

## *Editors' Summary*

THE BROOKINGS PANEL ON Economic Activity held its eighty-third conference in Washington, D.C., on March 29 and 30, 2007. This issue of the Brookings Papers includes the papers and discussions presented at the conference. The five papers in this issue span a range of domestic and global issues of current importance. The first paper uses risk analysis to assess the large foreign reserves holdings of emerging market countries. The second paper models the role of information technology in the rapid productivity growth of the past decade. The third paper examines the relation between foreign capital inflows and growth among nonindustrialized economies. The fourth paper estimates the effects of dividends on consumption and the implied effects of the 2003 dividend tax cuts. The final paper looks for explanations for the failure of long-term rates to respond as expected to the Federal Reserve's monetary tightening of recent years.

IN 2005 THE INTERNATIONAL reserves of emerging market countries reached nearly \$2 trillion, four times their level of the early 1990s relative to GDP. Some commentators have argued that such massive reserves are prudent insurance against the increased volatility of capital flows accompanying financial globalization, but others claim they are excessive and that these countries could put the funds to better use toward their economic development. In the first article of this issue, Olivier Jeanne examines whether the risk of currency crisis or a sudden cutoff in lending justifies the level of reserves held by these countries.

Historically, various rules of thumb have been used to make this assessment. Jeanne is much more ambitious. He calibrates a formal model of optimal reserves holdings to country-specific empirical estimates of the parameters crucial to determining this optimum. He finds that the reserves held by most Latin American countries can be rationalized with this model, but the risks of a capital account crisis confronting the Asian emerging

economies in his sample are much too small to justify their current holdings. This conclusion has implications both for the need for further accumulation and for the management of existing reserves. In the final sections of the paper, Jeanne considers the opportunities for emerging market countries to diversify their investment, the implications of such a change in investment behavior for global financial markets, and possible international arrangements for using this wealth to more efficiently insure against future crises and to promote financial development.

Jeanne begins by documenting the recent buildup in international reserves by thirty-two emerging market countries. In the decade before 1990 these countries held reserves averaging 5 percent of their GDP, a ratio comparable to that in industrial economies. Since then, however, this average has grown by a factor of four, to roughly 22 percent of GDP by 2005. More than half of the dollar growth in these countries' reserves during the period took place in Asia following the 1997–98 crisis in that region. China's ratio of reserves to GDP is similar to that of other emerging Asian countries, but because of its size China now holds the largest stock of international reserves in the world.

Growth in a country's international reserves is necessarily financed by some combination of current account surpluses and capital inflows. Among the Asian emerging market countries, current account surpluses accounted for 63 percent of their reserves accumulation during 2000–05. In contrast, the Latin American economies ran sizable current account deficits, and capital inflows amounted to 137 percent of their reserves accumulation. The experience of the Asian economies dominates the overall picture: for all emerging market countries combined, current account surpluses account for almost 60 percent of reserves accumulation during this period. Jeanne notes that the flow of capital from developing to developed economies was equivalent to more than 40 percent of the current account deficit of the United States in the same period. He also documents various salient features of the gross capital flows underlying the net flows to and from countries in his sample. Reserves accumulation constitutes more than 60 percent of the gross foreign asset accumulation of these countries. In contrast, foreign direct investment (FDI) accounts for almost 70 percent of the increase in their gross liabilities. This investment pattern is not new and is evident in the balance sheets of these economies. The share of reserves in gross foreign assets is almost nine times as large in the emerging market countries as in the industrial countries, whereas the share of FDI in their liabilities is almost twice as large.

How much of the dramatic growth in emerging market countries' reserves can be explained as a prudent response to the increased risks associated with their increased trade and financial integration with the rest of the world? Jeanne reports that the recent buildup is not consistent with any of three conventional reserves adequacy measures: the ratios of reserves to imports, to short-term external debt, and to the M2 measure of the money supply. By any of these yardsticks, reserves ratios have increased enormously since 1990. For example, reserves in the entire sample of emerging market countries in 1990 were roughly in line with the informal rule that a country's reserves should equal or exceed three months of its imports, but reserves today are close to seven months of imports. The Greenspan-Guidotti rule, which states that reserves should equal or exceed short-term debt, was developed to better capture the risks stemming from the capital account after the crises of the 1990s. Reserves in emerging market countries have now risen to five times their short-term debt.

Jeanne notes that reserves also significantly exceed the predictions of regression-based empirical models that do a good job of predicting reserves holdings before the Asian crisis, based on a few key variables that should affect the desired level of precautionary reserves. He observes that such regressions fail to capture the profound effect that the 1997–98 crisis has had on countries' perception of risk. That crisis marked a watershed, after which emerging market countries became painfully aware that even sound macroeconomic policies did not insulate them from contagion and sharp reversals in capital flows. Jeanne points out that capital account instability and large capital account shocks, or "sudden stops," are relatively new phenomena in Asia, where capital outflows and reserves losses in the crisis period were unprecedented.

Neither the informal rules of thumb nor the earlier regression studies that found reserves to be adequate were based on an explicit model of precautionary reserves. Jeanne employs an empirical version of such a model to judge the extent to which reserves in his sample of countries are currently excessive. His analysis builds on a long line of literature on optimal reserves, inspired by a seminal 1996 contribution by Robert Heller, although the traditional models used an objective function only loosely related to domestic welfare. Jeanne specifies a model that is specifically tied to welfare but retains some of the simplicity of the earlier literature.

The model is of a small, open economy vulnerable to crises, where a crisis is defined by the loss of access to external credit and the associated

loss of output. Holding reserves involves an opportunity cost but decreases the probability of a crisis and dampens the damaging effect on consumption when a crisis does take place. To capture the intertemporal nature of reserves decisions in the simplest way possible, the model has only three periods. The model economy is populated by a representative consumer who holds a certain amount of foreign assets, or “sovereign wealth,” in period 0. This wealth can be invested either in liquid international reserves or in an illiquid asset that has a higher expected return over two periods but is not available for use until the last period. In exchange for lower returns, investing in reserves provides two benefits: it lowers the probability of a crisis occurring in period 1, and it insures consumption in that period against the risk of a crisis, an event that entails a large capital outflow and a loss of current output.

Jeanne captures the first benefit by making the probability of a crisis a decreasing function of the ratio of reserves to short-term debt. When he later calibrates the model, he uses a probit specification, implying that the function is the cumulative distribution of the normal function. The second benefit, that of insuring consumption when a crisis occurs, depends on the amount by which resources available for first-period consumption are reduced in the event of a crisis and on the curvature of the utility function. In a crisis the resources available to support consumption are reduced for two reasons: output is reduced by some fraction of potential output, and the country is unable to roll over the external short-term debt that is due, which reduces imports. This is what the literature calls a “sudden stop” of capital inflows. Subject to this constraint from reduced resources and the inability to borrow, the representative consumer decides what portion of available resources to consume in the first period and what portion to carry into the second period. In the absence of a crisis, the consumer simply decides on current consumption and on the change in debt and reserves, constrained by potential output. Reserves and debt have the same rate of interest between periods 0, 1, and 2, which for two periods is less than that on the illiquid investment. Welfare in period 2 is characterized simply as period-2 wealth, equal to period-2 income plus the balance of assets minus liabilities carried into the period and the interest accrued on that balance. This assumption implies constant marginal utility of consumption in period 2, making solution of the model quite simple.

The qualitative characteristics of optimal reserves holdings are intuitive, increasing in the size of the crisis shock and with the degree of risk aver-

sion, and decreasing with the opportunity cost of reserves. Jeanne shows that if the consumer has constant relative risk aversion and the probability of a crisis is independent of the quantity of reserves, optimal reserves are related to a simple function of the opportunity cost divided by the probability of crisis, and the size of the resource shock. In the event of a crisis all of the reserves are consumed. He observes that the optimal level could be either higher or lower than under the Greenspan-Guidotti rule: higher because of the value of smoothing a fall in output as well as the capital outflow, or lower because of the opportunity costs ignored in that rule. When an increase in reserves reduces the probability of crisis, the optimal level may even exceed the maximum size of the shock that could occur so that some reserves may be saved in the event of a crisis.

Jeanne next turns to calibration of the model, beginning with the role of reserves in crisis prevention. Analysts have distinguished two types of crises, currency crises and sudden stops of foreign capital inflows. Jeanne estimates the effect of reserves on the probability of each type of crisis, which he identifies by four different criteria, covering the broad range of definitions used by various authors. Currency crises are identified by the rate of domestic currency depreciation, or the sum of that rate and the rate of reserves loss, using various thresholds and timings. Sudden stops are identified by a substantial drop in net capital inflows relative to GDP, accompanied or not by declines in output and the current account, again with varying thresholds and timings. Jeanne estimates a univariate probit regression for each definition of crisis for each of five different measures of reserves adequacy.

