

Perspectives on Research Issues In Consumer Behavior

Based on the speech President Santomero delivered to the Conference on Consumer Transactions and Credit, March 22-23, 2001

BY ANTHONY M. SANTOMERO

In this issue's *Third Dimension*, President Santomero discusses the subject of consumer financial assets and its omission from the research agenda of the academic community. He also presents a list of issues that need to be studied to further our understanding of this rapidly growing part of the financial landscape.

Generally, economic forums have tended to ignore the broader issues in consumer credit, preferring to focus on the valuation of more standard corporate financial claims. When the consumer is discussed, it is typically consumer consumption and savings decisions that are studied and analyzed. This neglect of consumer credit seems remarkable, given that debt owed by households represents over 25 percent of total credit market debt outstanding; that consumer credit, excluding mortgages, makes up over 10 percent of commercial bank credit; and that the outstanding volume of consumer credit, including mortgages, exceeds the volume of U.S. government debt.

There is no single reason for the omission of consumer financial assets from the research agenda of the academic community, but I can offer a few rationalizations. The first is that macroeconomists – even those who specialize in monetary theory – have had little interest in the detailed behavior of asset choice at the consumer level. They prefer to concentrate on consumption theory and the associated

empirical tests of these theories, rather than analyze the allocation of savings and wealth dynamics. Second, specialists in finance have tended to concentrate on firm-level behavior because firms are viewed as more rational players than consumers and firms' behavior is of more economic value. Asset sizes are bigger; representative agents are more easily modeled; and market discipline seems to force the decision maker closer to the optimal economic choice.

Only recently has this begun to change. With the emergence of the asset-based security markets, financial theorists and empiricists have begun to examine the behavior of financial assets that have resulted from the aggregation of consumer debt. This is most obvious in the mortgage market, where the emergence of various types of securitized mortgage instruments fostered research on their valuation and time-series dynamics. This interest has recently been expanded into other types of asset-based securities, such as CARDs, CARs, and CLOs. In each case, to analyze the underlying asset, the researcher has to examine optimal

decisions of a representative agent and the impact of aggregating individual claims on instrument behavior.

Another development that has helped spur interest in the micro-economics of consumer financial decisions is the intellectual shift that has taken place in macroeconomics toward a well-specified microeconomic foundation for macroeconomic theories. It is now acknowledged that to understand consumption and savings decisions on the macro level, we must model the behavior of individual agents. And this interest in microfoundations has been accompanied by the development of new data sets, such as the Federal Reserve's triennial Survey of Consumer Finances, the Consumer Expenditure Survey, and the Panel Study of Income Dynamics, that present information at a disaggregated level to allow for testing of these micro theories.

Yet, despite recent interest in consumer debt instruments, there is



Anthony M. Santomero, President,
Federal Reserve Bank of Philadelphia

much work to be done. This conference is just the first of many efforts that we at the Federal Reserve Bank of Philadelphia plan to make to advance the consumer credit research agenda. We hope to shed light on the state of research and to spotlight areas of potential future contributions.

In my comments I will try to put the current literature on consumer finance into context and explain why consumer credit should have a place in academic research between standard macro modeling and the valuation of standard financial assets. I will also try to set out a list of issues that must be studied to further our knowledge and understanding of one part of the financial landscape that continues to grow at double-digit rates.

CONSUMPTION THEORY AND THE RESEARCH QUESTIONS IN MACROECONOMICS

I'll begin with the standard view of consumer choice presented in macroeconomic theory. From the traditions of classical macro theory to the work of Keynes, standard macroeconomic models involved the representative economic agent who maximizes an expected utility function that summarizes well-behaved preferences, in a world with positive marginal productivity and a fixed discount rate. The agent's choice is constrained by a multiperiod budget or endowment sequence. Nearly all of us as graduate students could write down this problem's multiperiod first-order conditions and the appropriate transversality conditions associated with optimal behavior. The representative agent maximizes expected utility by equating marginal utility across periods, subject only to a borrowing and lending condition that bounds the problem. Such models yield paths of consumption and savings over time that achieve maximum expected utility subject only to the aggregate interest rate, which itself can be solved

for simultaneously in a general equilibrium setting.

Once consumption behavior had been characterized, economists turned to characterizing savings behavior in a world where the agent's life span is uncertain. Here, the profession's attention centered on models that featured various types of risk aversion and their impact on both wealth accumulation and decumulation. This work proved central to understanding the role of pensions on intertemporal consumption behavior and remains an important part of our understanding of life-cycle savings behavior. The credit market imperfections leading to liquidity constraints and the significance of the bequest motive in driving savings decisions have been studied.

On the public policy stage the research offers important insights into the debate over both Social Security and the entire area of private pension programs. Others have also used it to

relegated to macrofinance theorists, who began where the monetary theorists left off. In these macrofinance models the representative consumer maximizes his expected utility, which is represented by a concave function of exogenously given wealth. The models are usually single period. Multiperiod considerations are sometimes addressed with some intertemporal endogeneity, leading to some of the rich models of asset choice in a multiperiod consumption framework.

Still, we are a long way from a sophisticated model of households' ultimately joint decisions about how much to save, how to divide savings among different types of assets, how much to borrow, and what types of debt to incur. The characteristics of the specific assets are typically not modeled beyond the first two moments of their underlying distributions. Uncertainty is characterized only in this most general way, usually by reference to normality and stationarity. Anything more specific

Today an increasing share of household wealth is held in financial assets, and a smaller share is invested in the household's primary residence.

understand the effects of the private sector's evolution to defined contribution plans and their associated actuarial risks to covered workers. Some have even discussed and modeled the role of inflation uncertainty in this multiperiod asset-choice problem.

However, few researchers studying this problem attempted to explain the specific portfolio of risky assets that should be selected by economic agents over their life cycle or the appropriate vehicles to use to accumulate debt during periods of excess consumption. These issues were

is relegated to applied finance and its interest in state-dependent payoffs and their effect on valuation and time-series dynamics.

In short, the detailed payoff patterns associated with the debt instruments issued by the agent were entirely too mundane for the theorists' consideration. So, a student of consumer credit and consumer choice was sent to the applied finance literature for the modeling and valuation of real consumer debt instruments.

To some extent, the underdevelopment of this line of research

probably reflects the relative simplicity of most U.S. households' portfolios before the 1990s. But an explosive period of financial innovation in the last two decades and the rapid growth in consumers' wealth in the 1990s have introduced many more households to many more financial options. For example, today an increasing share of household wealth is held in financial assets, and a smaller share is invested in the household's primary residence. More households participate in equity markets, with equity making up over half of households' financial holdings.

These changes in the structure of household portfolios can be attributed to the stock market boom of the 1990s, the relatively slower growth in housing prices during the same period, a highly competitive mutual fund industry, and demographic factors. While the shift in household portfolios in the 1990s has been dramatic, what might be more surprising is that the shift hasn't been even larger. Households still seemingly prefer relatively simple and safe portfolios.

Exactly why this is the case, however, is unclear, for as noted above little economic research has centered on consumer asset choice. Still less has investigated the characteristics unique to each financial product available to smooth consumption over a multiperiod uncertain horizon. To learn more about these issues the researcher is directed to the field of finance.

FINANCE AND THE VALUE OF CONSUMER FINANCIAL CLAIMS

Interestingly, applied finance is up to the task of modeling consumer assets. Its general theoretical construct of stochastic behavior, using expected present value in a state-contingent framework, is well suited for the underlying uncertainty associated with consumer credit. With the rise of discrete-time models and binomial tree modeling techniques, the standard

financial valuation techniques can be applied to consumer debt instruments characterized by time-specific default probabilities and multiperiod cash-flow uncertainty.

