort, it does not have a direct impact on schooling per se.

Second, there is full insurance in consumption. The ability to (i) borrow and lend without restriction and (ii) differentiate the payments d_e and d_u, based on whether the student finds a job right after school, allows for perfect insurance. Perfect insurance is the equilibrium and optimal allocation since, in this environment, there is no reason to distort consumption over time and states of the world.

Implementing the first-best allocations may require large and stark transfers of resources between borrowers and lenders. First, a high-ability individual with low or no family wealth would need to borrow heavily at $t = 0$. In expectation, the lender would require a large net repayment, but this repayment would be very different depending on labor market outcomes. On the one hand, if the person fails to find a job, the lender might be required to make a transfer to the unemployed person. The optimal loan policy requires $d_u < 0$ (i.e., the lender making a payment to the student).

In a first-best world, therefore, student loan programs must also include an unemployment insurance mechanism. This type of unemployment insurance would differ greatly from conventional ones, since it is inherently associated with school investments, not previous employment. Also, note that existing mechanisms in U.S. student loan programs fall short of the insurance provisions in this environment. In practice, U.S. students can delay repayments using either deferments and/or forbearance for several different reasons. Moreover, students can opt for an income-contingent repayment program that fixes or limits the repayments to a fraction of their earnings. These provisions fall short of the prescriptions of the first-best allocations that might require compensation from the lender to the borrower, not just a limit or postponement of the payments from the borrower to the lender.

On the other hand, if the student finds a job right after school, then loan repayment must be expected over periods $t = 1$ and $t = 2$. This repayment is larger, in present value terms, than the original loan because it has to cover the expected cost of the unemployment insurance payments. Indeed, the high repayment of those loans can be a serious deterrent to student credit because of limited commitment and enforcement. This aspect has been extensively studied by itself (Lochner and Monge-Naranjo, 2011, 2012) and in relation to its interactions with other frictions (Lochner and Monge-Naranjo, forthcoming).

In the next two sections, I focus on moral hazard problems underlying the specific problem of finding a job right after school. Dealing with this problem explicitly provides guidance on how to optimally design repayment (and granting) of student loans.

4 OPTIMAL LOANS WITH MORAL HAZARD

A key assumption in the first-best allocation is the ability of the lenders to directly recommend and control the effort exerted by the borrowers. Since there is full insurance for the borrower, the optimal arrangement results in lenders being the only party really interested in the student finding a job right after school. Not surprisingly, a moral hazard problem arises if the student's effort cannot be perfectly monitored. In this section and in the next, I assume that effort cannot be observed or controlled by the lender in any way.

Under these circumstances, a contract that offers full insurance over youth unemployment is destined to fail unless it is based on the student exerting the minimum effort e to find a job. Under either continuous or discrete effort choice, $V_e = V_u$ unambiguously implies zero effort; $e = 0$. Clearly, to induce any positive effort, $e > 0$, it is necessary to break full insurance coverage by imposing $V_e > V_u$. Furthermore, the larger the difference between V_e and V_u, the larger would be the effort exerted by the student. This is called an incentive compatibility constraint (ICC) in any optimal arrangement with hidden action.

The optimal credit arrangement (and student loan contract) for a person with given ability a and family resources W maximizes the lifetime expected utility of the agent by choosing consumption levels for all dates and states of the world and investments in human capital and effort such that lenders break even (BE) in expectation and the ICC ensures that the borrower exerts the effort expected by the lender.

For brevity, in this section I consider only the discrete case. I set up the problem in terms of future utilities for the borrower, V_e and V_u, for students finding a job right after school or only at maturity. While this formulation requires additional explanation, it is very useful in Section 5, where I extend the model to a multiperiod setting.

First, note that if $e = 0$ is the effort sought by the lender, then the optimal contract is simply the first-best allocation described previously. Next, consider the optimal allocation if high effort, $e = 1$, is sought by the lender to maximize the chances of employment. Then, the optimal contract is given by the following loan program:

$$U^*(a,W;e=1) = \max_{c_0,h,V_e,V_u} u(c_0) - v(1) + \beta\left[p_H V_e + (1-p_H)V_u\right],$$

subject to the ICC that the expected gains from exerting effort more than compensate the cost

$$\beta(p_H - p_L)(V_e - V_u) > v(1) - v(0)$$

and subject to the BE constraint

$$-c_0 - h + W + \beta\left\{p_H\left[ah^\alpha(1+\beta G) - C(V_e)\right] + (1-p_H)\left[\beta ah^\alpha - C(V_u)\right]\right\} \geq 0,$$

where $C(\cdot)$ is the cost function that determines the cost for the lender to deliver a continuation utility level for the borrower. The function $C(\cdot)$ is strictly increasing and strictly convex as it is derived from the inverse of $u(\cdot)$. For constant relative risk aversion (CRRA) preferences, I can explicitly derive the form of this cost function as[11]

$$C(V)=(1+\beta)\left[\frac{(1-\sigma)V}{(1+\beta)}\right]^{\frac{1}{1-\sigma}}.$$

There are many alternative formulations of this principal-agent problem, all of which are equivalent. For example, in this case the lender takes possession of the labor earnings and commits to deliver continuation value V_e, letting $\mu \geq 0$ denote the Lagrange multiplier for the ICC and $\lambda > 0$ denote the multiplier for the BE constraint.[12]

After some straightforward simplifications, the FOCs that characterize the optimal allocations are

$$\begin{aligned}
&[c_0]: u'(c_0)=\lambda;\\
&[V_e]: \lambda C'(V_e)=1+\mu\left(1-\frac{p_L}{p_H}\right);\\
&[V_u]: \lambda C'(V_u)=1-\mu\left[1-\left(\frac{1-p_L}{1-p_H}\right)\right];\\
&[h]: \beta\left[p_H(1+\beta G)+(1-p_H)\right]\alpha a h^{\alpha-1}=1.
\end{aligned}$$