Since Jeanne uses four different combinations of time and country dummy variables, each measure of reserves adequacy is used in sixteen probit regressions for each type of crisis. Several facts stand out from the results. First, the ratio of reserves to short-term debt is the best predictor of a currency crisis and is both statistically and economically significant in all sixteen probit regressions. In the range of estimated coefficients on this reserves ratio, doubling the Greenspan-Guidotti ratio of 1 reduces the probability of a crisis from roughly 10 percent to 6 percent, assuming a prevention benefit parameter of 0.3. However, there are diminishing returns: increasing the reserves ratio from 5 to 6 lowers the probability by less than 1 percentage point. Two other reserves adequacy measures, the ratio of reserves to imports and that of reserves to GDP, are significant in a majority of the regressions at the 10 percent level or better, but given the superior

performance of reserves relative to short-term debt, Jeanne uses the coefficients from those equations to calibrate his later simulations.

None of the measures of reserves adequacy perform well at predicting sudden stop. The ratio of reserves to M2 performs best: it is significant in four of the sixteen equations. Jeanne notes that these results are consistent with the empirical literature, which finds an ambiguous benefit of reserves in reducing the probability of sudden stops. Jeanne also offers an important caveat even to the conclusion that greater reserves decrease the probability of currency crises. Existing studies, including this one, cannot distinguish the prevention of a crisis from its mere postponement. And in some cases trying to maintain a high level of reserves in the face of a loss of confidence in domestic policies may actually hasten a crisis.

Even if reserves have no influence on the probability of a crisis, they have an important role in mitigating the effects of a crisis on output and domestic absorption. Jeanne describes some ways in which international reserves can be used to reduce the loss of output. One way is to intervene in foreign exchange markets to mitigate a rapid depreciation of the currency and its disruptive effects on balance sheets; another is to allow the monetary authorities to provide liquidity to domestic financial markets, the banking sector, and even exporters. He points out that providing such liquidity is especially valuable if there is significant dollarization of bank deposits and other domestic liabilities.

A fall in domestic absorption is an important indicator of the welfare costs of a sudden stop. Domestic absorption, the sum of consumption and domestic investment, is also the sum of domestic output, capital inflows, net income from abroad, and reductions in reserves. Reserves can be used to offset a drop in output or capital inflows, stabilizing absorption. Jeanne examines the behavior of each of these in a five-year window surrounding sudden stops for each country in his sample. Not surprisingly given the criteria he uses to identify sudden stops, net capital inflows drop almost 10 percent relative to GDP in the year of the sudden stop, regaining about 60 percent of that fall in the next two years. Strikingly, a decrease in reserves offsets most of the capital account reversal, with a sudden stop causing domestic absorption to fall by only 3 percent of GDP in the year before the stop. GDP falls only slightly in the year of the shock but grows more slowly than trend in the following two years. It seems likely that the observed fall in output following a sudden stop is less than would have occurred if reserves were not used for stabilization. Since, in his model, a

period corresponds to more than a year, Jeanne finds it reasonable to assume in his baseline calibration that the drop in output following a shock is 10 percent of GDP. But he also notes that in some cases the capital flight in sudden stops may itself reflect, rather than cause, a loss of reserves and the associated damage to the economy.

The opportunity cost of reserves—the difference between the return on reserves and that on alternative investments—is of first-order importance in determining optimal reserves, but it is difficult to measure. One approach is to assume that physical capital is the alternative investment. In that case, with an estimated average real rate of return of almost 8 percent a year for seventeen of the countries in Jeanne's sample, and with an average of U.S. real short-term interest rates taken as the return on reserves, the opportunity cost of reserves is 6 percent. But investment in physical capital carries more risk than many other long-term assets, and its return should be adjusted downward to compensate. Using the interest rate on central banks' domestic currency assets (that is, the cost of sterilization) gives a much lower number, and in the case of countries with very low interest rates, such as China recently, the implied opportunity cost would be negative. Jeanne argues that this measure ignores expected depreciation and that although this rate is relevant to central bank profits, it may be a poor measure of the opportunity cost to the country as a whole. A third measure of opportunity cost is the spread between the interest rate on external debt and the return on reserves; by this measure the annual opportunity cost averaged 8.4 percent during 2000–05 for emerging market countries in the aggregate, but it is much higher than the average in Latin America, where it grew rapidly during 2000–02 and then gradually declined. The level and variation of the Latin American spread undoubtedly reflect the risk of default. Jeanne observes that if the default risk premium is a fair reflection of the probability of less than full repayment, it should not be included in the opportunity cost.

Jeanne calculates the forgone earnings on the 2000–05 reserves holdings of emerging market countries using three different estimates of the opportunity cost. Using only the 2 percent term premium, the average annual costs for the Asian countries are 0.45 percent of GDP and a more modest 0.22 percent of GDP for Latin America, where reserves are much lower relative to GDP. Using the return to capital to capture the opportunity cost more than doubles these losses, to 1 percent of GDP in the case of the Asian economies.

How do the reserves actually held by emerging market countries compare with optimal levels? Jeanne begins by calculating the levels implied by

his optimizing model, using baseline values suggested by the above discussion: 10 percent of GDP for the size of a sudden stop and the loss of output, and a 3 percent annual opportunity cost of reserves. He considers two alternative measures of the probability of a crisis. The first assigns a fixed probability of 10 percent, consistent with the historical frequency of sudden stops for his sample of countries; in this case optimal reserves depend only on the ratio of the opportunity cost to such an exogenous probability. The second is an endogenous probability that depends on the ratio of reserves to short-term debt as estimated by Jeanne's univariate probit regressions. The only other parameter of the model, the degree of relative risk aversion, is set at 2, a conventional value. The benchmark calibration with a fixed probability of a crisis implies that optimal reserves are 7.7 percent of GDP. This is close to the average ratio of reserves to GDP observed in Jeanne's sample for 1980–2000, but significantly lower than in recent years, especially in Asia. If reserves are credited with a reduction in the probability of a crisis in line with the probit estimates, optimal reserves are much greater, over 20 percent of GDP.

Jeanne shows that the sensitivity of optimal reserves to these parameters is substantial. Doubling the baseline opportunity cost or halving the baseline probability cuts optimal reserves to approximately zero; cutting the opportunity cost to 1½ percent or raising the probability of crisis to 20 percent almost doubles optimal reserves. Similarly, doubling the size of a sudden stop or the output cost of a crisis significantly increases optimal reserves. An implication of these sensitivities is that fixed rules such as the Greenspan-Guidotti rule, although perhaps useful in particular situations, are unlikely to be appropriate for all countries at all times. But it is also clear that the recent levels of reserves in many emerging market countries can only be justified by parameter values that are far from the baseline assumptions.

This conclusion leads Jeanne to examine more closely the differences in countries' vulnerability to sudden stops and currency crises, as evidenced by the experience before the recent buildup of reserves. He runs separate probit regressions for the probability of sudden stops and crises for the period 1980–2000 on the pooled sample of countries, experimenting with a variety of economic indicators as explanatory variables. In his preferred specifications, the deviation of the real exchange rate from trend is significant in both sudden stops and currency crises, with a weak currency lowering the probability. The current account and the public debt, both relative



to GDP, and foreign liabilities relative to money are also significant for sudden stops. Consistent with the earlier results, the ratio of reserves to short-term debt is highly significant for currency crises. All these results are robust with respect to the inclusion of fixed country or time effects. Tracking countries, Jeanne finds large differences in the probability of crises over time. The probability of a sudden stop in any given year of the early 1980s reached 14 percent for both the Asian and the Latin American countries. It then declined, more rapidly for the Asian economies, to a common low of about 4 percent around 1990. Since then the experience of the two regions has diverged strongly. Except for an upward blip, the probability of a sudden stop stayed low for the Asian economies in the 1990s, but it rose rapidly in Latin America, with the GDP-weighted average returning to the heights of the early 1980s. The probability of currency crisis followed the same pattern until the 1990s. At the beginning of that decade, the probability rose rapidly in Latin America but then slowly drifted down, paralleling the pattern in Asia but always substantially higher. The estimated probabilities for China are noteworthy: 2.7 percent for a sudden stop, and less than 0.2 percent for a currency crisis in 2000.

The model's estimate of optimal reserves for insurance against sudden stops, using the country-specific probabilities and the baseline values of the other parameters, confirms the view that the aggregate reserves of the emerging market countries in Jeanne's sample are far more than they need for that purpose, and that this is almost entirely a reflection of the large reserves holdings of the Asian economies. According to the model, optimal reserves in 2000 for the Asian countries in the sample totaled \$24 billion, whereas actual reserves were \$406 billion. The optimal level according to the model is actually zero in China, Korea, and Malaysia, because the probability of a sudden stop for those countries is so low. To rationalize the levels actually held in these countries would require extremely high risk aversion, a low opportunity cost of reserves, or an output loss in the event of a sudden stop in excess of 30 percent. By contrast, the observed level of reserves is actually slightly below the estimate of optimal reserves for the Latin American economies.