But, as indicated earlier, consumer assets do tend to have small market values, and the field was traditionally viewed as extraordinarily "unsexy." As a result, the new techniques being applied to analyze corporate finance instruments have only rarely been applied to consumer debt, and there has been little interest expressed by the research community in understanding the nuances of the individual consumer markets. As a result, the characteristics and time-series price dynamics associated with different types of consumer debt instruments have never been seriously investigated.

This void in the literature is important. Consumer debt of various types exceeds total U.S. Treasury debt, and it is likely to be the only liquid market of the early 21st century. The economics and finance profession has been slow in concentrating on both the theoretical and empirical issues surrounding these instruments – with one exception, the mortgage market.

Residential mortgages currently account for the majority of consumer debt, with an outstanding value of over \$5 trillion, and represent the only area of consumer debt that has received significant attention. Fostered by the development of various types of mortgage pass-through instruments, a large literature developed about their valuation. The aggregation of individual instruments added scale but also complexity. Prepayment risk, default modeling, and state-dependent cash flows led to the development of a rich literature.

Interestingly this literature has not crossed over to other areas of consumer debt, even as the percentage of consumer debt represented by mortgage assets continues to decline. In

fact, there is very little in the literature on various types of asset-backed securities, in spite of their growing importance, and we know little about the difference in yields across various types of asset-backed securities.

ADDRESSING THE VOID IN THE LITERATURE

How can this seemingly important void in the literature be addressed? The research plan is fairly straightforward if one proceeds from first principles. The securitization vehicle can be understood as the aggregation of individual economic decisions made by a series of representative agents. Using consumption theory in a stochastic environment, one could model cash-flow dynamics, including default and prepayment variation. Present value valuation techniques could then be applied to the underlying aggregation, subject to the subtle but important covenant constraints that differentiate each instrument.

Collateral, recourse, and seniority become key elements in the distinction between repayment timing and ultimate default probabilities associated with different consumer debt types. Issues of adverse selection, switching costs, borrower heterogeneity, and liquidity constraints also come into play, making the pricing of consumer debt instruments, such as credit cards, not a trivial exercise. However, with sufficient care, the techniques applied to mortgage debt, which were derived from multiperiod consumption behavior, could be applied to CARs, CARDs, and CLOs. Price dynamics would follow directly.

Any such effort would immerse the research in the public policy debates that dictate terms and conditions. Just as "due on sale" clauses substantially altered mortgage dynamics in the 1970s, and regulations surrounding adjustable rate mortgages altered the valuation of these instruments substan-

tially, consumer protection policies have substantial implications for consumer credit instruments. For example, bankruptcy regulation alters repayment probability and valuation, as do access to collateral, privacy rules, and public disclosure regulation. All of these policies alter credit risks and credit spreads.

The public debate surrounding these consumer protection attributes would be enhanced by both the estimation and discussion of the implied costs of such regulations. The beneficial effect of the consumer rights acquired through such legislation must be weighed against the implied cost to the lender and the derived increase in credit spreads to the entire consumer class.

THE UNFINISHED AGENDA

There is much to be gained from the application of standard economics and finance to the valuation of consumer debt instruments. By modeling the uncertain cash flows and creditor rights associated with different debt instruments, we can obtain proper pricing and risk assessment. This will also facilitate arbitrage across markets and reaffirm the law of one price across various types of instruments subject to similar uncertainty.

Such research will also illuminate the current policy debate regarding debtor versus creditor rights at the basis of various types of consumer right proposals. Valuation techniques can derive the costs of such legislation against which the benefits can be contrasted.

THE FEDERAL RESERVE BANK OF PHILADELPHIA'S ROLE

This is a rich agenda that all of us share as researchers. And the Fed has

a unique role to play, as both a center of financial knowledge and a supporter of various types of economic research. The Fed has traditionally devoted significant resources to banking and financial market research, at a number of the Reserve Banks.

We here in Philadelphia have a particular interest in this area. The Third District is the home of the credit card industry, with approximately 40 percent of all consumer credit cards emanating from the state of Delaware.

To underscore our interest in this area and to address what we see as a mandate to understand this industry, we have established a Payment Cards Center within the Bank to investigate issues central to this part of the financial services sector. The value of various types of consumer debt instruments and their relative values under different


By modeling the uncertain cash flows and creditor rights associated with different debt instruments, we can obtain proper pricing and risk assessment.

public policies, including but not limited to bankruptcy rights, collateral access, and various types of consumer protection legislation, are clearly crucial issues, and they definitely belong on the Payment Cards Center's research agenda.

Payments issues, as well, are on the Center's agenda. Clearly these too are central to understanding the payment cards industry and the evolving financial services sector. Just as one can tie the study of consumer debt instruments to a traditional literature,

the study of consumer transactions media finds a historical precedent in the traditional theory of money demand. The traditional theories of money demand — for example, the Baumol-Tobin inventory model — weren't specific about the medium used for transactions. And, as we know, financial innovations and deregulation effectively destroyed the empirical relationship between money and income.

We need to develop new theories if we hope to explain the economic rationale for and the impact of various transactions media, like credit cards, debit cards, and smart cards, which are much more complicated than our traditional characterization of money. Debit cards and smart cards are relatively new developments, but there is much we still do not know about credit cards, which have been around

for years: why do so many consumers, even wealthy ones, borrow at the high rates on credit cards? Given the large number of credit card issuers, why are credit card rates so high and sticky? Why was there a large increase in delinquencies and bankruptcies when economic conditions were so favorable? A well-developed theoretical framework for solving these and related puzzles is a pressing task, which promises an interesting and rich research agenda. 

A Summary of the Conference On Consumer Transactions and Credit

In March of this year, the Federal Reserve Bank of Philadelphia and the Wharton School of the University of Pennsylvania, in association with the *Journal of Financial Intermediation*, sponsored a conference on research issues involving consumer transactions and credit. This article offers a summary of the papers presented at the conference.

Consumers today have more financial options for saving, borrowing, spending, and investing than ever before. Yet little is known about consumers' decisions about how much to save, which types of assets to invest in, how much to borrow, which types of debt to incur, and which instruments to use to make payments. Similarly, little is known about how firms price the financial instruments consumers choose to use. This intellectual neglect of consumer finance, and especially consumer credit, is somewhat surprising, given that debt owed by households represents over 25 percent of total credit-market debt outstanding and that the outstanding volume of consumer credit, including mortgages, exceeds the volume of U.S. government debt.

To begin to address these issues and to encourage more research in the field of consumer finance, the Federal Reserve Bank of Philadelphia and the Wharton Financial Institutions Center, in association with the *Journal of Financial Intermediation*, sponsored a

one-day conference at the Philadelphia Reserve Bank on March 23. Five research papers and two addresses were presented to an international group of economists, who discussed consumer credit and transaction behavior. *

In his opening address to conference participants, Federal Reserve Bank of Philadelphia President Anthony M. Santomero suggested that the lack of research on consumer financial behavior in part "reflects the relative simplicity of most U.S. household portfolios before the 1990s. But an explosive period of financial innovation in the last two decades and the rapid growth in consumers' wealth in the 1990s have introduced many more households to many more financial options."

President Santomero said more research is needed. "This conference is just the first of many efforts that we at

* The papers are available on our web site at www.phil.frb.org/econ/conf/program.html.

the Federal Reserve Bank of Philadelphia plan to make to advance the consumer credit research agenda," he said. "We hope to shed light on the state of research and to spotlight areas of potential future contributions."