These simple conditions deliver a very sharp set of implications for the behavior of consumption over time and across states and for investments in human capital h. First, note that $C'(V)=\left[\frac{(1-\sigma)V}{(1+\beta)}\right]^{\frac{\sigma}{1-\sigma}}$. Second, in this environment the cheapest way to deliver the continuation utilities V_e and V_u is with different, but constant over time, consumption flows c_e and c_u; that is, $V_e=(1+\beta)\frac{c_e^{1-\sigma}}{1-\sigma}$ and $V_u=(1+\beta)\frac{c_u^{1-\sigma}}{1-\sigma}$. Then, $C'(V_e) = (c_e)^{\sigma} = 1/u'(c_e)$ and $C'(V_u) = (c_u)^{\sigma} = 1/u'(c_u)$. In this case, the result was directly obtained because there was a closed form for $C(\cdot)$. However, the result that $C'(V') = 1/u'(c')$ holds much more generally as a direct application of the envelope condition on the value function $C(\cdot)$. Then, the first three FOCs can be succinctly summarized as

(5) $$u'(c_0)=\left[1+\mu\left(1-\frac{p_L}{p_H}\right)\right]u'(c_e),$$

(6) $$u'(c_0)=\left[1-\mu\left(1-\frac{1-p_L}{1-p_H}\right)\right]u'(c_u).$$

First, note that if the ICC does bind (i.e., when $\mu > 0$), the term in brackets in the first expression is necessarily greater than 1; the opposite holds for the second expression. Obviously, this implies that $u'(c_e) < u'(c_0) < u'(c_u)$ and, because of strict concavity, these inequalities necessarily imply the following ordering for the consumption levels of initially employed c_e and initially unemployed c_u graduates, relative to their consumption while in school:

$$c_u < c_0 < c_e.$$

The crucial inequality is $c_u < c_e$, which is needed for the student to find it optimal to exert costly effort to find a job right after graduation. The inequalities with respect to c_0 are driven by my assumption that the implicit interest rate coincides with the borrower's discount factor. Of more interest, note that the wedges between consumption levels are determined by the ratio p_L/p_H, the likelihood ratio of finding a job without exerting any effort versus exerting it, and $(1 - p_L)/(1 - p_H)$, the likelihood ratio of not finding a job while not exerting effort versus not finding a job while exerting effort. When either (i) the ICC is not binding ($\mu = 0$) because effort is costless, $v(1) = v(0)$, or because the low effort is pursued by the credit arrangement or (ii) failing to find a job is uninformative about the effort exerted by the borrower, $p_L = p_H$, then consumption will be perfectly smooth. Otherwise, consumption will always be shifted in favor of those who find a job right after school as a mechanism to reward the exertion of effort.

Finally, despite the presence of moral hazard, there is a stark implication for the investment levels. Conditional on effort e, the investment in human capital h is always at the efficient level $h^*(a,e=1)=\left\{\beta\left[p_H(1+\beta G)+(1-p_H)\beta\right]\alpha a\right\}^{\frac{1}{1-\alpha}}$. That is, as long as the contract arrangement finds it optimal to request a high level of effort, the distortions on consumption not only can deliver the optimal investment, but also find it optimal to do so. The result, however, is conditional only on effort. In general, the sets of abilities a and family wealth levels W for which the individual exerts high effort are different in the first-best allocations than when moral hazard is a concern.

4.1 Implied Credit and Repayments

I have solved the optimal contract formulating the problem in what is called the *primal approach*—that is, looking at the allocations of consumption and human capital while imposing the ICC and the BE constraint for the lender. Underlying these constrained optimal allocations are transfers of resources—credit and repayment—between the borrower and the lender.

Let us first consider transfers when no effort is required (i.e., $e = 0$). As indicated previously, the ICC is not implemented, $\mu = 0$, and the perfect smoothing in consumption holds. The investments in schooling are

$$h_L^*(a)=\left\{\beta\left[p_L(1+\beta G)+(1-p_L)\beta\right]\alpha a\right\}^{\frac{1}{1-\alpha}},$$

where the subscript L stands for low effort, $e = 0$.

Initial wealth W implies a total lifetime of net resources $W + (1/\alpha - 1)h_L^*(a)$ and constant consumption levels $c_L^*(a,W) = [W + (1/\alpha - 1)h_L^*(a)]/(1 + \beta + \beta^2)$. With this information, at $t = 0$, an amount of borrowing $b_L(a,W)$

$$b_L(a,W)=\left(\frac{1/\alpha+\beta+\beta^2}{1+\beta+\beta^2}\right)h_L^*(a)-\left(\frac{\beta+\beta^2}{1+\beta+\beta^2}\right)W,$$

is needed to consume and invest $c_L^*(a,W) = h_L^*(a)$. Borrowing is always decreasing with own resources W. It is increasing with ability a not only because students with higher ability demand more current resources to invest, but also because higher future incomes are expected. Borrowing can be negative (savings) if the person is rich relative to his or her ability—that is, high W and/or low a.

The repayment of this loan is structured as follows: If the person is employed at $t = 1$, then he or she must make a repayment

$$p_{L,1}^e(a,W) = a\left[h_L^*(a)\right]^\alpha - \frac{\left[W + (1/\alpha - 1)h_L^*(a)\right]}{1+\beta+\beta^2}$$

and another one equal to

$$p_{L,2}^e(a,W) = Ga\left[h_L^*(a)\right]^\alpha - \frac{\left[W + (1/\alpha - 1)h_L^*(a)\right]}{1+\beta+\beta^2}$$

at $t = 2$. However, with competitive post-schooling borrowing and lending markets, the timing of the payments does not matter as long as the present value of the repayments as of $t = 1$ is equal to a debt balance equal to

$$D_L^e(a,W) = (1+\beta G)a\left[h_L^*(a)\right]^\alpha - \frac{1+\beta}{1+\beta+\beta^2}\left[W + (1/\alpha - 1)h_L^*(a)\right].$$

In contrast, if the person remains unemployed right after college, he or she would receive a transfer from the lender. The repayment, $p_{L,1}^u(a,W)$, would be negative:

$$p_{L,1}^u(a,W) = -\frac{\left[W + (1/\alpha - 1)h_L^*(a)\right]}{1+\beta+\beta^2};$$

as for $t = 2$, the payments would be

$$p_{L,2}^u(a,W) = a\left[h_L^*(a)\right]^\alpha - \frac{\left[W + (1/\alpha - 1)h_L^*(a)\right]}{1+\beta+\beta^2}.$$

Note that these payments incorporate not only an element of insurance, but also the returns to the person's own initial wealth W.