Higher reserve levels would be called for if reserves lower the probability of a crisis. Using the estimate of this effect from the probit regressions explaining currency crises, Jeanne calculates the expected benefit, net of costs, for the reserves buildup of 2000–05 assuming a 10 percent output loss in the event of crisis and an opportunity cost of 3 percent. For the Asian

emerging market countries as a group, he finds the costs to be more than five times the benefit. The probabilities of a crisis were already so low in 2000 that the further reductions in probability are trivial. But in Latin America, where the probabilities were significant in 2000, the benefit of additional reserves outweighs their cost. Jeanne cites Mexico as an example, where doubling the reserves held in 2000 would have reduced the probability of a crisis that year from 9.6 percent to 5.6 percent.

Given the apparent failure of the precautionary model to explain the accumulation of reserves by the Asian economies, what is likely to be the true explanation? Jeanne's main alternative explanation is that the accumulation is the unintended consequence of a mercantilist policy of pursuing large current account surpluses. Such a policy requires that the central bank accumulate reserves to avoid an appreciation of the currency, together with capital controls and financial repression to avoid inflation. In support of this explanation, he reports cross-country correlations showing that, during 2000–05, reserves accumulation was associated with current account surpluses and unrelated to changes in gross external liabilities. The change in the ratio of reserves to GDP is also correlated with capital account restrictions. These correlations are hard to reconcile with a precautionary explanation, for which the risk of capital outflows is central.

Jeanne notes that a mercantilist strategy involves, at least implicitly, an accumulation of claims against the rest of the world, but also that it need not involve levels of highly liquid international reserves beyond those required for precautionary reasons. He observes that emerging market countries have started to mitigate the cost of holding excess reserves by transferring a portion of their foreign exchange reserves from the central bank to new “sovereign wealth funds,” which are not unlike the natural resource–based stabilization funds set up by a number of commodity exporters. He cites the sovereign wealth funds of Korea and China as recent examples, and he reports forecasts that the holdings of such funds, already more than \$2 trillion, could exceed \$12 trillion by 2015.

Would such a development lead to disruptions in exchange rates and in the relative prices of financial assets in global markets? Jeanne documents the significant difference in size between the portfolios of the official and the nonofficial sectors and considers, as an example, the likely effect of shifting \$450 billion of the roughly \$2 trillion in foreign assets into a “world” portfolio. His back-of-the-envelope calculations show that although such a shift plays against the dollar, and in particular against fixed-

income dollar assets, changes in net demands would be relatively small fractions of outstanding stocks, with moderate effects on prices and exchange rates. But Jeanne also points out that a substantial portion of outstanding U.S. debt is held by the foreign official sector, and that the effect on the interest rate the U.S. government pays may be nonnegligible, depending on the substitutability between U.S. Treasury securities and other forms of dollar debt in the portfolios of global investors. He also cautions that if the pace of diversification is rapid, the effects might be much larger than would be expected in the long run.

Jeanne argues that the abundance of reserves held by emerging market countries reduces the need for collective insurance provided at the global level by the International Monetary Fund (IMF) and various regional insurance arrangements, noting that the increase in reserves in the Asian emerging market countries over 2000–05 was more than four times the IMF's usable assets at the end of 2005 and more than twenty times the bilateral swap agreements to date under the Chiang Mai Initiative launched by a group of Asian countries in 2000. He suggests various other potential uses of the large accumulated stocks, such as collective insurance for risks other than capital account crises; such risks might include severe shocks to countries' terms of trade or output, as well as natural disasters. He also suggests that sovereign wealth could be used to promote the development of regional financial markets, which might in turn serve as a catalyst for private investment in Asian issues and help enable emerging market countries develop investment instruments with long maturities, denominated in domestic currency, which would be safer for borrowers.

Finally, Jeanne observes that even if the recent high rate of reserves accumulation were to abate, the public sectors of many emerging market economies will have to manage stocks of foreign financial assets of unprecedented size for some time to come. Although it may be a challenge to ensure that the diversification of those assets avoids large or abrupt changes in the relative prices of financial assets or exchange rates, the growth in sovereign wealth of these countries provides many new opportunities to benefit the developing world.

THE ACCELERATION OF PRODUCTIVITY GROWTH was an economic hallmark of the second half of the 1990s. After twenty years during which labor productivity growth averaged only around 1.5 percent a year, it rose at an average annual rate of over 2.5 percent between 1995 and 2000. That

development had important consequences for the nation's economic performance, not only through the direct effect of raising national income at given levels of employment but also by its effect on the conduct of monetary policy: then-Federal Reserve Chairman Alan Greenspan's early recognition that productivity was speeding up is credited with the continuation of monetary accommodation, which in turn permitted a strong expansion of employment and reduction of unemployment for another five years. Most analysts at the time attributed an important part of the productivity acceleration to the continuing boom in information technology (IT) hardware and software. When IT investment declined sharply after 2000, that diagnosis of the boom period implied that the productivity trend would also slow. It did not. In the second article of this issue, Stephen Oliner, Daniel Sichel, and Kevin Stiroh use the surprising developments of both the pre- and post-2000 periods to refine our understanding of productivity growth, project the likely future path of productivity, and assess the uncertainty surrounding that projection.

As the framework for their analysis, the authors extend the growth accounting model they developed previously to study the role of IT capital in the nonfarm business sector. As in standard growth accounting, growth in aggregate labor productivity in their model can be decomposed into the contribution from the increase in capital per hour worked, the contribution from improved labor quality, and growth in multifactor productivity (MFP), each measured as an aggregate average. The authors extend this standard model in several directions. To focus on the role of IT, they disaggregate total output into five IT-producing sectors and a sixth, non-IT-producing remainder. Data for four of the IT-producing sectors—computer hardware, software, communications equipment, and semiconductors—are taken from the National Income and Product Accounts (NIPAs). Exports of semiconductors are classified as final output in the NIPAs, and the remaining output in this sector is treated as intermediate inputs to the others. The fifth IT sector produces intangible IT capital; since this output is not measured in the NIPAs, the authors introduce their own estimate. This intangible output is treated as investment, whose services are inputs to future production. The authors also take account of cyclical influences on productivity growth (changes in factor utilization) and of the adjustment costs associated with new capital goods; both these factors could have had significant effects on observed productivity movements over the past decade.

Using their six-sector disaggregation, the authors estimate the effect of the business cycle on measured MFP growth, which they take as propor-

tional to the length of the workweek, and of the costs of adjustment, which they take as proportional to investment. Each of these effects drives a wedge between measured MFP growth and the true pace of improvements in technology and similar underlying effects affecting long-run growth. Lengthening the workweek is expected to boost measured MFP growth as firms get more output from their capital and labor, and raising investment is expected to lower MFP growth as firms divert resources from producing market output to installing new capital. Aggregate MFP growth and the aggregate effects of these adjustments are given by the share-weighted sum of the sectoral MFP growth rates, where the weights are shares of each sector's gross output in aggregate value added. Because not all the data required for a thoroughgoing bottom-up estimation of these sectoral relations are available, the authors rely in part on estimates made at the aggregate level and distributed to the sectors, and on estimates of certain parameters from earlier research by themselves and others.

The authors first undertake a growth accounting that does not incorporate the estimates of intangible IT capital. This accounting, using only the familiar standard NIPA data, shows how the effects of the business cycle and adjustment costs on investment influence the accounting for the productivity surge in the late 1990s and after 2000. Average annual growth of labor productivity picked up from about 1.5 percent during 1973–95 to about 2.5 percent in the following five years, and to 2.9 percent during 2000–06. Of the initial 1.0-percentage-point speedup, 0.6 percentage point reflected the quickened growth of IT capital, and another 0.5 percentage point came from the acceleration of MFP in IT-producing sectors, so that IT accounted for slightly more than the entire acceleration in labor productivity. The cyclical and adjustment cost effects were each noticeable. Compared with the pre-1995 long-run average, the strong expansion added 0.16 percentage point to annual MFP growth during 1995–2000, and the adjustment costs associated with high investment reduced MFP growth by 0.15 percentage point in this period. The two effects having opposite signs, their combined effect was negligible.

The accounting for the further acceleration of MFP after 2000 is quite different. The cyclical and adjustment cost effects are again substantial, but of the opposite sign from the previous five years, and again roughly cancel each other out. But compared with 1995–2000, in the post-2000 period IT capital deepening and MFP growth in IT-producing industries are both slower and together contribute 0.7 percentage point a year less to labor pro-

ductivity growth, whereas in the non-IT sectors capital deepening and MFP growth both contribute substantially more, accounting for more than the entire further acceleration of labor productivity in this period. Averaged over the whole 1995–2006 period, total IT contributions account for roughly two-thirds of the acceleration in labor productivity. However, comparing just the post-2000 period with 1973–95, the total contribution of IT accounts for only about a quarter of the 1.4-percentage-point acceleration.