As he noted, the Philadelphia Fed has a particular interest in the area, since some 40 percent of consumer credit card activity emanates from Delaware banks in the Third Federal Reserve District. To underscore this interest, the Bank has established a Payment Cards Center that will serve as a focal point for investigating issues central to this dynamic sector of the financial services industry. The conference, proceedings of which are summarized below, represents an important first step in that direction.

CONSUMER RESPONSE TO CHANGES IN CREDIT SUPPLY

The credit card market remains a relatively understudied area of consumer finance. Nicholas S. Souleles, of the Wharton School, presented some interesting findings on the behavior of credit card borrowers, concluding that liquidity matters. His paper, "Consumer Response to Changes in Credit Supply: Evidence from Credit Card Data," co-authored by David B. Gross, formerly of the Graduate School of Business, University of Chicago, is based on analysis of a unique data set of several thousand individual credit card accounts followed monthly for 24 to 36 months. Souleles reported that according to their empirical work, "increases in credit limits for credit card borrowers generate immediate and significant increases in debt," especially for people who are already close to their limit. This

suggests that these credit limits are a binding liquidity constraint. However, even people who have not borrowed enough on their credit cards to be near their limits start borrowing more when credit card issuers increase their credit limits.

This finding is consistent with theories about precautionary savings. Consumers worry not only about credit constraints that are already binding but about the possibility that they will face a binding constraint in the future and will not be able to borrow and consume as much as they would like. Consumers want to keep some cash on hand, including some of their available credit, to act as a buffer against unexpected emergencies. But when their credit limits rise, only part of the increase is reserved for the buffer; it is optimal for them to consume some of the increase.

Souleles indicated that some of their other results, however, are not consistent with current theories of consumer savings behavior. For example, many of the people who are borrowing on their credit cards hold relatively large balances in their low-interest checking and savings accounts. Gross and Souleles found that “one-third of borrowers have over one month’s worth of income in liquid assets, which is more than typically needed for cash transactions.” These funds could be used to pay off high-interest credit card debt without sacrificing much interest income.

Another interesting result is that in contrast to most other studies, their research also indicates that consumers’ credit card debt is particularly sensitive to changes in credit card interest rates, especially to large declines in rates. This might explain the widespread use of teaser rates. The discussant, Paul S. Calem, of the Federal Reserve Board, suggested that further analysis of the effect of teaser rates on credit card borrowing would be an interesting avenue for future

research. Souleles said their analysis also showed that consumers seem to respond to a reduction in a credit card rate by switching balances from other cards to the low-rate card. So consumers appear to be sensitive enough to interest rate changes to overcome the costs associ-

tude and significance of switching costs. In their paper, the authors develop an empirical model that is able to quantify the importance of switching costs as well as customers’ probabilities of switching from one firm to another, even when customer-specific data are absent.

In contrast to most other studies, [Gross and Souleles’] research also indicates that consumers’ credit card debt is particularly sensitive to changes in credit card interest rates, especially to large declines in rates.

ated with switching balances from one card to another. Such switching costs were the subject of another paper presented at the conference.

ESTIMATING SWITCHING COSTS

One of the factors often pointed to in explaining why credit card interest rates are relatively sticky and generally only somewhat responsive to changes in market interest rates is that there appear to be large costs associated with switching accounts. But these switching costs have been difficult to quantify. In “Estimating Switching Costs and Oligopolistic Behavior,” Moshe Kim and Doron Kliger, both of the University of Haifa, and Bent Vale of the Central Bank of Norway, present a method for estimating switching costs for bank customers. Their research is based on aggregate data, since customer-specific information is typically not available.

In his presentation, Kim discussed the theoretical literature on switching costs, which indicates how such costs can influence firms’ pricing behavior. But because micro-level data on individual transactions are nearly impossible to come by, researchers have had difficulty in estimating the magni-

The model, which was estimated using aggregate data on banks that operated in Norway from 1988 to 1996, focused on the market for bank loans. Bank lending is a good candidate for study, since long-term relationships and repeated contacts among banks and their customers — factors that characterize bank lending — may be a source of switching costs.

The study’s empirical results confirm the importance of switching costs in bank lending, with the estimated magnitude of switching costs differing across various subsamples of the banks. For the entire sample, switching costs average 4.1 percent, which is about one-third of the market average interest rate on loans. But switching costs are found to decrease with bank size, down to 2.1 percent for banks with 60 or more branches. This decrease in the size of switching costs may occur because the customers of large banks tend to be large companies. These firms are often publicly traded and enjoy greater market mobility than small retail customers. Consistent with this result is the finding that in the sample, customers’ relationships with their banks ranged between 16.7 years at small banks down to 11.3 years at large banks. Kim and his

co-authors also find that the customer lock-in generated by the switching costs is valuable to banks: locked-in customers add nearly 25 percent to banks' value.

The discussant, Steven A. Sharpe, of the Federal Reserve Board, suggested that the authors try to measure the loan price-cost margin more precisely, since it is a crucial variable in their model. In particular, Sharpe was concerned that the measure used could reflect differences in loan risk, which could differ by bank size, as well as differences in market power, and that the imprecise measurement of the price-cost margin could be influencing the results.

BANK CONSOLIDATION AND CONSUMER LOAN INTEREST RATES

Another paper at the conference also examined banks' loan pricing behavior. In "Bank Consolidation and Consumer Loan Rates," Charles Kahn and George Pennacchi, of the University of Illinois, and Ben Sopranzetti of Rutgers University, examine whether banks change their pricing of consumer loans after bank mergers and whether the pricing behavior differs for different types of consumer loans.

The recent wave of mergers in the banking industry has spurred a number of researchers to examine the impact of mergers on potentially vulnerable bank customers. But as Sopranzetti explained, to date, most studies have focused on the supply and pricing of small-business loans and consumer deposits. In their paper, the authors shed new light on the effect of bank consolidation on the pricing of auto loans and unsecured personal loans.

In particular, while rates on personal loans tend to rise at banks in the market following a bank merger, rates on automobile loans tend to fall. The authors attribute this difference to

the scale economies that exist in the auto loan market and the fact that there is strong competition from nonbank lenders for auto loans; hence, the merger does not represent an increase in banks' market power in the auto loan market. Thus, consumers in the market for new auto loans are likely to benefit from a merger, since prices fall, while those seeking unsecured personal loans are not likely to see better pricing options.

Sopranzetti reported on other findings of their analysis: (1) Leader-follower pricing behavior is more widespread in automobile loan markets than in personal loan markets. In other words, if one bank changes its auto loan rate, other banks are likely to follow it by changing their rates. This is consistent with the greater competition found in auto loan markets. The authors suggest that higher consumer switching costs in the personal loan market might make the personal loan rate set by a given bank less responsive to its competitors' rates. (2) Personal loan rates are stickier than automobile loan rates, in the same

While rates on personal loans tend to rise at banks in the market following a bank merger, rates on automobile loans tend to fall.

way that consumer deposit rates tend to be sticky, that is, not very responsive to changes in the overall level of market interest rates. And, consistent with empirical research on consumer deposits, personal loan rates are more rigid in more concentrated markets. (3) Both automobile and personal loan rates tend to respond asymmetrically to increases and decreases in market rates. Banks

are slower to lower consumer loan rates when warranted by declines in other market rates than they are to raise consumer loan rates when other market rates rise.