Now, consider the more interesting case in which high effort ($e = 1$) is induced. For brevity, I use the subscript H here and define

$$M \equiv \left[1 + \mu\left(1 - \frac{p_L}{p_H}\right)\right]^{\frac{1}{\sigma}} > 1$$

and

$$m \equiv \left[1 - \mu\left(1 - \frac{1-p_L}{1-p_H}\right)\right]^{\frac{1}{\sigma}} < 1.$$

From the FOCs (5) and (6), it can be verified that $c_e = Mc_0$ and $c_u = mc_0$, so what remains is the determination of $c_{0,H}(a,W)$, the initial consumption of someone who exerts high effort in school. Recall that (i) the first-best level of investment for $e = 1$ attains in this case and is equal to $h_H^*(a)=\left\{\beta\left[p_H(1+\beta G)+(1-p_H)\beta\right]\alpha a\right\}^{\frac{1}{1-\alpha}}$ and (ii) the net present value of lifetime resources at risk-neutral prices is $W + (1/\alpha - 1)h_H^*(a)$. Then, the budget constraint for the consumption profiles of the borrower is reduced to

$$c_0+\beta\left[(1+\beta)p_H c_e+(1+\beta)(1-p_H)c_u\right]=W+(1/\alpha-1)h_H^*(a)$$

or

$$c_0\left[1+\beta(1+\beta)\left(p_H M+(1-p_H)m\right)\right]=W+(1/\alpha-1)h_H^*(a).$$

Thus, the terms m and M act as shifters in the relative probability weights for the consumption levels of individuals who find employment early and late, respectively. Since $m < 1 < M$, the shift increases the implicit probability weight for early employment, and since $c_e > c_0$ and $c_u < c_0$, the initial consumption c_0 is *lower* than in the first-best. In any event, the solutions for the consumption levels are

$$c_{0,H}(a,W)=\frac{1}{1+\beta(1+\beta)\left[m+p_H(M-m)\right]}\left[W+(1/\alpha-1)h_H^*(a)\right],$$

$$c_{e,H}(a,W)=\frac{M}{1+\beta(1+\beta)\left[m+p_H(M-m)\right]}\left[W+(1/\alpha-1)h_H^*(a)\right],$$

$$c_{u,H}(a,W)=\frac{m}{1+\beta(1+\beta)\left[m+p_H(M-m)\right]}\left[W+(1/\alpha-1)h_H^*(a)\right].$$

Therefore, the borrowing and repayments for a person exerting effort $e = 1$ are as follows: At $t = 0$, the consumption $c_{0,H}(a,W)$ and investments $h_H^*(a)$ require borrowing

$$b_H(a,W)=\frac{\frac{1}{\alpha}h_H^*(a)+\beta(1+\beta)\left[m+p_H(M-m)\right]\left[h_H^*(a)-W\right]}{1+\beta(1+\beta)\left[m+p_H(M-m)\right]}.$$

Similar relationships with respect to a and W hold as in the case with low effort. However, note that the levels are different. First, high effort $e = 1$ implies higher investments, $h_H^*(a) > h_L^*(a)$. From here, for any W and a, borrowing with high effort is higher—potentially much higher—than borrowing with low effort. Second, a more subtle mechanism is at work: The weights for consumption with low effort are simply $1 + \beta + \beta^2$, but once the ICC is imposed to incentivize effort, it is $1 + \beta(1 + \beta)[m + p_H(M - m)]$. The difference between the two, of course, is determined by the magnitude of the multiplier μ.

With respect to repayments, a person who finds employment at $t = 1$ and signed a contract that implements high effort must repay

$$p_{H,1}^{e}(a,W)=a\left[h_H^*(a)\right]^{\alpha}-c_{e,H}(a,W)$$

at $t = 1$ and

$$p_{H,2}^{e}(a,W)=Ga\left[h_H^*(a)\right]^{\alpha}-c_{e,H}(a,W)$$

at $t = 2$. As before, with competitive post-school borrowing and lending markets, the timing of the payments is immaterial as long as the $t = 1$ present value of payments from $t = 1$ and $t = 2$ adds up to

$$D_H^{e}(a,W)=(1+\beta G)a\left[h_H^*(a)\right]^{\alpha}-(1+\beta)c_{e,H}(a,W).$$

On the contrary, if the borrower remains unemployed after school, then the optimal contract would prescribe a *negative repayment* for $t = 1$:

$$p_{H,1}^{u}(a,W)=-c_{u,H}(a,W).$$

Unambiguously, the lender would have to transfer additional resources to pay for the consumption of the unemployed borrower. For the second period, the payments must be

$$p_{H,2}^{u}(a,W)=a\left[h_H^*(a)\right]^{\alpha}-c_{u,H}(a,W),$$

which might be positive. In any event, as before, with postcollege borrowing and lending, the key is that the debt balance is set to

$$D_H^{u}(a,W)=\beta a\left[h_H^*(a)\right]^{\alpha}-(1+\beta)c_{u,H}(a,W),$$

which might be positive or negative.

The optimal contract has clear prescriptions for the level of credit granted and the investments made by different individuals according to their initial wealth and ability. According to these characteristics, a level of effort would be implemented and repayments would be set contingent on whether the individual finds a job right after college. Such repayments necessarily entail elements of insurance, and unless the initial wealth W is very high, the lender must end up making a net transfer for the unemployed. While successful individuals—that is, those who are employed right after school—are required to make larger repayments, the repayments are limited by the need to reward effort in the first place.