To better understand these developments and how they may inform projections of future productivity growth, the authors next construct a measure of intangible IT capital and introduce it into their growth accounting. In earlier work, Carol Corrado, Charles Hulten, and Sichel argued that any intangible asset that provides services beyond the current period should be considered as part of the capital stock, with its production counted as investment in current-period output. Using a wide range of data sources, they estimated that such intangible investment, ranging across such categories as research and development, computerized information, and firm-specific organizational capital, totaled roughly \$1 trillion a year over 2000–03, or almost as much as business fixed investment.

To obtain up-to-date estimates of the intangible capital that can be associated with IT, the authors turn to a model developed by Susanto Basu and coauthors and based on the idea that intangible capital is a complement of IT capital as measured. That model specifies a function in which tangible and intangible IT capital are treated as complementary factors, with firms optimizing the ratio of the two given their relative costs. Total output then depends on this combined IT input, labor, and non-IT capital. The optimization of tangible and intangible IT capital inputs assumes an elasticity of substitution between them of 1.25, a value that best approximates the trend in the income shares of the two estimated by Basu and coauthors. The estimates are anchored to the values in the original Corrado-Hulten-Sichel data by making the average income share of intangible capital over 1973–2003 equal to the income shares of those intangibles that are IT related in those authors' original estimates. The present authors thus generate a time series for intangible capital and, using a standard perpetual inventory relation, a series for intangible investments associated with it. Although intangible capital investment is linked to IT investment in these constructed series, the price declines in IT capital cause its user cost to trend lower, leading intangible capital to grow less rapidly than IT capital in all periods. This difference is more pronounced after 1995 than before.

Incorporating intangible capital raises both the input and the output sides of the production accounts compared with the standard NIPA treatment. And the augmented framework changes both the pattern of labor productivity over time and the accounting for its growth. In the baseline case, which corresponds to the timing assumptions in the model of Basu and coauthors, labor productivity growth speeds up by 1.4 percentage points in 1995–2000, even more than the 1-percentage-point speedup in the NIPA data, reflecting both the more rapid growth of output when investment in intangible capital is added to the NIPA estimates, and greater capital deepening. Labor productivity then grows by 2.4 percent a year on average in the following period, or 0.4 percentage point less than in the NIPA estimates. Including intangible investment in output has a noticeable effect on MFP growth. Whereas the NIPA data implied speedups in MFP of 0.7 percentage point after 1995 and a further 0.6 percentage point after 2000, the augmented data imply a 0.9-percentage-point speedup after 1995 and only a negligible further change after 2000. The contribution of capital deepening from intangible capital speeds up by 0.1 percentage point after 1995 and then declines by 0.4 percentage point after 2000.

The baseline estimates assume a contemporaneous relation between the growth of IT capital and the growth of intangible capital associated with it. To check the robustness of their results, the authors present three alternative series for intangibles that allow for a smoothing of their relation to IT capital and, alternatively, a one-year lag of intangibles growth behind IT capital growth. The smoothing is accomplished by using a three-year or a five-year centered moving average for the growth rate of IT capital and its user costs in the equations used to generate estimates of intangible capital. In each of these alternatives, after 2000 there are noticeable declines in intangible capital investment and slight declines in the capital services from intangibles. The range of estimates across these four series indicates that the results are quite robust to the variations in timing that are considered. Annual labor productivity growth speeds up by between 1.4 and 1.1 percentage points after 1995 and slows by between 0.5 and 0.1 percentage point after 2000. The contribution of intangible capital deepening speeds up by 0.1 percentage point after 1995 in all alternatives and slows down by between 0.3 and 0.4 percentage point after 2000. And the contribution of aggregate MFP to labor productivity speeds up by between 0.7 and 0.9 percentage point after 1995 and by a further 0.04 to 0.4 percentage point after 2000.

The authors turn next to productivity at the industry level. They rely, alternatively, on annual measures of value added and gross output from the Bureau of Economic Analysis (BEA) and combine each of these with data on hours from the Bureau of Labor Statistics (BLS) to calculate two alternative measures of average labor productivity by industry. The value added-based measures of productivity by industry combine directly to an aggregate measure of labor productivity for all private industry. The gross output-based measures require an allowance for changes in intermediate inputs to be reconciled with that aggregate. The authors also calculate a proxy for capital services, using BEA data on nonresidential capital that is disaggregated by type and industry. The available data permit calculation of average labor productivity growth for sixty industries for the period 1988–2005. The value added-based estimates allow the authors to form a panel of fifteen broad sectors and a private industry aggregate made up of their sixty industries. Labor productivity growth rates for this aggregate differ somewhat from those for the more inclusive BLS data covering the private nonfarm business sector. But the authors capture the pickup in productivity growth after 1995 and the smaller further pickup after 2000 that characterizes the BEA aggregates. Among the fifteen broad sectors, productivity speeded up after 1995 in eight, which together account for 73 percent of value added. After 2000, productivity continued growing faster than it had during 1988–95 in six of these sectors: durable goods; wholesale trade; retail trade; professional and business services; education services, health care, and social assistance; and agriculture. Each of these sectors is a relatively large part of the economy, except agriculture, which is very small. And productivity accelerated sharply in the information industry, leaving its productivity growth over the entire 1995–2005 period well above its previous rate.

To examine the contribution of IT, the authors classify individual industries according to whether they are IT producing, IT using, or neither. IT-producing industries are the four industries classed as such by the BEA, which together produced nearly 5 percent of aggregate value added in 2005: computer and electronics products, publishing including software, information and data processing services, and computer system design and related services. IT-using industries are defined as those whose use of IT exceeds that of the median industry (excluding IT producers). The authors then decompose the changes in labor productivity for their private industry aggregate into the direct contributions of IT-producing industries,



IT-using industries, and other industries, and the contribution coming from reallocations across industries. Using the industry value-added data, they measure the contribution from input reallocations by changes in hours worked across industries. This contribution is positive when hours grow relatively more in high-productivity industries. When the industry gross output data are used, the contribution from reallocations includes, in addition, changes in intermediate inputs across industries. This contribution is positive when gross output is rising faster than purchased material inputs.

Both the gross output-based and the value added-based measures assign a dominant role to IT in the acceleration of labor productivity after 1995. Of the 0.96-percentage-point acceleration in the aggregated industries measure during 1995–2000, the analysis using the gross output data attributes 1.45 percentage points to the direct effect of productivity acceleration in IT-producing and IT-using industries, and a small decline to other industries. The reallocation effects are  $-0.48$  percentage point from intermediate inputs and 0.13 percentage point from hours. The value-added data show direct effects of 1.16 percentage points from IT industries and a small decline in the others, along with a 0.13-percentage-point hours reallocation effect. In both cases the large direct IT contribution comes mainly from IT-using industries. The small further acceleration of 0.32 percentage point in aggregate productivity after 2000 is accounted for very differently in the two data sets. The growth output data attribute  $-0.94$  percentage point to the direct effects of industry productivity growth, with IT-producing, IT-using, and other industries all contributing to the slowdown. The hours reallocation contributes 0.31 percentage point and the intermediate inputs reallocation contributes 0.94 percentage point. The value-added data show no change in direct effects, with a small deceleration in IT producers offsetting a small acceleration in IT users. In both the gross output and the value-added data, over the entire post-1995 period, the direct contribution of IT-using and IT-producing industries accounts for most of the speedup in aggregate labor productivity; the direct contribution of other industries is slightly negative.

The authors perform regressions on their industry-level data to further explore the acceleration of productivity. They construct two alternative measures of IT intensity by industry: one is the share of IT capital services in total capital services, and the other is a dummy variable equal to 1 if an industry's use of IT capital as a share of its total capital is above the median for all industries. Using the gross output data, they find that IT-intensive

industries, by either definition, experienced faster productivity growth after 1995, but they find no further significant effect on productivity growth starting in 2000. The value-added data yield less significant results.

Some analysts have speculated that productivity growth after 2000 may have been driven by unusual pressures on firms to restructure so as to cut costs in a more competitive environment. To explore this and related hypotheses, the authors examine the relation between the change in industry profit share from 1997 to 2001 and the growth of hours and labor productivity from 2001 to 2004 across their sixty-industry sample. On average, industries with below-median changes in profitability experienced 2 percentage points slower hours growth and 3 percentage points faster productivity growth than those with above-median changes in profitability. The authors go on to run separate regressions using the change in profits between 1997 and 2001 to explain growth rates in hours, intermediate inputs, labor productivity, and output over 2001–04, the period of extremely rapid productivity growth. The regressions include a number of additional variables to control for demand effects, longer-run trends, and IT intensity. With either the gross output or the value-added data, they find significant negative effects of changes in profits on growth rates of hours and significant positive effects on labor productivity, no effects on intermediate inputs or gross output, and marginally significant effects on value added. These results are consistent with the idea that competitive pressures and restructuring help explain the post-2000 productivity gains. The regressions also reveal no significant effects from IT intensity on productivity or output, which the authors interpret as support for the idea that IT was not an important factor in the post-2000 changes in productivity.