The discussant, Leonard I. Nakamura, of the Federal Reserve Bank of Philadelphia, indicated that personal loans are much more heterogeneous across banks than are auto loans and that this might be influencing the results. He also suggested that having a theory about how banks' pricing behavior would change after a merger would be helpful in interpreting the results. Conference participants suggested that a more complete consideration of market competition would be a useful extension. For example, how do the special characteristics of auto finance subsidiaries affect the bank market for auto loans? How does unsecured credit card debt come into play in understanding the market for personal bank loans?

PERSONAL BANKRUPTCY AND THE LEVEL OF ENTREPRENEURIAL ACTIVITY

Another area relevant to the study of consumer finance is the issue of bankruptcy. In "Personal Bankruptcy and the Level of Entrepreneurial Activity," Wei Fan of the University of Michigan and Michelle J. White of the University of California San Diego examine the effect of the provisions in bankruptcy law on entrepreneurial activity. Their paper is fitting for a conference on consumer finance, since, in many cases, it is very difficult to disentangle a small business's finances from that of the owner's.

Small businesses, as well as consumers, can file for bankruptcy under Chapter 7 of the federal bankruptcy code. Debts of noncorporate firms are considered personal liabilities of the entrepreneur/owner in the event of a business failure. The law requires that the entrepreneur give up assets above a

fixed bankruptcy exemption level for repayment to creditors, but all their future earnings are exempt from the obligation to repay. This “fresh start” provision lowers the ultimate risk of starting a business, since more of the entrepreneur’s assets will be protected in the event the business fails. So the exemption provides a form of “wealth insurance” to the business owner, and higher exemption levels could potentially encourage the formation of more new businesses. The fact that the exemption levels are set by states and vary widely, especially the exemption for the debtor’s house (the homestead exemption), provides a natural laboratory for studying whether bankruptcy exemptions have a significant economic effect on entrepreneurship.

The authors find empirical support for the idea that the bankruptcy system is a factor in a worker’s decision to be self-employed rather than to work for others. Fan and White find: (1) The probability that families that own homes are self-employed is 35 percent higher if families live in states with unlimited homestead exemptions rather than low exemptions. (2) Families that are homeowners are 22 percent more likely to start businesses if they live in states with higher or unlimited, rather than low, homestead exemptions. And they are more likely to organize their businesses as noncorporate rather than corporate. (3) One possible negative effect of higher exemptions is that they may encourage more bankruptcy filings, but the authors do not find that entrepreneurs are more likely to terminate businesses if they live in states with unlimited rather than low homestead exemptions.

In her presentation, White commented on the proposed new federal bankruptcy legislation and possible implications based on this research. The proposed legislation focuses on reducing abuses by relatively well-off individual debtors. However,

Fan and White’s research suggests an unintended consequence of adopting these reforms could be a reduction in the level of self-employment by U.S. households.

The discussant, Mitchell Berlin, of the Federal Reserve Bank of Philadelphia, generally applauded the paper for extending what is generally considered a consumer issue to the environment of small business and entrepreneurship. At the same time, he also made the point that while high exemptions may support small-business formation, they may just as likely reduce the supply of credit to borrowers, which would mitigate any positive effect on small-business formation. In addition, he pointed out that it was still an open question whether self-employment necessarily means more economic growth or whether it reflects employment redistributed from larger firms.

THE COSTS AND BENEFITS OF TRANSACTIONS PRIVACY

Another issue that’s become increasingly important to consumers is whether the privacy of their transactions is being protected. Rapid advances in information technology have dramatically lowered the cost and increased the speed of record keeping and transmission of information. The Internet has not only affected the cost of transmitting information but also broadened the nature of potentially available information, including information stored on personal computers. All of these factors have led to undeniable increases in convenience and welfare to consumers, but they have also fueled the public debate on privacy, particularly Internet privacy.

In their paper, “A Theory of Transactions Privacy,” Charles M. Kahn, of the University of Illinois, James McAndrews, of the Federal Reserve Bank of New York, and William Roberds, of the Federal Reserve Bank of Atlanta, develop a model to examine

the tradeoffs between the costs and benefits of transactions privacy. In their model, privacy means the concealment of potentially useful information, but concealment also potentially bestows benefits. As long as contracting is flexible and the initial rights to the

Rapid advances in information technology have dramatically lowered the cost and increased the speed of record keeping and transmission of information.

information are clearly assigned, Coase’s theorem would suggest that privacy laws would not be necessary. Once property rights are initially assigned to a party (either one), the parties will bargain, making appropriate side payments to one another, so that the outcome chosen regarding whether to reveal or conceal information will be the one that has the largest total benefit to both parties. (The initial assignment of rights will affect the distribution of those benefits.)

However, the authors argue that there are good reasons to believe that the assumptions of Coase’s theorem wouldn’t apply in our current transactions environment. For example, it is difficult to commit to not using information once it becomes known, and currently, neither the law nor technology clearly assigns rights to transactions information. The authors show that in the current environment, the initial assignment of rights over private information could have economic consequences.

As McAndrews pointed out in his presentation, murky rights to trans-

action information provide incentives for parties to develop technologies to control information. This could result in a race to create technologies to conceal information (for example, anonymous electronic money) or to reveal information (for example, Internet “cookies”). Clearly defining rights to transaction information would forestall wasteful investments in technology to control information.


William L. Lang, of the Office of the Comptroller of the Currency, complimented the authors on developing a formal model with which to


address the privacy issue. He did question whether the model was rich enough for policy analysis. For example, in the model, parties are assumed to be aware of when transaction information is disclosed and what information is disclosed. However, in many cases, people do not know that information has been disclosed. While the Gramm-Leach-Bliley Act requires financial institutions to disclose information sharing, nonfinancial firms are not required to do so. Lang also thought that the paper overstated the rights of consumers under Gramm-Leach-Bliley.

PAYMENT SYSTEM COMPETITION

Jean-Charles Rochet, of Toulouse University, France, presented the conference’s keynote luncheon address, “Payment Systems Competition.” Rochet’s talk was based on his ongoing research on payment systems, which he is carrying out with his colleague Jean Tirole. His remarks focused on their research-in-progress on competition between different types of payment systems: those structured as open associations of banks, like Visa and MasterCard, and those structured as closed systems, like American Express and Discover. An important and still unanswered question is whether such competition will lead to more efficient usage of payment cards. The Rochet-Tirole model is an important contribution to the study of payment systems, and it is applicable to other environments, such as competition between interbank large-value payment systems like CHIPS and Fedwire, or competition between credit cards and debit cards.

SUMMARY

The research on payment systems by Rochet and Tirole, along with the other work presented and discussed at the conference, represents an important step in meeting the challenge posed by President Santomero in his opening address to the conference. In discussing the current state of the literature, President Santomero emphasized that “we need to develop new theories if we hope to explain the economic rationale for and the impact of various transactions media, like credit cards, debit cards, and smart cards, which are much more complicated than our traditional characterization of money.” We hope that this conference and the work of the Federal Reserve Bank of Philadelphia’s Payment Cards Center will inspire other researchers to join in this effort. 



Changes in the Use of Electronic Means of Payment

BY LORETTA J. MESTER

This article presents an update of an article first published in the March/April 2000 *Business Review*, which presented data from the 1995 Federal Reserve Survey of Consumer Finances. Loretta Mester, author of the original article, has compiled information from the 1998 survey to keep our readers up-to-date.