I wish to comment on simple yet stark implications of my analysis. First, by the proper structure of consumption distortions, the optimal contract attains the first-best investment in human capital. Once ability and effort are controlled for, the investment of individuals should be independent of family wealth and background. Second, the structure of repayments is very different from the existing mechanisms of deferments, forbearances, and the emerging income-contingent repayment loans that limit the repayment to a fraction of the income accrued and, therefore, forgive debts of the unemployed. The optimal contract can indeed go beyond this by prescribing an unemployment insurance transfer.

4.2 Discussion: Incentives Versus Insurance

Hidden action problems while studying in college or later, when searching for a job, lead to a form of the proverbial trade-off between insurance and incentives. On the one hand, consider the case in which lenders, including the government, ignore the incentive problems. Aiming for the best competitive contract in their perceived environment, lenders would offer the first-best allocations. If, however, *none* of the borrowers, not a single one—regardless of their wealth W or ability a—would exert effort, the unemployment rate of this cohort of students would be $1 - p_L$, possibly much higher than the lenders' expectations. This higher unemployment rate would also be the culprit behind financial losses for lenders, as they receive lower repayments from the fewer borrowers who found employment and make higher transfers to those who remained unemployed. Over time, the lack of self-sustainability implies that private lending would simply disappear and government lending programs would have consistent deficits that should be financed with taxes.

On the other hand, credit arrangements that ignore the desire for insurance would be suboptimal, not only because of the lack of consumption insurance, but also because of suboptimal investments in human capital. For instance, consider a credit arrangement that requires the same repayment regardless of whether the person finds work. Trivially, unless there is a post-schooling credit market, the only arrangement is the autarkic one discussed previously since unemployed youth are unable to repay anything in this environment. With post-schooling credit markets, the set of student loans with constant repayments across youth unemployment or employment is significantly larger. But in any such case, the consumption across states of the world would be disrupted, especially for those with high ability and low family wealth, who might have invested and borrowed larger amounts. In response, the borrowing and investment in human capital of those individuals could be severely limited. See the appendix for further discussion.

Interestingly, in this environment under the optimal contract, the proper balance between incentives and insurance enables efficient investments in human capital. This efficiency result is conditional on effort and, thus, the distortions are on consumption and/or the exertion of effort, but, given the latter, not on the schooling investment level. For this efficiency result, the optimal contract arrangement may require additional transfers from the lender to the borrower in case of unemployment, and if so, high payments in case of employment. However, the efficient repayment must always reward the consumption levels of the successfully employed relative to those who are unemployed after school.

5 MULTIPERIOD UNEMPLOYMENT SPELLS

The stylized three-period model analyzed earlier (see Section 2.1) assumes that all college-educated workers eventually find employment. Unemployment may be avoided, but if it happens, it is only at the beginning of the labor market experience of college-educated workers. Within this model, I derived some basic implications for the optimal design of student loans and their repayment. In particular, I found that the efficient investment levels can be delivered as long as the repayment of the loans is designed with the proper balance between incentives and insurance. I also highlighted the fact that the optimal student loan program should incorporate an unemployment insurance mechanism as an integral component.

In this section, I explore an important aspect of such an unemployment insurance mechanism. To do so, I need to incorporate into the model the fact that finding a job may require repeated search efforts. The unemployment insurance mechanism should provide the right incentives and not dissuade new graduates from seeking employment. To explore this important dimension, in this section I change the specification for the post-schooling lifetime. Specifically, I now consider an environment in which the labor market participation is open-ended—that is, I consider a simple infinite horizon model for the job search after school. This version of the model is a simplification of the already simple and well-known Hopenhayn-Nicolini (1997) model of unemployment insurance. The model is highly stylized; in Section 6 I discuss a number of life cycle issues that would be desirable to explore.

Consider the following environment. As before, knowing a young person's ability a and initial wealth W, the person must decide how much to invest in human capital h and how much effort e to exert. Ability and schooling investments determine the income $y = ah^{\alpha}$ of a worker *once he or she is employed*. Similarly, a level of effort e during school determines the student's probability of finding a job right after school. Otherwise, if the student finishes school with no job, he or she can also search for employment after school. For concreteness, I restrict attention to the discrete case, $e \in \{0,1\}$; that is, where the student exerts effort either fully or not at all.[13] For simplicity, as in Hopenhayn-Nicolini (1997), I assume that once a worker has found a job, regardless of whether it is right after school, the worker will remain employed forever. Then, if unemployment happens, it is in the initial postcollege periods. The length of this initial unemployment spell is endogenous.

The time interval length for schooling differs from the length of each period in the labor market. College can be considered a period of four years; however, for modeling the job search a more meaningful breakdown of time is biweekly or monthly time intervals. To this end, the discount factor between the flow utility and the continuation utility for $t = 0$ is β, as before.[14] For all subsequent periods, the discount factor between the current flow and continuation utilities is $\delta \in (\beta,1)$. Likewise, the probability of getting a job while still in school is p_H if $e = 1$ or p_L if $e = 0$. After school, if the graduate is unemployed, the probabilities of finding a job are q_H or q_L, depending on the effort. Finally, the costs of exerting effort are given by the increasing functions $v(e)$ and $s(e)$ for the periods $t = 0$ and $t \geq 1$, respectively. For simplicity, I normalize the search costs so that $s(0) = 0$. In any event, the expected lifetime utility of a person as of $t = 0$ is given by

$$U = u(c_0) - v(e_0) + \beta E\left[\sum_{t=1}^{\infty} \delta^{t-1}\left[u(c_t) - s(e_t)\right]\right],$$

where $E[\cdot]$ is the expectation over consumption and search efforts and $u(\cdot)$ is the CRRA utility function used throughout the article.

Relative to the stylized model of previous sections, the enhanced search dynamics in this environment do not add anything interesting unless there is an incentive problem. All individuals would be fully insured in their consumption levels regardless of the length of unemployment spells. Conditional on search effort levels, which would depend on ability *a* and family wealth *W*, the investments in human capital *h* would maximize expected lifetime earnings.