Following their econometric analysis, the authors go on to put recent productivity gains in historical perspective. They show that labor productivity growth in the private sector has averaged 2.2 percent a year since 1909. Dividing that nearly century-long period into several subperiods, they show that annual productivity growth averaged 2.9 percent during 1950–73, the “golden era” of productivity growth, and 2.8 percent during 1995–2006. It averaged around 1.5 percent a year in 1909–28 and 1973–95, the two remaining periods. Thus the recent past is one of unusually rapid growth relative to historical experience.

Looking ahead, the authors report a range of steady-state growth rates that would be consistent with their growth accounting model under various parameter assumptions. The key parameters affecting aggregate labor

productivity are the rate of improvement in labor quality and the rate of advance in technology outside the IT-producing sectors. Capital deepening is endogenously determined by the rate of growth of MFP outside IT. Using bounds for these parameter values based on historical experience, they calculate upper and lower bounds for projected steady-state annual productivity growth of 3.1 percent and 1.5 percent, respectively. The center of the range is  $2\frac{1}{4}$  percent. At the lower-bound estimate, capital deepening contributes 0.75 percentage point and MFP growth 0.56 percentage point to productivity growth. At the higher bound, capital deepening contributes 1.39 percentage points and MFP 1.55 percentage points. The authors show that their central projection of  $2\frac{1}{4}$  percent is very near the projections of a number of other studies, and close as well to their own ten-year projections using a Kalman filter model that allows estimation of a stochastic process for productivity growth. But they stress that the uncertainty around their own projections, and presumably around the others as well, is considerable.

ONE PERSISTENT DIFFERENCE BETWEEN the poor and the rich countries of the world is that rich countries have more real capital per worker. All else equal, the marginal return to capital should be lower where capital is more abundant. It follows that investment funds should profitably flow from richer to poorer countries. Since Robert Lucas pointed out in 1990 that it typically does not, analysts have tried to understand the reasons for this seemingly perverse phenomenon. Some invoke financial market inefficiencies that interfere with the productive investment of capital; others point to political or policy instabilities in poor countries that make investments in them exceptionally risky. Recently, Pierre-Olivier Gourinchas and Olivier Jeanne have expanded the Lucas paradox by showing that capital flows to nonindustrialized countries have not been concentrated among those growing relatively fast, even though successful growth would presumably indicate that capital is productive in those countries and that they are relatively creditworthy. Understanding why foreign capital inflows have not been associated with faster growth and, more generally, how international capital flows influence the development process are key issues for developing countries, which must decide whether and how to open their economy to global investors. In the third article of this volume, Eswar Prasad, Raghuram Rajan, and Arvind Subramanian take a further look at these issues.

The authors first lay out several broad trends in international capital flows over the past few decades, focusing on national current accounts, which are the counterpart to the broadest measure of net foreign finance. A current account surplus equals the sum of net outflows of private and official financial capital, net accumulation of reserves, and net errors and omissions. The last of these reconciles direct measures of current account flows, such as exports and imports, with direct measures of capital flows, such as equity and debt purchases, and is generally thought to reflect capital moving through unmeasured channels. Thus a positive number for errors and omissions is customarily treated as a capital outflow, just as a purchase of foreign securities or a deposit in a foreign bank would be. The authors show that, at least since the early 1990s, the sum of current account surpluses in economies with surpluses has risen sharply as a share of world GDP. This might have been expected given the growing globalization of international finance. What is more striking is that, since the early 1980s, the average relative income of the countries with current account surpluses (the capital exporters) has trended down even as the average relative income of countries with current account deficits has trended up. Much of this difference in income trends remains even when the United States and China, two large outliers, are omitted from the calculations.

Going a step beyond the simple association of growth and current accounts, the authors divide their sample of some sixty nonindustrialized economies into four groups depending on whether a country's current account balance, as a ratio to GDP, was above or below the group median and whether its ratio of real investment to GDP was above or below the median. Averaged over 1970–2004, higher investment ratios and larger current account surpluses are both associated with faster growth in GDP per capita. Among countries with high investment ratios, those with less reliance on foreign capital grew faster, by an average of 1 percentage point a year.

To examine these connections more rigorously, the authors turn to regressions using annual data covering 1970–2004 for their sample of nonindustrialized countries. They first present cross-sectional regressions that, following the work of Barry Bosworth and Susan Collins, explain purchasing power–adjusted growth rates in GDP per capita over this period with the ratio of the current account to GDP and a number of control variables: the logarithm of initial GDP per capita, initial-period life expectancy, trade openness, the fiscal balance, a measure of institutional quality, and

dummy variables for oil exporters and countries in sub-Saharan Africa. Their baseline regression finds a significant positive association of growth with current account balances, indicating that countries that rely less on foreign finance grow faster. Among the control variables, institutional quality, initial GDP per capita, and life expectancy are all significantly associated with growth. Omitting data for three countries that are outliers has no effect on these results. But when countries receiving aid flows that average more than 10 percent of their GDP are also omitted, the coefficient on the current account balance doubles, and the institutional quality variable becomes highly significant. This indicates that the results are not driven by large aid recipients. When accumulated foreign assets or liabilities are substituted for the current account balance, they are usually not significant, and the regressions do not explain growth quite as well.

It might be thought that slow growth reflects low investment and that current account deficits reveal a lack of domestic resources that constrains investment. However, when the ratio of domestic investment to GDP is added to the regression with the current account, investment is not significant and has little effect on the other variables. The authors take this as evidence that the correlation of the current account balance with growth does not arise because investment is constrained by a lack of domestic resources. By contrast, when the ratio of domestic saving to GDP is added instead of the investment ratio, it is highly significant and the current account balance loses significance. The authors interpret this as evidence that the positive association between current account balances and growth stems largely from a positive relation between domestic saving and growth.

Some further regressions investigate the robustness of these basic results. When estimated over 1985–97, a period when international finance grew rapidly and international financial crises were relatively unimportant, the association of growth with the current account balance is even stronger. When industrialized countries are added to the sample (and distinguished by a dummy variable), the coefficient for the nonindustrialized countries is unaffected, whereas the dummy for industrialized countries has a significant net negative coefficient, indicating that current account balances are negatively related to growth in these countries. The same is true of the countries in transition from socialism when they are added to a regression estimated over 1990–2004. The authors also add the share of the working-age population in total population, an exogenous demographic variable that is expected to be associated positively with saving. Its regres-

sion coefficient is significantly positive, and its inclusion reduces the current account balance variable to insignificance. This, the authors reason, supports the idea that the association of growth with the current account balance arises from the relation of domestic saving to growth. Furthering this argument, they average the experience of countries in the years surrounding spurts in their growth rates and show that such growth spurts have led to higher saving, a result consistent with a role for habit formation in consumption.

The authors extend their analysis with panel regressions using five-year averages to capture the changes in a country over time, along with many of the same controls and specifications as in the cross-sectional regressions. In these regressions the estimated effect of the current account balance reaches statistical significance only when countries receiving substantial aid are dropped from the sample, along with the three large outliers. The investment-GDP and saving-GDP ratios are each significant when added separately to the equation, as is the working-age share of the population. For industrial countries the relation between the current account balance and growth is again negative. The panel results thus largely support the cross-sectional results, although with weaker evidence of a positive relation between the current account balance and growth for nonindustrial countries.

The interrelations among the macroeconomic variables being analyzed make it difficult to establish causality. The current account balance, investment, and saving are all endogenous, with shocks to any one potentially affecting the others. Because developments in individual industries have only modest effects on aggregates, estimation with industry-level data is less subject to such endogeneity problems. The authors therefore turn to industry-level data to examine some hypotheses about their aggregate findings. They first examine whether manufacturing industries that are relatively dependent on outside financing for investment, as opposed to financing from internal cash flows, grow faster when foreign capital is more available. Following earlier work by Rajan with Luigi Zingales, they run separate regressions explaining industry growth for each of five measures of foreign capital inflow: the stock of inward foreign direct investment (FDI), the stock of FDI and portfolio investment combined, the net flow counterparts of these two measures, and the current account deficit, each relative to GDP. (Each of these measures is interacted with a measure of the industry's dependence on outside financing.) In separate regressions the

above openness measures are also interacted with an indicator of level of financial development. All regressions include the relative size of the industry in the country's manufacturing value added, and all are cross sections run separately for average values of the 1980s and the 1990s. The regressions include country and industry fixed effects, in effect explaining within-country differences in industry growth rates by the industry's dependence on outside financing and the country's use of foreign financing.