In “The Changing Nature of the Payments System: Should New Players Mean New Rules?” (*Business Review*, March/April 2000), I presented some data from the 1995 Federal Reserve Survey of Consumer Finances on the use of electronic banking. This survey of more than 4000 households, which is designed to be representative of all households in the U.S., is redone

every three years. Attached are updates of Tables 2 and 3 from the *Business Review* indicating how the use of various means of electronic payment have changed over the three years.

As seen in Exhibit 1, use of electronic forms of payment, including ATMs, debit cards, direct deposits, automatic bill paying, and smart cards, has risen from about 77 percent of households to over 85 percent of households. A major increase was seen in debit card use, which nearly doubled, to about a third of households. Increases were seen in all categories by age, income, and education. Use of direct deposit and automatic bill paying also showed sizable increases. More than two-thirds of households have an ATM

card, but smart card use remains very low.

As seen in Exhibit 2, households that do business with at least one financial institution have shifted from paper-based methods of conducting this business to automated methods.* A sizable fraction of households, over 75 percent, still report that one of the main ways they deal with at least one of their financial institutions is in person. However, this fraction declined between 1995 and 1998, even in the population over 60 years old. Use of electronic means of doing business, including ATM, phone, direct deposit and payment, and computer, increased

*Table 3 in the *Business Review* indicated the percentage of households that reported each method as the main way of doing business with at least one of their financial institutions. In Exhibit 2 here I am reporting the percentage of households that reported each method as among one of the main ways of doing business with at least one of their financial institutions. The reason for the change is that there was a slight change in the question between 1995 and 1998, which made a comparison of the main way of doing business much less informative. In particular, the 1995 survey did not mention “check” as a possible answer, but recorded it if the respondent said it. In contrast, the 1998 survey did list “check” as a possible answer. Hence, there was much greater reporting of “check” as the main way of doing business in 1998 than in 1995. This means that a comparison between 1995 and 1998 of the percentage of households reporting a method as the main way of doing business is not informative; a comparison of the percentage of households reporting a method as among the main ways is less affected by the change in the survey question.



Loretta Mester is senior vice president and director of research at the Federal Reserve Bank of Philadelphia.

EXHIBIT 1

Percent of U.S. Households That Use Each Instrument: 1995 vs. 1998^a

	ATM ^b		Debit Card		Direct Deposit		Automatic Bill Paying		Smart Card		Any of These	
	1995	1998	1995	1998	1995	1998	1995	1998	1995	1998	1995	1998
All Households	61.2%	67.2%	17.6%	33.8%	46.8%	60.5%	21.8%	36.0%	1.2%	1.9%	76.5%	85.5%
By Age:												
Under 30 years old	71.1%	75.6%	24.5%	45.0%	31.1%	45.2%	17.9%	30.5%	1.8%	2.6%	75.2%	80.2%
Between 30 and 60 years old	67.2%	75.9%	19.7%	38.6%	42.9%	58.0%	24.5%	38.6%	1.5%	2.3%	77.4%	87.4%
Over 60 years old	43.1%	41.6%	9.6%	16.0%	63.2%	74.8%	18.2%	33.0%	0.3%	0.5%	75.2%	83.7%
By Income ^c :												
Low income	36.0%	45.7%	7.1%	19.7%	32.7%	44.3%	9.8%	17.1%	0.8%	1.5%	54.5%	69.1%
Moderate income	60.1%	64.1%	16.0%	31.6%	43.1%	58.8%	17.7%	30.5%	0.6%	3.1%	77.0%	87.2%
Middle income	69.4%	71.9%	20.3%	36.6%	48.3%	66.1%	23.4%	42.8%	1.3%	2.0%	83.6%	89.4%
Upper income	76.6%	82.1%	25.0%	43.8%	58.3%	70.4%	32.0%	49.3%	1.8%	1.7%	89.1%	94.8%
By Education												
No college degree	52.8%	59.9%	14.3%	29.2%	40.4%	54.4%	18.2%	30.2%	0.8%	1.8%	69.8%	80.7%
College degree	80.1%	81.9%	25.2%	43.1%	61.0%	72.6%	30.1%	47.7%	2.1%	2.0%	91.5%	95.1%

^aThe percentages reported are based on the population-weighted figures. (For further discussion see the Survey of Consumer Finances codebooks at www.bog.frb.fed.us/pubs/oss/oss2/98/scf98home.html and www.bog.frb.fed.us/pubs/oss/oss2/95/scf95home.html.)

^bThe question on ATMs asked whether any member of the household had an ATM card, not whether the member used it. The other questions asked about usage. Source: 1995 and 1998 Survey of Consumer Finances, Federal Reserve System.

^cSee note on Exhibit 2.


sharply: whereas half the households used an electronic method as one of their main ways of conducting business in 1995, over three-fourths did so in 1998. Use of ATMs became more popular over the three years, and the phone showed an even greater increase in popularity. Youth, high income, and a college degree continue to be associated with a higher incidence of computer banking, but, perhaps surprisingly, the computer remains a less popular means of doing business with financial institutions. 



EXHIBIT 2

Percent of U.S. Households with at Least One Financial Institution Using Each Method Among the Main Ways of Conducting Business with at Least One of Their Financial Institutions^a

	In Person		Mail		ATM		Phone		Computer		Electronic ^b	
	1995	1998	1995	1998	1995	1998	1995	1998	1995	1998	1995	1998
All Households	85.4%	76.3%	56.5%	34.0%	33.8%	52.3%	25.7%	51.0%	3.7%	4.0%	49.7%	78.2%
By Age:												
Under 30 years old	76.9%	69.1%	57.8%	32.2%	53.0%	67.9%	20.8%	52.8%	5.2%	5.6%	62.3%	78.7%
Between 30 and 60 years old	86.7%	78.6%	62.0%	37.6%	37.8%	61.2%	28.1%	58.3%	4.6%	5.0%	53.8%	81.6%
Over 60 years old	86.6%	74.5%	44.0%	25.9%	16.2%	22.3%	23.0%	32.4%	1.2%	0.6%	35.0%	69.6%
By Income ^c :												
Low income	81.2%	65.1%	32.6%	20.3%	19.6%	34.0%	13.4%	28.5%	1.3%	1.2%	29.8%	61.9%
Moderate income	85.8%	77.2%	48.6%	28.0%	29.8%	47.8%	18.7%	42.2%	1.9%	2.3%	41.1%	76.7%
Middle income	85.6%	79.2%	56.8%	33.2%	37.8%	53.4%	22.8%	57.6%	4.0%	2.9%	53.3%	81.1%
Upper income	87.6%	81.6%	74.3%	45.5%	42.3%	65.1%	37.8%	66.0%	5.9%	6.9%	63.9%	87.7%
By Education												
No college degree	85.7%	76.0%	49.4%	29.7%	27.5%	44.7%	19.7%	43.3%	2.8%	1.8%	41.3%	73.1%
College degree	84.6%	76.9%	71.0%	42.0%	46.8%	66.5%	38.1%	65.5%	5.6%	8.0%	66.9%	87.8%

^aReferring to each financial institution with which the household does business, the survey asked: "How do you mainly do business with this institution?" Respondents could list multiple methods, with the main method listed first. This table reports on all the methods a respondent listed for each of the household's financial institutions. The percentages reported are based on the population-weighted figures. Note, the percentages do not add up to 100 percent across columns, since households could list more than one method and more than one financial institution.

^bIn 1995, electronic refers to ATM, phone, payroll deduction and direct deposit, electronic transfer, or computer. In 1998, electronic refers to ATM, phone (via voice or touchtone), direct deposit, direct withdrawal/payment, other electronic transfer, computer/Internet/online service, or fax machine.