Much more interesting is the behavior of consumption, effort, and investments when effort is a hidden action. I now explore the optimal contract with moral hazard, in which effort, if desirable, must be induced by the proper ICC. My treatment of the problem in the previous section can be conveniently adapted to the current environment: Consider the utility optimization of a student with ability *a* and wealth *W*, for whom the lenders aim to induce high effort. The optimization problem is very similar to the previous case:

$$U^*(a,W;e=1)=\max_{c_0,h,V_e,V_u} u(c_0)-v(1)+\beta\left[p_H V_e+(1-p_H)V_u\right],$$

subject to exactly the same ICC as before:

$$\beta(p_H-p_L)(V_e-V_u)>v(1)-v(0),$$

and to a different BE constraint:

$$-c_0-h+W+\beta\left\{p_H\left[\frac{ah^\alpha}{1-\delta}-\frac{\left[(1-\sigma)V_e\right]^{\frac{1}{1-\sigma}}}{1-\delta}\right]+(1-p_H)\left[-C\left(V_u,ah^\alpha\right)\right]\right\}\geq 0.$$

The first term inside the square brackets is the next payoff for the lenders if the student is employed right away. Here $\frac{ah^\alpha}{1-\delta}$ is the present value of the resources generated by the individual, while the term $\frac{\left[(1-\sigma)V_e\right]^{\frac{1}{1-\sigma}}}{1-\delta}$ is the present value of the cost of providing a (constant) consumption flow that delivers a lifetime utility equal to V_e. The second bracketed term, $-C(V_u,ah^\alpha)$, is the cost of delivering an expected continuation utility V_u for an unemployed person whose earnings potential, once employed, is equal to ah^α every period.

The cost function $C(V_u,ah^\alpha)$ is given by the following Bellman equation:

$$C\left(V_u,ah^\alpha\right)=\min\left\{C\left(V_u,ah^\alpha;e=0\right),C\left(V_u,ah^\alpha;e=1\right)\right\};$$

that is, the lowest cost between loan programs with low and high job search efforts. These programs are defined as follows: If the program does not require effort, the minimization is

$$C\left(V,ah^\alpha;e=0\right)=\min_{c,V_e',V_u'}\left\{c+\delta\left[q_L\left(\frac{ah^\alpha}{1-\delta}-\frac{\left[(1-\sigma)V_e'\right]^{\frac{1}{1-\sigma}}}{1-\delta}\right)+(1-q_L)C\left(V_u',ah^\alpha\right)\right]\right\},$$

subject to the promise-keeping constraint

$$V=u(c)+\delta\left[q_L V_e'+(1-q_L)V_u'\right].$$

If the program requires effort, the minimization is

$$C\left(V,ah^{\alpha};e=1\right)=\min_{c,V_e',V_u'}\left\{c+\delta\left[q_H\left(\frac{ah^{\alpha}}{1-\delta}-\frac{\left[(1-\sigma)V_e'\right]^{\frac{1}{1-\sigma}}}{1-\delta}\right)+(1-q_H)C\left(V_u',ah^{\alpha}\right)\right]\right\},$$

subject to the promise-keeping constraint

$$V=u(c)-s(1)+\delta\left[q_H V_e'+(1-q_H)V_u'\right],$$

and the ICC:

$$\delta(q_H-q_L)(V_e'-V_u')\geq s(1).$$

A number of interesting implications arise from this program. First, the low-effort level $e = 0$ would be optimal for very high and very low levels of V. On the one hand, for very high levels of utility V, the cost of compensating the agent for exerting effort might be too high in terms of consumption goods, given the decreasing marginal utility of consumption. That is, inducing effort from the very rich might be too costly. On the other hand, for very low levels of utility, if the coefficient of relative risk aversion $\sigma < 1$, then a low level of utility V may leave little room to generate the needed gap between V_e' and V_u' to generate effort. That is, it is very difficult to induce the very poor to exert the costly effort $s(1)$ since for them, there might be little to gain.

In either of these cases, from the FOCs of V_e' and V_u' and the envelope condition on $C(V,ah^{\alpha}; e = 0)$, I can show that

$$V_e'=V_u'=V;$$

that is, the person's continuation utilities would be the same as the initial one, regardless of whether he or she finds a job. In both extremes of utility entitlement levels, as long as a young person remains unemployed, he or she would receive a constant unemployment compensation (UE) equal to

$$UE(V)=u^{-1}\left[(1-\delta)V\right].$$

Upon obtaining a job, the person must repay a constant payment (p) to the lender equal to

$$p=ah^{\alpha}-u^{-1}\left[(1-\delta)V\right].$$

More interesting is the case when high effort $e = 1$ is implemented. In this case, the ICC requires that $V_e' > V_u'$. Indeed, the ICC will hold with equality because the borrower is risk averse and the lender is risk neutral, so $V_e'=V_u'+\frac{s(1)}{\delta(q_H-q_L)}$. Using the envelope condition on $C(V,ah^{\alpha}); e = 1)$, it can be shown that

$$V_u'<V.$$

That is, if the person fails to find a job in this period, his or her continuation utility will be lower next period. The optimal program prescribes a decreasing sequence of utilities $\{V_t\}$ for as long as the person remains unemployed.

Here I emphasize that both continuation utilities for those employed and those unemployed, V'_e and V'_u, respectively, decline over time. Therefore, unemployment compensation declines with each period the person fails to find a job until, possibly, hitting a threshold after which no effort is induced. But the fact that V'_e also declines with the unemployment spell indicates that the amount to be repaid forever after finding a job has to be increasing; that is,

$$p_t = ah^{\alpha} - u^{-1}\left[(1-\delta)V_t\right]$$

is increasing as V_t declines. As shown by Hopenhayn and Nicolini (1997), most of the insurance and efficiency gains arise not only from a declining path in the unemployment compensation transfers, but also from an earnings tax that grows with the length of time a worker needs to finds a job.