The regression results generally support the idea that countries with greater financial development respond differently to external finance. This result is robust for the 1990s, where all five measures of capital inflow interacted with dependence on external finance are significantly negatively related to industry growth for less financially developed countries, and positively related for more financially developed countries. Roughly similar results are obtained for the 1980s. The findings are again the same in panel regressions that include variations across the two decades. The authors discuss in some detail the various ways in which lower financial development might interact with other characteristics of economies to affect growth. Although they can only conjecture about specifics, they reason that improving financial development brings substantial benefits.

The authors recognize that reliance on foreign capital may also affect growth through entirely different channels that change a country's international competitiveness. To explore this possibility, they use regressions to estimate the relations among currency overvaluation, capital flows or stocks, and growth, using a measure of overvaluation developed by Simon Johnson, Jonathan Ostry, and Subramanian that adjusts exchange rates for purchasing power parity. Cross-country regressions using data averaged over 1970–2004 explain overvaluation with the working-age population and various measures of capital stocks and flows. All show a positive association between the presence of foreign capital and overvaluation. The best-fitting regressions are for net FDI flows and net private inflows. And by far the best fits are obtained when industrialized economies are distinguished from non-industrialized economies. In these regressions, capital flows to industrial economies are more or less unrelated to overvaluation, whereas the relation is significantly positive for nonindustrialized economies. Additional regressions show that overvaluation has a significantly negative relation to growth. Finally, regressions using industry data that distinguish industries by their potential for exporting show that the negative effects of overvaluation on an industry's growth depend on that potential.

The authors' conclusions are modest, reflecting the uncertain causality behind the relations they examine and the many factors that influence growth. It is clear that nonindustrial countries that have relied relatively heavily on foreign capital have not grown faster than those that have not. And their results suggest that the less successful developing countries have limited capacity to utilize foreign resources, either because their financial markets are underdeveloped, or because they are prone to overvaluation in the face of rapid capital inflows, or both. Although the authors' results caution against encouraging some forms of capital inflow to reduce the risk of overvaluation, financial openness may be needed to spur development of the domestic financial system. Given this dilemma, the authors reason that policymakers need to take account of country-specific issues when opening the current account and need to be creative and flexible in managing it.

A CENTRAL FEATURE OF the Job Growth and Taxpayer Relief Reconciliation Act of 2003 was a large cut in the tax on dividends, lowering the maximum rate from over 38 percent to 15 percent. A natural prediction is that these lower rates will lead to an increase in dividends over time, both by increasing the distribution of current earnings and by causing a shift from debt to equity financing of corporations. A number of studies have attempted to quantify the effects of the tax cut on dividend payments, but relatively little recent work has been done on how dividends themselves affect consumption. Estimating these effects from aggregate time series is confounded by the multitude of factors correlated with dividends that plausibly influence consumption. In the fourth article of this issue, Malcolm Baker, Stefan Nagel, and Jeffrey Wurgler utilize cross-sectional data on households to infer the effect of dividends on consumption, how those effects differ from the effects of capital gains, and by implication how increases in dividends resulting from tax cuts are likely to affect the saving available for investment.

The authors begin by examining the cross-sectional relationship between dividends and consumption using data from the Consumer Expenditure Survey (CEX) for 1988–2001. Their sample includes interviews of several hundred households per year, each household for five consecutive quarters. The surveys elicit information on households' consumption, income, wealth, and financial returns as well as important demographic characteristics. The authors restrict their attention to households with positive wealth, nonzero holdings of stocks or mutual funds, and unchanged marital



status and household size over the period they were interviewed. Households in this sample had, on average over 1988–2001, after-tax income of \$56,566, total consumption expenditure of \$48,076, and nondurables consumption of \$15,042, all measured in December 2001 dollars. Their financial wealth averaged \$67,700 and constituted about a third of their average total wealth. Dividends and interest income averaged \$935 and \$1,264 a year, respectively. As expected, the distributions of financial assets and returns are very unequal, skewed strongly to the right. For example, mean financial wealth is roughly 1.75 times median wealth, and median dividend income is zero, even though all of the households in the sample hold some stock, either directly or through mutual funds.

Variation in household consumption and dividend income in a cross section of households undoubtedly reflects a variety of households' non-financial characteristics in addition to income and other financial variables. The authors are agnostic about what specific model is the most accurate representation of consumption behavior and instead attempt to isolate the causal effect of dividends by controlling for a wide variety of demographic characteristics. In their basic specification explaining the level of either total or nondurables consumption, the demographic controls include age of household head, family size, and education. These variables are allowed to enter linearly, quadratically, and in interactions with each other and the financial variables. The authors also include a set of year-month fixed effects to absorb seasonal variation in consumption as well as variation in macroeconomic factors affecting the consumption of households sampled at different times. In their words, "In the end the levels specification boils down to asking whether two consumers in the same financial situation, with similar income, similar household characteristics, and similar total return on financial assets, but different *compositions* of total returns across dividends and capital gains, have different consumption."

The authors find that an extra dollar of dividend income has substantial and statistically significant effects on consumption: the marginal propensity to consume (MPC) is roughly three-quarters for total consumption and one-sixth for nondurables consumption. The estimated effects of capital gains on consumption (given by the coefficient on total returns, since dividends are entered separately) are near zero and insignificant. In specifications in which interest is lumped with dividends and included in total returns, the MPCs are somewhat smaller than those for dividends alone but also highly significant. The authors include a dummy variable for non-dividend-

receiving households to ensure that their results are not driven by the large number of zero-dividend observations. They experiment with controlling for age, noting that dividends are a higher percentage of income for older households, perhaps because those households are more likely to be in retirement. The MPC out of dividends for total consumption among those over 65 is found to be substantially lower than for those younger than 65, but this result is of borderline statistical significance.

The authors recognize that the correlation of consumption with dividends across households may reflect in part the effect of some unmeasured differences in households' characteristics that influence both consumption and dividends, or it may reflect a reversal of causation, with households with high consumption needs choosing high-yielding stocks. Although the authors believe the controls they use should do a reasonable job of minimizing this bias, they recognize that it is difficult to fully rule out some remaining, unobserved differences between households that hold high- and low-dividend-paying stocks.

Differencing the consumption levels equations reported above would remove any fixed household effects correlated with the level of dividend income, providing a check on whether the estimated MPCs in those equations are biased by endogeneity. However, the limited information on the change in a household's consumption and asset returns in the CEX makes it impossible to run a true differenced form of the equations. The CEX survey does provide information for each household on the difference between quarterly consumption at the second and the fifth interviews, and the difference in dividends in the twelve months preceding each of those interviews. The authors therefore regress the change in consumption between these two quarters on the change in dividends, the level of total returns, the age of the household head, and other demographic variables.

In the equations explaining total consumption, the resulting MPC is 0.057. Assuming that the change in annual dividends affects four quarters of consumption equally, not just consumption in the quarters reported, this translates to an MPC of 0.228, roughly a third of the MPCs found in the levels equations. The authors see this reduction as consistent with some upward bias in the levels equations, but they note that it could also reflect the noise introduced through the imperfect matching of dividends and consumption measurement periods. If lagged consumption is included in this equation, the MPC for total consumption is substantially greater, but the coefficient on lagged consumption is negative, suggesting that more than 60

percent of the increase would be reversed in the following period. Since this is a cross-sectional regression, this negative coefficient could arise if wealthy households with large financial asset holdings and large absolute changes in dividends are relatively less responsive to changes in dividends. The authors add that the lagged level of consumption may be absorbing some of the noise resulting from the imperfect matching of the dividend and consumption measurement periods.

For nondurables consumption the estimated effect of the change in dividends is significant only when the level of lagged consumption is included; the MPC is again about a third of that in the corresponding levels equation and includes a large negative coefficient on lagged consumption. In the log difference equations, which should be free of scale effects, dividend changes are entered simply as a dummy variable, equal to 1 when positive. The estimated effects of dividend increases are large: households with an increase in dividends have a 7- to 8-percentage-point greater increase in total consumption than those without an increase. The corresponding point estimates for growth in nondurables expenditure are also substantial, although not statistically significant.

To supplement these results, the authors turn to a second micro dataset containing portfolio information for a large number of household accounts at a large discount brokerage firm. Although this dataset does not include data on consumption, assets other than those held in the brokerage account, or household demographic characteristics, it provides much more accurate and detailed information about dividends, capital gains, and stock transactions. And although the data do not allow estimation of the effects of dividends on consumption, they do allow investigation of whether the receipt of a dividend leads to a withdrawal from the account, which would be necessary if the dividend were to be used directly for consumption. The brokerage data also have the advantage that they follow individual households over long periods. This provides some information on how households respond over time to dividends and capital gains, whereas in the CEX data this behavior had to be inferred entirely from cross-sectional differences. The authors restrict themselves to households that had an open brokerage account in 1991, and they exclude margin accounts, IRA and Keogh accounts, accounts that are not individual or joint tenancy accounts, and accounts whose value falls below \$10,000. They also exclude household-month observations for which they cannot identify the mutual funds or common stock in the portfolio if the unidentifiable assets constitute at least

75 percent of the account's value, and they exclude observations with extreme values of withdrawals. Their final sample includes 92,412 household-months. Because the brokerage data do not explicitly report dividend income, the authors estimate it by matching portfolio holdings with publicly available data on the dividend distributions of specific stocks and mutual funds. Net monthly withdrawals or additions to accounts are then calculated as the difference between the end-of-period account value and the appreciated (end-of-month) value of the assets held at the beginning of the month plus dividends on those assets during the month.