^cLow income is defined as less than 50 percent of the median household income; moderate income is 50 to 80 percent of the median; middle income is 80 to 120 percent of the median; and upper income is greater than 120 percent of the median. Median income was \$32,264 in 1994, the year to which the 1995 survey questions refer. So for the 1995 survey, low income is less than \$16,132; moderate income is \$16,132 to \$25,811; middle income is \$25,811 to \$38,717; and upper income is over \$38,717. Median income was \$37,005 in 1997, the year to which the 1998 survey questions refer. So for the 1998 survey, low income is less than \$18,503; moderate income is \$18,503 to \$29,604; middle income is \$29,604 to \$44,406; and upper income is over \$44,406.

Source: 1995 and 1998 Survey of Consumer Finances, Federal Reserve System.

“We Control the Vertical”: Three Theories of the Firm

BY MITCHELL BERLIN

The Gramm-Leach-Bliley Act has expanded the set of products and services banks and other financial firms can provide to their customers. But how will financial institutions organize their production? In this article, Mitchell Berlin discusses three broad approaches to vertical integration and the pros and cons of a firm’s providing all stages of production and distribution.

How can a bank (or any firm) decide how much of the chain of production and distribution it should carry out on its own? For example, should a bank that offers a line of profitable credit cards handle its own back-office operations, a move that economists call *backward integration*? And when (if ever) should a technology firm that has been content to provide information-processing services to retail financial firms (such as banks) decide to *integrate forward* and provide financial services directly to the public? Questions such as these have become particularly pressing for bankers and

their competitors now that the Gramm-Leach-Bliley Act of 1999 has expanded the products and services that banks and other financial firms can provide.

One standard answer is that complementary products and services are more profitably produced by a single firm, but the real answer is not that simple.¹ True, complementary activities are often carried out by vertically integrated firms,² but they are also carried out by separate firms specializing in single stages of production or distribution. A prominent example is life insurance. Underwriting insurance and

selling insurance are almost surely complementary activities. But insurance companies market policies through two different channels. They use agents who sell only their own company’s products, and they also use independent agents who sell the products of multiple insurance companies.³

Although this article focuses on vertical integration, particularly forward integration into retailing, many of the same issues arise when a firm decides whether to expand its product line to include products or services that are not vertically related.⁴

Economists have been puzzling over this issue for more than 70 years under the general rubric the *theory of the firm*.⁵ In a classic article in 1937, Nobel Prize winner Ronald Coase first posed the (seemingly) simple question: When will a transaction be carried out within a single firm rather than by two separate firms transacting in the market?

Although economists have proposed a multitude of theories to answer this question, there is some common ground: The best way to structure a transaction depends on how it affects the incentives of the parties to the transaction. (But see *Other Models*



Mitchell Berlin is research officer and economist in the Research Department of the Philadelphia Fed.

¹Activities are *complements* when doing one reduces the cost of doing the other. For example, originating a loan and providing credit insurance to the borrower *may* be complementary financial activities, since information about the borrower is reusable.

²A vertically integrated firm is one that carries out two or more stages of production or distribution by itself. For example, a firm that both produces and sells its own products is vertically integrated.

³Allen Berger, J. David Cummins, and Mary Weiss’s article presents empirical evidence that the coexistence of these two delivery systems is not merely because one inefficient system survives alongside an efficient one.

⁴See my earlier article for a review of some of the recent empirical evidence on the costs and benefits of product line specialization and diversification.

⁵Bengt Holmstrom and John Roberts’ literature review is an excellent critical discussion of the theory of the firm.

of *Vertical Integration* for explanations that don't focus on incentives.) The types of incentives that theories of the firm have emphasized include individuals' willingness to cooperate in response to unforeseeable events, their willingness to provide maximum effort, and their willingness to allocate their time and attention in a sensible way across a multitude of competing tasks.

Three broad approaches have had the most influence in recent years: the transactions cost approach, the property rights approach, and the multitask approach. In this article I will use each of these approaches, in turn, to examine a firm's decision to sell through independent sales representatives or through an in-house sales group. In other words, I will look at a firm's decision about whether to integrate forward into retailing.

IN-HOUSE SALES FORCE OR INDEPENDENT REPS?

AdaptorDie Corp. (AC) is a manufacturer of electronic components with a broad menu of products.⁶ Like many other firms in the components business, AC uses two different channels for selling its products.

Twenty percent of its components are sold through an in-house sales force that is paid a fixed salary plus a modest commission for each sale. Following an almost universal pattern, AC's own sales employees sell only AC components.

The remaining 80 percent of AC's components are sold through independent manufacturers' representatives such as DirectCell Corp. (DC). DC's relationship with AC is not exclusive. DC offers a sales package that includes AC's products and also those of

Other Models of Vertical Integration

B

roadly, the industrial organization literature has proposed three motives for vertical integration, in addition to theories that rely explicitly on incentive or bargaining considerations.^a

One classic motive is firms' desire to avoid *double marginalization* (or the *chain of monopolies* problem). Consider a monopolist manufacturer of electronic components selling to a retailer who is also a monopolist. If the manufacturer can't charge a fixed fee, in addition to its price per unit sold, it will charge the retailer a price above the marginal cost of production. Thus, the retailer faces a higher price for the component than it would if it could purchase the good in a competitive market. In turn, the retailer charges final customers the monopoly price (the price it paid for the component, plus its own monopoly markup).

The problem with this outcome is that firms produce too little output at too high a price, from the standpoint of both the firm and the consumer. Firms would increase their joint profits and increase consumer satisfaction as well if they integrated and sold the good at a lower price. The integrated firm would maximize profits by charging consumers a monopoly markup over the marginal cost of producing the good, a price lower than the one that includes the double markup charged by successive monopolies.^b

The main criticism of double marginalization as a convincing motive for vertical integration is that there are straightforward contractual solutions to the problem that don't involve vertical integration. For example, the manufacturer could set its price equal to the marginal cost of producing each unit of the good and charge the retailer a fixed fee (as compensation for giving up monopoly profits). To be a fully convincing motive for vertical integration, double marginalization requires some reason such contracts are infeasible, for example, the contractual difficulties highlighted by the theories in the text.

A second classic motive for vertical integration is *vertical foreclosure*. Consider the case of a monopolist supplier of an input necessary to competitive manufacturers. One of the manufacturers could integrate with the monopolist supplier to gain a competitive advantage in the market for the final good. Indeed, the integrated firm might even be able to drive other manufacturers out of the market. Although this is a plausible motive for integration, the conditions for vertical foreclosure — a monopolist provider of an input — are clearly restrictive and can't explain many cases of vertical integration observed in the marketplace. For example, General Motors produces more of its auto components internally than does Ford, but not because the nature of the inputs is fundamentally different.

A third classic motive for vertical integration is *supply assurance*.^c According to this motive, a manufacturer may wish to guarantee an adequate supply of an input in the face of uncertainty about its own requirements (say, because of fluctuating demand for its own manufactured good) or in the face of uncertainty about the total supply of the input. Like double marginalization, supply assurance appears to be a sensible explanation for vertical integration, but only in the presence of the types of contracting and bargaining problems that take center stage in the approaches discussed in this article. Without such problems, firms could write relatively simple contracts to ensure an adequate supply of inputs. For example, uncertain demand for electricity has not led to vertical integration but to long-term contracts between electric utilities and coal suppliers.

^aJean Tirole's excellent textbook has a full chapter devoted to models of vertical integration.