In sum, once I incorporate the possibility of multiperiod search with hidden action, the optimal arrangement extends the results of the previous section. Although the optimal arrangement still implies unemployment insurance right after school, over the unemployment spell there is a clear downward trend in unemployment transfers and an upward trend in the flow of payments once a job is found. The first result is well known and the second is highlighted by Hopenhayn and Nicolini (1997) as a crucial aspect in improving the provision of insurance and incentives.

Adapting some aspects of the optimal declining unemployment compensation/increasing debt repayments may lead to substantial gains in the efficiency of human capital investments and welfare generated from student loan programs. Unfortunately, to my knowledge, the use of variations in the balance of the debt as a function of the length of the initial unemployment (or in richer models, in any subsequent spell) is not a feature of any current student loan programs.

CONCLUSION

Unemployment is a major risk for college investment decisions. Even if a college graduate eventually finds a job that matches his or her qualifications—thereby enabling the repayment of loans—the possibility of long spells of unemployment, underemployment, and difficulty repaying student loans may limit and even dissuade productive investments in human capital.

In this article, I explored the optimal design of student loans in the face of a higher incidence of unemployment for the earlier periods of a person's labor market experience. Using a highly stylized model, I derived three main conclusions. First, the optimal student loan program must incorporate an unemployment compensation mechanism as a key element, even if unemployment probabilities are endogenous and subject to moral hazard. Second, even under moral hazard, a well-designed student loan program can deliver efficient levels of investments. Distortions in consumption may remain, but the labor market potential of any individidual, regardless of family background, should not be impaired as long as he or she is willing to put forth the effort. Third, the provision of unemployment benefits and debt balances must be set as functions of the unemployment spell to provide the right incentives for youth to seek employment.

These conclusions are valid for all forms of college education investment. However, the brief data exploration in this article uncovered substantial differences across college majors in terms of unemployment, underemployment, and levels and growth of earnings after graduation. In principle, the schooling costs of the different majors can also vary widely, partly because the effective time for graduation in different majors can vary but also because the cost of equipment and the salaries of professors and other instructors can vary widely. A natural extension of the analysis presented here would quantitatively examine how these differences should translate into the credit provision and repayment schedules for students opting for various majors. A subsequent step would be to examine the implications for income distribution and social mobility of reforming the current student loan system for schemes that efficiently cope with the different incentive problems involved in financing a country's higher education. Pursuing this line, of course, would require accounting for the general equilibrium implications of such a policy change, including the labor market tightness and job-finding rates for graduates in all the different fields. ■

APPENDIX: A WORLD WITH NO STUDENT LOANS

As additional benchmarks, in this appendix I explore the allocations when student loans are not available. I first consider the case in which there are no financial markets at all. Then, I consider the case in which financial markets are available for the post-schooling stage for those active in labor markets.

Case 1: Complete Autarky

Schooling investment and effort levels are the choices to be made by a young person in the first period, $t = 0$. Since no financial markets are available, the first-period consumption is simply $c_0 = W - h$; that is, initial resources minus investments h. After school, the inability to transfer resources across periods implies that consumption would always be equal to labor earnings. Then, the optimal schooling choices solve the problem

$$\max_{e,h} u(W-h)-v(e)+\beta\left\{p(e)\left[u\left(ah^{\alpha}\right)+\beta u\left(Gah^{\alpha}\right)\right]+\left(1-p(e)\right)\left[u(b)+\beta u\left(ah^{\alpha}\right)\right]\right\}.$$

Given an effort level e and the CRRA utility function, the optimal investment $h(a,W;e)$ is uniquely pinned down by the FOC[15]:

$$(W-h)^{-\sigma}=\beta\alpha a^{1-\sigma}\left[p(e)\left(1+\beta G^{1-\sigma}\right)+\left(1-p(e)\right)\beta\right]h^{\alpha(1-\sigma)-1}.$$

Given the lack of credit markets, schooling investments are limited by the individual's own wealth W; it is straightforward to show that schooling investments h are always increasing in W. Similarly, schooling investments are always increasing in the exertion of effort e, as it increases the probability of youth employment, with the potential added benefit of the labor market experience gains G.

From the point of view of efficiency, perhaps the most relevant relationship is the one between an individual's ability a and investments in schooling h. In this environment, this relationship is entirely driven by the opposing *wealth effects* (a smarter individual has more future earnings for the same investment) and *substitution effects* (a smarter individual gains more future income for each additional investment). The relative strength of these effects is governed by $1/\sigma$, the intertemporal elasticity of substitution (IES). If $\sigma < 1$, the IES is large and a positive relationship between ability and investment arises. The more problematic case, in terms of efficiency and also relative to empirical evidence, is when $\sigma > 1$, which is the most common condition for quantitative work. In such a case, the IES is low and a negative relationship between a and h is implied by the model. Here, the smarter a person is—that is, the higher a—the higher would be his or her future consumption for any investment level. In addition, his or her optimal investment would be lower to enhance consumption at time $t = 0$, as it is entirely limited by W. It is worth noting the relationship between schooling investment h and the gains from labor market experience G. These results are explored further in Lochner and Monge-Naranjo (2011).

In the case of $\sigma = 1$ (i.e., log preferences), the wealth and substitution effects cancel each other, leading to a simple formula:

$$h = \frac{\beta\alpha\left(p(e)+\beta\right)}{1+\beta\alpha\left(p(e)+\beta\right)}W;$$

that is, investment should always be a constant fraction of the individual's wealth. That fraction depends on (i) the discount factor, (ii) the elasticity of income relative to investment α, and (ii) the individual's exerted effort in school and the early labor market e, but not on his or her ability a or experience gains G.