A simple scatterplot of household-month net withdrawals against contemporaneous total dividends, while showing a large dispersion of net withdrawals for a given level of dividends, clearly reveals bimodal behavior: a significant fraction of households at every level of dividends withdraw either exactly that amount or exactly zero. A plot of median or average net withdrawals versus dividends for deciles of total return suggests a very high propensity to withdraw dividends during the month in which they occur; similar plots indicate at most a very small withdrawal propensity from total returns. These graphical results are confirmed by regression. A linear regression of withdrawals on dividends and total returns, with all variables expressed as a percentage of previous-period account value, yields marginal propensities to withdraw dividends and total returns of 0.35 and 0.02, respectively. Both propensities are highly significant statistically. On average, dividends are at reasonable levels, with ordinary monthly dividends averaging 0.12 percent of beginning-of-month wealth. However, the maximum dividends are extremely high, amounting to nearly 30 percent of beginning-of-month wealth in the case of mutual funds and more than 100 percent in the case of special dividends. The estimated 2 percent withdrawal of total returns may seem small, but capital gains, which are much more variable than dividends, may have a nonnegligible effect on withdrawals.

The authors show that the results for total dividends mask an important difference between the propensity to withdraw mutual fund dividends and that for the other two types, ordinary and special dividends. The propensity to withdraw mutual fund dividends, reflecting in part more frequent use of automatic reinvestment options, is only about half that of the other two, and the estimate of the propensity to withdraw ordinary dividends is approximately 0.8 for all but the highest yields, which is considerably larger than the estimates for the aggregate of dividends.

The authors run a second version of these equations that includes a twelve-month lag of dividends. The results provide some clues about the dynamics of withdrawals and the importance of what the authors call “*ex ante*” effects, that is, the possibility that dividends are endogenous. If unobserved household characteristics, such as retirement, lead some households to choose high-dividend portfolios to support their consumption, causation is reversed, with consumption and withdrawals causing dividends rather than the other way around. The fact that the data include both cross-sectional and intertemporal observations makes it difficult to distinguish the effects of unmeasured household differences from withdrawal dynamics for a given household. The authors report that a one-year lag of ordinary dividends explains 57 percent of current dividends, and together with a three-month lag it explains 81 percent of current dividends. This may reflect the relative permanence of dividends paid on individual stocks, but also unmeasured and relatively permanent household characteristics. The effects of including lags in the mutual fund equation are similar, but smaller. In both cases the coefficient estimates are imprecise. In the case of special and other dividends, which are likely to be unexpected and therefore unlikely to involve reverse causation, the lagged effect has a significant negative coefficient. This may help explain why, in the regressions for total dividends, where the type of dividend is not distinguished, lagged dividends are not significant. The authors conclude that although there is some evidence of reverse causation, it is likely to play a fairly modest role in the case of ordinary and mutual fund dividends, and even less of a role for special dividends. They believe that all of the results are consistent with an important element of causality running from dividends to withdrawals and, based on their analysis of the CEX data, to consumption.

Having concluded that the MPC out of dividends is substantial and much greater than for capital gains, the authors discuss various possible explanations for this behavior. They argue that borrowing constraints, sometimes used to rationalize a high sensitivity of consumption to current income, do not explain the large gap between the MPC from dividends and that from capital gains. While acknowledging that, in principle, transaction costs could explain that gap, they note that the propensities are essentially the same for households that are saving, who could simply save less without incurring additional transaction costs, and that the propensity is the same or higher for households with high portfolio turnover, who also could reinvest dividends with little additional cost.

The authors are skeptical of tax explanations, pointing out that high-tax-rate households withdraw much more than needed to meet the associated tax obligation and that, throughout their sample period, the tax rate was much higher on dividends than on capital gains, suggesting that, if anything, consumption needs should be met by selling stock. They acknowledge there is empirical support for the idea that aggregate consumption responds more to permanent than to transitory changes in asset values, and they recognize that changes in dividends may provide information about the permanent component of stock returns. But they observe that their results are driven largely by cross-sectional, not intertemporal, variations in returns and that the use of time fixed effects in their regressions absorbs aggregate movements in asset values, so that differences in capital gains across household-month observations are relatively permanent. Although the remaining variation in idiosyncratic household stock returns may have a transitory component, they find it difficult to explain the large difference between the consumption of dividends and that of capital gains by its presence.

The authors suggest that mental accounting theories may provide the best explanation of their results. They give as an example a model developed by Hersh Shefrin and Richard Thaler in which households place wealth into one of three mental accounts—current income, current assets, and future wealth—and follow the popular advice to “spend income, not principal.” Indeed, Shefrin and Thaler explicitly predict that the propensity to consume wealth categorized as current income, such as dividends, is greater than that from wealth categorized as assets. Not only is this prediction consistent with the authors’ main results, and with the fact that their estimate of the MPC of dividends is similar to the MPC of labor income, but mental accounting can also rationalize other features of their results, such as the fact that the MPC for special dividends falls in between that for ordinary dividends and that for capital gains. They recognize that the underlying psychology behind this sort of mental accounting is an important open question, and they suggest that self-control and prospect theory are potential psychological roots.

Many proponents of the Job Growth and Taxpayer Relief Reconciliation Act of 2003 argued that the dramatic tax cut on dividends, by lowering the cost of capital, would stimulate investment and saving. The authors’ results suggest that the cuts might indeed have stimulated the economy, but by leading to an increase in consumption rather than in investment and saving.

According to some, the cut would simply raise the price of stocks without necessarily increasing dividend payouts. But even without higher payouts, the reduction in taxes on existing dividends could have significant effects on consumption. The middle range of the estimates of the MPC from before-tax dividends implies an MPC from after-tax dividends above 0.7. Using the IRS figure for individuals' dividend income in 2002 of \$103 billion and James Poterba's estimate of the reduction in the average marginal tax rate on dividends resulting from the tax cut, the authors calculate that consumption would increase by about \$8.5 billion.

Increases in dividends, induced by the more attractive tax treatment of dividend income, would substantially add to that effect. The authors report that evidence on the importance of the cut tax for dividend payouts is mixed. Survey evidence suggests a relatively minor role, but a study led by Raj Chetty and Emmanuel Saez that compares dividend payouts of firms with different tax incentives credits the tax cuts with virtually all of the increase in dividends from 2002 to 2003. This would imply an additional \$5.8 billion in consumption, for a total effect of \$14.1 billion. Using the authors' high-end estimate of the MPC would raise that estimate to \$23.8 billion. If the cut were credited with the increase in dividends over two years, these effects would be roughly doubled. Although these numbers do not look large relative to total consumption expenditure in 2003 of \$7.7 trillion, the authors observe that they are not small relative to the \$66 billion standard deviation of consumption increases over the previous five years.

In conclusion, the authors observe that their findings that the composition of financial returns is of first-order importance in explaining consumption has implications for a range of questions in corporate finance, macroeconomics, behavioral economics, and tax policy. They stress that the interesting result is not that the propensity to consume capital gains is low, but that the propensity to consume dividends is so high. And they suggest that this difference may, at least in part, reflect mental accounting processes of the sort summed up in the adage "consume income, not principal."

HISTORICALLY, INCREASES IN short-term interest rates in periods of monetary tightening have been accompanied by increases in longer-term yields. Yet between June 2004 and February 2005, a period in which the Federal Open Market Committee raised the federal funds rate by 150 basis points, the ten-year yield on government bonds *fell* 70 basis points and the ten-year

forward rate by more than 100 basis points. Described at the time as a conundrum by then–Federal Reserve Chairman Alan Greenspan in testimony before Congress, this pattern of interest rates perplexed many economists and financial analysts. By June 2005 the ten-year forward rate had fallen a total of 170 basis points, despite further Federal Reserve tightening. Since then the forward rate has rebounded, but by only about 50 basis points. Observers have offered a range of possible reasons for this puzzling behavior, invoking changes in expectations of long-term growth of output or inflation, global increases in saving, declines in macroeconomic and financial uncertainty, better allocation of risk, and developments that change the demand for or the supply of long-term fixed-income securities. In a report that concludes this issue, David Backus and Jonathan Wright discuss the empirical basis for such explanations and the findings of bond pricing models that decompose forward rates into term premiums and expected future short-term rates. Versions of these models that distinguish between the real and the inflation components of interest rates find that declining term premiums, rather than changes in expected real rates or inflation, appear to be the major element in the surprising behavior of long-term rates. This leads the authors to explore the possible influence of various changes in the macroeconomic environment on the term premium itself.