^bDouble marginalization is really a special case of a more general class of distortions covered at length by Tirole under the general heading "the basic vertical externality."

^cSome prominent formal models of the supply assurance motive include the article by Dennis Carlton and the one by Patrick Bolton and Michael Whinston.

⁶My description of AdaptorDie and DirectCell draws heavily on Erin Anderson and David C. Schmittlein's article. Both firms are fictional.

other electronics manufacturers, although, following common practice, it doesn't offer competing versions of the same component. DC rents its own office space and hires its own sales staff. All components are sold on straight commission, that is, AC pays DC a percentage of the sale price for each component sold. However, the unsold products remain the property of the manufacturer.

AC's use of two distinct sales channels raises some obvious (and not so obvious) questions. What types of components are sold using each channel? Why does AC use different compensation schemes for the two channels? (Its internal sales force works for fixed salaries while independent reps receive a percentage of sales.) A less obvious question — but only because we may not think to question a practice that is so common — is why AC's own sales force is not permitted to sell other firms' components while DC's sales force has a nonexclusive sales relationship? The answers to questions like these may shed light on the broader question: What are the relative advantages and disadvantages of each type of sales channel?⁷

IT'S THE TRANSACTION, STUPID

The *transactions cost* approach argues that the answers to these questions can be found by looking at

the details of individual transactions between different firms, for example, a contractual agreement by DC to sell AC's capacitors as part of its sales package.⁸ This approach says that costs inevitably arise as firms bargain and disagree in the normal course of conducting business in a rapidly changing marketplace. Transaction costs are distinct from production costs such as AC's manufacturing costs or DC's sales costs. Broadly, transaction costs include all expenses and foregone opportunities that arise because of actual bargaining and dickering as well as expenses borne to avoid potential disagreements. These costs range from lost sales when disagreements lead to delays to lawyers' fees when negotiations become so contentious that the courts (or even the threat to go to the courts) come into play.

The transactions cost approach begins with a straightforward thought experiment: For a particular transaction — such as the sale of capacitors — we can evaluate the transaction costs that would arise if separate firms carried out

component will be sold through AC's internal sales force instead, because incentives to disagree can more easily be mitigated or overcome within a firm.⁹

Asset Specificity Creates Transaction Costs. The most important determinant of transaction costs is the so-called *degree of asset specificity*: the extent to which the transacting firms invest in assets whose value depends on the business relationship's remaining intact.

Both AC and DC have made numerous investments in the course of their business relationship. For example, DC invested in office equipment, including a personal computer for each of its sales agents. A computer is a nonspecific asset that can be used to store and process information about accounts for any manufacturer; that the computers currently store information about orders for AC doesn't affect the value of the computer. And if AC were to replace DC and find a new manufacturers' rep, it wouldn't spend much time wondering whether it could find a firm with equally powerful computers.

If [transaction costs] are high enough, the theory predicts that the component will be sold through [an] internal sales force instead, because incentives to disagree can more easily be mitigated or overcome within a firm.

the transaction, that is, if the capacitor were produced by manufacturer AC and sold by independent sales representative DC. If these costs are high enough, the theory predicts that the

However, firms also invest in assets that would lose much, if not all, of their value if the business relationship broke down. For example, some

⁷The real world organizational choice is somewhat more complicated than my description of a choice between an internal sales force that sells the firm's products exclusively and the independent sales rep that has nonexclusive relationships with many firms. Traditional franchises, such as gas stations, are independent retailing firms that have exclusive sales relationships with gasoline producers. Also, Dell's internal sales personnel sell the products of other computer companies, for example, Hewlett Packard printers. These organizational forms may be thought of as intermediate contractual solutions.

⁸Seminal contributions in the transactions cost approach have been made by Oliver Williamson and by Benjamin Klein, R.A. Crawford, and Armen Alchian. See Williamson's 1985 book for a critical summary of the literature in this tradition.

⁹This is too simple, since difficulties also arise when a transaction is handled within a single firm. We discuss these below, but the transactions cost approach has a less distinctive and less complete analysis of the various costs of keeping transactions carried out within the firm.

electronic components are not standardized and must be modified to suit a particular customer's needs. Customization is a process that requires an understanding of the customer's needs, an understanding of which modifications are feasible, and a channel for communications between product designers and customers. Thus, the sales agent must have an intimate knowledge of the final customer's business and must also have a working relationship with the manufacturer's designers and engineers.

Knowledge about customers and working relationships, both of which take time and effort to build and nurture, are examples of (intangible) *relationship-specific*, or *idiosyncratic*, assets.¹⁰ If AC and DC parted ways in a dispute over the feasibility of a customer's demand for a product modification: (1) AC would lose its storehouse of knowledge about the customers who traditionally purchase through DC, and (2) the working relationships between DC's sales force and AC's engineers would go up in smoke.

With so much to lose on both sides, a complete breakdown in the business relationship is unlikely. But that won't prevent the firms from haggling over who receives the gains and who bears the brunt of making adjustments. And even if disagreements don't typically lead to a split, haggling can be time-consuming and expensive. Perhaps more significant, if everyone expects lots of disagreements, or if the adjustments lead to an unequal distribution of the net gains, firms may avoid making idiosyncratic investments in the first place; that is, the willingness to make

valuable investments is undermined because of individuals' unwillingness to cooperate.

For example, the head of AC's engineering division can centralize (and restrict) communications between AC's engineers and DC's sales force. If communication is difficult, working relationships between sales personnel and designers may never develop, and potentially profitable product adjustments may never get proposed. The lost profit from inflexible product design should be reckoned an indirect transaction cost.

In a predictable business environment, transactions will be largely routine, and firms may be able to write contracts that specify each party's rights and obligations.

Uncertainty Increases the Costs of Haggling. In a predictable business environment, transactions will be largely routine, and firms may be able to write contracts that specify each party's rights and obligations. In an unpredictable world, things are very different. Changing circumstances require adaptive responses, and it may be impossible to write contracts that are both flexible enough to permit adjustments and precise enough to give

adequate guidance when disagreements arise.¹¹ For example, for nonstandardized products no one can accurately predict which customers will seek customized variants and which changes they will demand. Thus, for nonstandardized components, uncertainty about the future is great, and transaction costs are likely to be high; time-consuming haggling will be a problem, and the risks of a bargaining breakdown are likely to be great.

Customized Products Are Sold by In-House Sales Forces. Let's push our thought experiment about AC further. Divide AC's products into two groups: those that are standardized and those that are customized to meet customers' demands. Since customized products require an unpredictable series of adaptations, and since these will require significant investments in idiosyncratic assets (knowledge about customer needs and relationships between design and sales personnel), transaction costs are likely to be high if such components are sold through independent sales rep DC. Theory predicts that these products are more likely to be sold through the in-house sales staff, while AC's standardized products will be sold through DC.

Indeed, this prediction is supported by the work of Erin Anderson and David Schmittlein, who examined a real-world electronics firm and found that it sells customized components through its in-house sales force and standardized components through manufacturers' reps.¹²

In-House Sales Divisions Have Advantages When Transaction Costs Are High. If difficulties can afflict transactions between separate firms, how can keeping the transaction

¹⁰Idiosyncratic assets needn't be intangible. The textbook example of an idiosyncratic asset is Fisher Body's plant for producing auto chassis located right next to Chevrolet's assembly operation.

¹¹In the economics literature, contracts that do not include detailed clauses to cover all contingencies are called *incomplete contracts*. Typically, contracts are incomplete because (1) it is very hard to specify all contingencies in advance, and (2) it is difficult to describe contingencies with sufficient clarity that a court can actually enforce the contract.