Last but not least, the optimal exertion of effort is determined by the difference in the value of the career of a young person who is employed right after school and the value of the career of a young person who is unemployed during his or her youth and works only during maturity. Let V_e and V_u denote, respectively, these post-schooling labor market career values:

$$V_e = u\left(ah^{\alpha}\right) + \beta u\left(Gah^{\alpha}\right)$$

and

$$V_u = u(b) + \beta u\left(ah^{\alpha}\right).$$

Then, the optimal exertion of effort is as follows. If effort is a continuous variable, $e \in [0,\infty)$, the optimal is given by

(7) $$v'(e) = \beta p'(e)\left[V_e - V_u\right].$$

That is, the higher the *absolute gap* between the two career outcomes, the more effort the individual would exert in seeking employment right after school. In the alternative case where effort is a discreet choice—that is, $e \in \{0,1\}$—for brevity, let $p_H = p(1)$ and $p_L = p(0)$. Obviously, $0 \leq p_L < p_H < 1$ and the optimal exertion of effort would be given by

(8) $$e = \begin{cases} 1 & \text{if } \beta\left[(p_H - p_L)(V_e - V_u)\right] \geq v(1) - v(0), \\ 0 & \text{otherwise.} \end{cases}$$

That is, a young person would be interested in assuming the higher cost of effort, $v(1) - v(0)$, only if the returns of that effort are more than compensated by the expected career gains, $\beta[(p_H - p_L)(V_e - V_u)]$. Here β captures the fact that (i) the gains occur in the future, (ii) $p_H - p_L$ is the increased probability of finding a job if effort is exerted, and (iii) $V_e - V_u$ is the net gain of being employed.

Case 2: Post-Schooling Borrowing and Lending

Credit cards, auto loans, mortgages, and many other forms of credit besides student loans might be available once a person has started participating in labor markets. While these forms of credit might not be available for individuals to finance their education, their presence may alter the human capital decisions, especially in the face of early unemployment.

To incorporate some of the main implications of these forms of credit into my simple model, consider the case in which students cannot borrow or save in $t = 0$ but between periods $t = 1$ and $t = 0$ they can fully smooth consumption by borrowing or lending. Given my assumption that $\beta(1 + r) = 1$, after realizing their employment status in $t = 1$, their consumption will be fully smoothed whether employed or unemployed between periods $t = 1$ and $t = 2$. If they are employed, then the present value of consumption will be $ah^{\alpha}(1 + \beta G)$, and consumption at $t = 1$ and $t = 2$ will be equal to $c_1 = c_2 = ah^{\alpha}\frac{(1+\beta G)}{1+\beta}$. The value of a career with early employment is then

$$V_e = \Theta_e \frac{\left[ah^{\alpha}\right]^{1-\sigma}}{1-\sigma},$$

where $\Theta_e \equiv (1 + \beta)^{\sigma}[1 + \beta G]^{1-\sigma}$.

Likewise, an unemployed young person would borrow to consume at $t = 1$ and repay at $t = 2$. The present value of resources βah^{α} will be equally consumed in both periods, leading to $c_1 = c_2 = \frac{ah^{\alpha}}{1+\beta}$,[16] and a value, V_u, as of $t = 1$ equal to

$$V_u = \Theta_u \frac{\left[ah^{\alpha}\right]^{1-\sigma}}{1-\sigma},$$

where $\Theta_u = (1 + \beta)^{\sigma}$.

With those values in place, I can succinctly define the problem for choosing effort e and schooling investments h as

$$\max_{e,h} \frac{[W-h]^{1-\sigma}}{1-\sigma} - v(e) + \beta\left[p(e)\Theta_e + (1-p(e))\Theta_u\right]\frac{\left[ah^{\alpha}\right]^{1-\sigma}}{1-\sigma}.$$

For a given choice of effort e in finding a job while young, the optimal investment in schooling h is given by

$$\left[W-h\right]^{-\sigma} = \beta\alpha a^{1-\sigma}\left[p(e)\Theta_e + (1-p(e))\Theta_u\right]h^{\alpha(1-\sigma)-1}.$$

Note that while the levels might differ, the implied relationship between investments h, and the individual's ability a and wealth W are both in the same direction as in the previous environment with no markets whatsoever.

Similarly, the optimal exertion of effort is given by the condition

$$v'(e) = \beta p'(e)[\Theta_e - \Theta_u]\frac{\left[ah^{\alpha}\right]^{1-\sigma}}{1-\sigma}$$

for the continuous effort case. For the discrete effort case, the condition is exactly the same for autarky. However, in this case I can write it more explicitly as

$$e = \begin{cases} 1 & \text{if } \beta\left[(p_H - p_L)(\Theta_e - \Theta_u)\right]\frac{\left[ah^{\alpha}\right]^{1-\sigma}}{1-\sigma} \geq v(1) - v(0), \\ 0 & \text{otherwise.} \end{cases}$$

I wish to highlight that although the patterns are similar, the allocations can differ substantially with or without credit markets after school. The ability to borrow from future income allows the person to partially insure against low consumption outcomes in case he or she does not find a job when young (and misses the experience gains for $t = 2$). But this insurance is partial at best. Moreover, the inability to borrow at time $t = 0$ can severely limit the investments of young persons, especially those with high ability but very little material support from their families.

NOTES

1 See Lochner and Monge-Naranjo (2015b) and references therein.

2 ACS data are available at https://www.census.gov/programs-surveys/acs/data.html.

3 The B&B:08/12 is available at https://nces.ed.gov/pubsearch/pubsinfo.asp?pubid=2014641rev.

4 The U.S. civilian unemployment rate during the period ranged from 5.8 percent in July 2008, before the Great Recession, to 9.5 percent in July 2009.

5 Lochner and Monge-Naranjo (2015a) use repayment measures based on individual loan records from the National Student Loan Data System accessed in both 1998 and 2003. The loan status represents the most recent available status date at the time. See Lochner and Monge-Naranjo (2015a) for further details.

6 The data in this section show substantial differences across majors in terms of unemployment and underemployment levels and earnings growth. Schooling costs for the different majors can also vary widely as college professors and even fixed-term instructors command very different wages depending on their field. Examining these differences and whether they should be incorporated in the form of differences in the student loan programs offered to prospective students deserves ample attention—enough to warrant a separate work—and is beyond the scope of this article.