Backus and Wright begin by presenting some basic facts about the conundrum and the behavior of variables often cited as central to the determination of the yield curve. The conundrum is real. They show that during the three preceding episodes of monetary tightening, starting in 1986, 1994, and 1999, the ten-year yield on U.S. Treasuries increased sharply, in stark contrast with the modest decline in long-term yields in the recent episode. Long-term yields are essentially an average of short-term rates and forward rates, and the authors show that the modest decline in long-term yields itself masks a sharp decline in forward rates, offset by the increase in short-term rates. Plotting monthly data on the ten-year Treasury forward rate against the unemployment rate from January 1985 to February 2007, they show that although forward rates in the recent recovery show the same countercyclical behavior as in earlier recoveries, the level is 1½ to 2 percentage points lower than past cyclical patterns would predict.

A change in inflation expectations, suggested by some as the explanation for the drop in long-term rates, should show up in the behavior of nominal, but not real, forward rates. Using the rate on Treasury inflation-protected securities (TIPS) as a measure of real rates, the authors show that the shifts



between June 2004 and June 2005 in the nominal and real forward rate curves were similar, although the shift in the nominal curve was greater. At a ten-year horizon, the nominal curve dropped by 172 basis points, whereas the real curve dropped by 96 basis points. In interpreting this difference, the authors note that the TIPS rate includes an inflation risk premium and a TIPS liquidity premium as well as an indication of expected inflation.

Are these shifts consistent with forecasts of inflation, GDP growth, and short-term interest rates in 2004–05? The authors report that the professional forecasts in the Blue Chip surveys for this period show little movement in expected inflation or GDP growth at a five- to ten-year horizon. The long-range forecast of the three-month Treasury bill rate was also flat. They conclude that the recent decline in long-term yields does not come from expected declines in any of these three variables.

Indicators of shifts in investors' perception and pricing of risk are more promising candidates. The authors report that, between 2004 and early 2005, there were sharp declines in the spread between Baa- and Aaa-rated corporate debt and in the implied volatility of one-year interest rate caps, a measure of short-term volatility in the six-month London interbank offer rate. Realized volatility, measured by the standard deviation of daily changes in forward rates and an index tied to the short-term volatility of S&P stock index options, also declined during the period. The authors find this persuasive evidence that financial market risk and risk premiums across a range of assets were substantially lower in 2005 than they had been two or three years earlier.

The source of the decline in asset market volatility is less clear. The authors regard macroeconomic uncertainty, particularly at long horizons, as inherently difficult to measure. Although recent Federal Reserve communications may have given markets more forward-looking guidance, which may have made the path of monetary policy more predictable, the broader impact on asset market volatility is less evident. Taking the view that the dispersion of long-horizon predictions may reflect intrinsic uncertainty, the authors show that the dispersions of predictions of consumer price inflation, real GDP growth, and the three-month Treasury bill rate have all trended noticeably lower. Most of the decline in the dispersion of interest rate and inflation expectations occurred in the early 1990s; both then rose in the last recession before declining again over the conundrum period. The authors note that movements in forward rates for the United States, Ger-

many, and the United Kingdom have been highly correlated in the last few years, and rates have declined in a number of middle-income countries as well. So explanations for the decline in U.S. forward rates should not be confined to U.S. developments and should account for the decline that has occurred more broadly.

Backus and Wright turn from this evidence—that lower risk or term premiums, rather than declines in expected inflation or future short-term real rates, are responsible for the decline in long-term yields and forward rates—to the results of more formal affine bond pricing models, all of which decompose long-term forward rates into expected future short-term rates and a time-varying term premium. Such models all generate yields and forward rates that are linear functions of a vector of state variables, which themselves follow a linear autoregressive process. The models differ in their choice of explanatory variables. The simplest use a small number of “latent” factors characterizing the level and shape of the current yield curve as state variables, but more recent versions add direct information about the dynamics of interest rates or other economic variables.

The authors present results for three such augmented versions. The first uses professional forecasts of Treasury bill rates to identify the expected short-term rate component of the forward rate. This decomposition suggests that nearly the entire decline in ten-year-ahead forward rates from June 2004 to June 2005 comes from a decline in the term premium. Since then the term premium has been roughly constant, with a small rise in the forward rate attributed to a comparable rise in expected future short-term rates. A second version adds the dynamics of inflation to the model, making it possible to price synthetic real bonds and decompose a nominal forward rate into an expected future real short-term rate, expected future inflation, a real term premium, and an inflation risk premium. The third version distinguishes the same four components of forward rates but makes direct use of the information in TIPS yields, which are assumed to equal the sum of yields on a synthetic liquid real bond and a TIPS liquidity premium. This model’s estimates of the ten-year forward rate suggest that the real term premium has fallen sharply since June 2004, and the inflation risk premium by a lesser amount. On the other hand, expected future real short-term interest rates have been flat, and expected future inflation has actually risen modestly.

A shortcoming of affine term structure models is the lack of an economic explanation of movements in the latent factors that serve as state variables.

The authors explore the possibility of the role of cyclical movements and inflation uncertainty by regressing the nominal term premiums estimated from the first model just described on unemployment and the dispersion of long-term inflation forecasts. They find that the coefficients in this regression are statistically significant, and the variables explain a substantial fraction of the decline in term premiums since 1990, their rise in 2000–02, and a modest portion of their decline since June 2004.

The authors view the combination of facts about risk spreads and volatilities, survey expectations, and affine model decompositions as persuasive evidence that a decline in term premiums, rather than a decline in future expected rates, is likely to be the principal explanation for the decline in long-horizon forward rates from June 2004 to June 2005. They recognize that the models do not provide conclusive answers. However, they stress that, within this framework, the arithmetic identity that the forward rate is the sum of these two components carries a strong implication. If term premiums did not decline, then the long-run expectation of the federal funds rate must have fallen by about 1.7 percentage points. They note that economic forces—for example, an increase in global saving or a scaling back of expectations about long-run productivity growth amid weak business investment—could have worked to lower the expected real short-term rate. But they doubt that such factors caused the large decline actually observed, given the robust growth during the period and the upward drift in inflation. The same arithmetic implies that if the term premium today equals its 2-percentage-point average over the last twenty years, then the current forward rate of about 5 percent implies an expected federal funds rate ten years hence of close to 3 percent.

What factors could explain the decline in uncertainty about future inflation, interest rates, and growth that Backus and Wright have identified as a partial explanation of the decline in term premiums? Those they identify include more credible and transparent monetary policy, a global trend toward central bank independence, and the greater integration of financial markets, which reduces the potential for short-term gains from adopting a more inflationary policy. As evidence supporting this view, they cite the especially sharp decline in forward rates in the United Kingdom around 1997, when the Bank of England was granted operational independence. But as contrary evidence they note that one would have to explain why it was only in the last few years, long after the disinflation engineered by the Federal Reserve under Chairman Paul Volcker in the early 1980s, that U.S.

investors have suddenly become confident that inflation will remain contained indefinitely. Still, the authors place reduced inflation uncertainty resulting from changes in central bank policy high on the list of plausible explanations.

Backus and Wright cite three other proposed explanations for the decline in term premiums in particular countries, but they note that these have more difficulty explaining the pattern across countries. The decline in real long-term rates in the United Kingdom could reflect the prospect that U.K. corporate pension reform will encourage pension funds to better match the durations of their assets and liabilities, causing a substantial increase in the demand for long-duration securities. But this explanation does not fit the United States, where pension reform has proceeded more slowly. The large purchase of U.S. Treasury securities by Asian central banks might have contributed to the fall in the U.S. forward premium but does not explain the parallel falls in Germany and the United Kingdom, unless the bonds of all three countries are close substitutes. Last, the authors discuss the suggestion that changing demographics may have contributed, as large cohorts nearing retirement shift their savings from risky equities to bonds and other relatively safe assets. Setting aside the question of whether the equity premium has increased coincidentally with a decline in the term premium, the authors note that demographics are slow moving and predictable, with no substantial or unexpected shift having occurred in or around June 2004.

The authors conclude that the evidence points to a declining term premium as the primary source of the recent fall in long-term rates, and that such a decline is broadly consistent with observed changes in risk spreads and other measures of uncertainty. But they recognize that the economics profession is far from having a complete understanding of the behavior of bond prices. In their view the next step should be the development of models in which macroeconomic policy and behavior can be tied more directly to the behavior of interest rates. They believe that recent lines of research give hope that eventually economists will better understand the connections between macroeconomic developments and financial market outcomes.