7 It is equivalent to assume that the effort must be exerted at the beginning of $t = 1$. Therefore, in the simple model I do not distinguish between whether the individual's effort is to succeed in school and find a job based on the recommendation of the school or to search early for jobs and interviews. That distinction is present in the multiperiod model later in the article.

8 See the discussion in Lochner and Monge-Naranjo (2012).

9 The FOC is

$$\text{(4a)} \qquad u'\left(\frac{W+(1/\alpha-1)h^*(a;e)}{(1+\beta+\beta^2)}\right)\left(\frac{1-\alpha}{\alpha}\right)\frac{\partial h^*(a;e)}{\partial e}=v'(e).$$

However, notice that $\frac{\partial h^*(a;e)}{\partial e}=\frac{(1+\beta(G-1))}{1-\alpha}\beta\alpha a\left[h^*(a;e)\right]^{\alpha}\times p'(e)$. From the FOC on h, $[\beta + p(e)(1+ \beta(G-1))]\beta\alpha a h^{\alpha} = h$.
Substituting the latter expression into the former, I obtain

$$\frac{\partial h^*(a;e)}{\partial e}=\frac{1}{1-\alpha}\left(\frac{1+\beta(G-1)}{\beta+p(e)(1+\beta(G-1))}\right)h^*(a;e)\times p'(e).$$

Substituting this last equation into the FOC (4a) yields the expression in the text.

10 The FOC is necessary only because multiple crossings can happen unless additional restrictions on $p(\cdot)$ and $v(\cdot)$ are imposed. However, around an equilibrium, the left-hand side of this equation must be decreasing and the right-hand side increasing, otherwise the second-order conditions would be violated.

11 First, note that given the concavity of the utility function $u(\cdot)$, the cheapest form of delivering a utility level V is $(1 + \beta)c$ for a constant flow $c > 0$, such that $V = (1 + \beta)u(c)$. The level c is given by $c=u^{-1}\left[\frac{V}{1+\beta}\right]$. Then, the cost for the lender of providing V is as given in the text.
If $\sigma < 1$, the domain for the function $C(\cdot)$ is $[0,\infty)$; if $\sigma > 1$, its domain is $(-\infty, 0)$.

[12] The Lagrangian is standard:

$$L = u(c_0) - v(1) + \beta[p_H V_e + (1 - p_H)V_u]$$
$$+ \lambda[-c_0 - h + W + \beta\{p_H[ah^{\alpha}(1 + \beta G) - C(V_e)] + (1 - p_H)[\beta ah^{\alpha} - C(V_u)]\}]$$
$$+ \mu[\beta(p_H - p_L)(V_e - V_u) - (v(1) - v(0))],$$

and it is straightforward to show that the FOCs are sufficient since $u(\cdot)$ is strictly concave, $C(\cdot)$ is strictly concave, and the feasible set for $\{c_0, h, V_e, V_u\}$ is convex.

[13] The continuous effort case is very similar. Indeed, the Hopenhayn-Nicolini (1997) model is a continuous effort model.

[14] For an explicit model of a sequential investment in college education, see Garriga and Keightley (2007).

[15] The left-hand side of this equation is obviously increasing in h, while the right-hand side is strictly decreasing since the term within square brackets is positive and the exponent $\alpha(1 - \sigma) - 1$ on h is always negative.

[16] It is assumed that $b = 0$ for simplicity. The assumption of employment during $t = 2$ and borrowing available for $t = 1$ implies that the problem is always well defined.

REFERENCES

Chatterjee, Satyajit and Ionescu, Felicia. "Insuring Student Loans Against the Financial Risk of Failing to Complete College." *Quantitative Economics*, November 2012, *3*(3), pp. 393-420; http://dx.doi.org/10.3982/QE100.

Garriga, Carlos and Keightley, Mark P. "A General Equilibrium Theory of College with Education Subsidies, In-School Labor Supply, and Borrowing Constraints." Federal Reserve Bank of St. Louis Working Paper No. 2007-051A, November 2007; https://research.stlouisfed.org/wp/2007/2007-051.pdf.

Hopenhayn, Hugo A. and Nicolini, Juan P. "Optimal Unemployment Insurance." *Journal of Political Economy*, April 1997, *105*(2), pp. 412-38; http://dx.doi.org/10.1086/262078.

Lee, Sang Y.; Shin, Yongseok and Lee, Donghoon. "The Option Value of Human Capital." Unpublished manuscript, Washington University in St. Louis, November 2015; https://pages.wustl.edu/files/pages/imce/yshin/hc_option.pdf.

Lochner, Lance and Monge-Naranjo, Alexander. "The Nature of Credit Constraints and Human Capital." *American Economic Review*, October 2011, *101*(6), pp. 2487-529; http://dx.doi.org/10.1257/aer.101.6.2487.

Lochner, Lance J. and Monge-Naranjo, Alexander. "Credit Constraints in Education." *Annual Review of Economics*, *Annual Reviews*, September 2012, *4*(1), pp. 225-56; http://dx.doi.org/10.1146/annurev-economics-080511-110920.

Lochner, Lance J. and Monge-Naranjo, Alexander. "Default and Repayment among Baccalaureate Degree Earners," in Brad Hershbein and Kevin M. Hollenbeck, eds. *Student Loans and the Dynamics of Debt*. Chap. 8. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research, 2015a, pp. 235-86; http://dx.doi.org/10.17848/9780880994873.ch8.

Lochner, Lance J. and Monge-Naranjo, Alexander. "Student Loans and Repayment: Theory, Evidence and Policy," in Eric A. Hanushek, Stephen J. Machin, and Ludger Woessmann, eds., *Handbook of the Economics of Education*. Chap. 8. Amsterdam: North-Holland, 2016, pp. 397-478; http://dx.doi.org/10.1016/b978-0-444-63459-7.00008-7.

Wang, Cheng and Williamson, Stephen. "Unemployment Insurance with Moral Hazard in a Dynamic Economy." *Carnegie-Rochester Conference Series on Public Policy*, June 1996, *44*(1), pp. 1-41; http://dx.doi.org/10.1016/S0167-2231(96)90002-0.