¹²Bengt Holmstrom and Paul Milgrom interpret this evidence differently as we will see below.

within a single firm be better? One possibility is that information may flow more easily within an organization than between organizations. Since members of a single organization all share in the same flow of profits, cooperation may be more the norm, and sales personnel and designers within the same firm may have fewer incentives to withhold information.¹³ This cooperation can be facilitated through judiciously designed internal compensation schemes. Furthermore, to the extent that the information flows are the result of conscious design, the firm's top management can encourage the free flow of useful information. This may be particularly important when disputes arise, as they inevitably will, between separate divisions of a single firm.

Apart from the value of a freer flow of information, the resolution of disputes between divisions may be eased by the use of *administrative fiat*; that is, top management can impose an outcome when the sales and engineering divisions can't come to agreement on their own. Firms often have specialized internal mechanisms for handling more serious disputes. For example, in a case study of a high technology firm's choice between purchasing inputs from external sources and producing the inputs internally — sometimes called the *make or buy decision* — Marc Knez and Duncan Simester found that disputes between the engineering and sales divisions were arbitrated by “chief technologists,” usually former engineers who acted as final judge. Knez and Simester also found that in this firm, arbitration was never used to resolve disputes between outside producers of inputs and the purchasing firm.

¹³The property rights theorists discussed in the next section disagree strongly with these claims.



ASSETS ARE POWER

The *property rights* approach also addresses the problem of designing organizations to mitigate the effects of the disputes that inevitably arise in the course of doing business in an unpredictable world. But the property rights view doesn't agree that keeping a transaction within the firm is more likely to lead to improved information flows or greater incentives for individuals to cooperate.¹⁴ Accordingly, the underlying thought experiment differs from that of the transactions cost approach. Property rights theorists begin with separate organizations — separate firms or even separate divisions of the same firm — and with a description of the various assets needed to carry out business

¹⁴The property rights approach ignores differences in the way information flows between firms and within firms and also ignores differences in the details of compensation schemes.

transactions. For a property rights theorist, the underlying question is: “Who should own which assets?”¹⁵

To illustrate this approach, let's go back to our example. There are three distinct organizations: AC's manufacturing division, AC's sales division, and independent manufacturers' rep, DC. To simplify, imagine that there are just two types of assets: the machines used to produce electronic components and the sales lists of customers who have bought each component in the past. In the property rights view, AC and DC would be viewed as independent firms only if AC owns the machines and DC owns the customer list. The fundamental feature that distinguishes DC from AC's internal sales division is that DC owns its

¹⁵This approach was originated by Oliver Hart and Sanford Grossman and was developed subsequently by Oliver Hart and John Moore in a series of articles. See Hart's 1995 book for an accessible introduction.

customer list while AC's sales division doesn't.

Owners Are Powerful. In the property rights view, if someone owns an asset, he or she has the power to exclude others from using it.¹⁶ For example, the owner of the customer list controls an essential link between AC and its potential customers; knowing who has purchased a product in the past is very valuable knowledge when AC wants to market an advanced version of an existing component. Without this knowledge AC would have to rely on expensive scattershot methods to inform potential customers of the new product (e.g., commercial time during the Super Bowl).

Consider the case where AC and DC are separate firms, that is, DC owns the customer list, and imagine that one of DC's important customers insists on a costly customization of a component. DC's credo is: "The customer is always right!" But AC's engineers are under severe pressure to keep costs in line after a number of embarrassing cost overruns. AC's engineers prefer that the customer accept a modest customization (at most) and that the costs of producing the tailor-made product be charged to the customer. DC's sales personnel argue that they have made promises that AC would work closely with the customer to adapt the component.

¹⁶Proponents of the property rights view are well aware that this definition is a simplification, but one chosen to facilitate formal analysis. In a recent contribution, Raghuram Rajan and Luigi Zingales have proposed a more nuanced view of ownership in which the owner can regulate access to an asset and also regulate access to those who work with the asset. This extension is important because, in many cases, it is the knowledge and expertise of the team of employees working with an asset that is most valuable rather than ownership of the asset itself. Phillippe Aghion and Jean Tirole propose another significant extension and distinguish formal ownership and real ownership — which requires the owner to be well informed enough to know how to make good use of the asset.

Who will win this dispute?

The owner of the customer list, DC, is likely to win because it has a lot of bargaining power. It holds the (mainly implicit) threat to walk away with the customer list and to peddle its services to

conflicts than in conflicts with an independent manufacturers' rep.

All Power to the People (Whose Efforts Are Most Valuable). If disputes, large and small, cordial and hostile, are the warp and the woof of

Conflicts are pervasive both within and between firms. But an internal sales division doesn't have the same bargaining power as an independent firm, since the head of the sales division can't walk away with the sales list.

one of AC's competitors.¹⁷ This threat is a powerful one, since AC may find it very difficult to quickly re-establish a channel to existing customers without the customer list.

After enough disputes like this, the head of engineering at AC might well ask herself whether it wouldn't be more profitable to sell a larger share of its products through its own in-house sales division. The property rights view emphasizes that the head of engineering would be naïve to expect that conflicts would be less pervasive within a single firm or that internal sales personnel would be more concerned about the engineering division's cost control efforts than would independent sales personnel. Conflicts are pervasive both within and between firms. But an internal sales division doesn't have the same bargaining power as an independent firm, since the head of the sales division can't walk away with the sales list.¹⁸ Thus, the head of engineering can reasonably expect to prevail more often in internal

business life, it is easy to see why the head of engineering at AC or the president of DC would be concerned about who owns which assets. However, from an efficiency standpoint — that is, if we are primarily concerned about increasing individuals' incentive to make jointly valuable investments — it is not immediately clear that it matters who prevails more often. Cutting costs and satisfying unique customer needs are both worthy business goals, and how the engineers of AC, the sales personnel of DC, and the sales personnel of AC's sales division divide up the profits doesn't seem to be an important issue to anyone but themselves.

In fact, relative bargaining strengths *do* matter because members of the different organizations invest time and effort that increases the value of their *joint* output, but each member's willingness to make such investments depends on his or her *own* expected return on the investment. The central idea of the property rights view is that bargaining power — and the assets that confer bargaining power — should be in the hands of those people whose efforts are most significant in increasing the value of the business relationship. Giving these people more bargaining power ensures that they receive more of the

¹⁷Economists call this the "hold-up problem."

¹⁸What's to keep the head of sales from walking away with the customer list? Noncompete clauses are customary in situations such as these.

rewards from investing time and energy and, thus, that they have a stronger incentive to make these investments.¹⁹

Who Should Own the Customer List? Two Examples.

Consider a component that comes in many varieties, each one tailored to a particular type of user, or one that requires extensive follow-up service. Here the relationship between the sales organization and the customer is paramount. Sales personnel must know their customers' needs; indeed, the seller may play a significant role in advising the customer, both at the time of sale and after.

For products such as these, the most important investments are made by the sales organization, and the sales organization should own the customer list to capture a larger share of the profits from providing excellent customer service. Theory would predict that these products should be sold by DC, rather than by AC's in-house sales organization.²⁰ Increasing DC's bargaining power can also increase AC's profits, even if AC's engineers often have difficulty winning disputes. DC's investments in customer service also increase the value of AC's investments in product design. Thus, giving DC lots

of bargaining power can significantly increase the sum of AC's and DC's firm-specific investments and the *total profits* to be divided.²¹

Consider another class of components in which AC is a leader in product innovation, but which are mainly standardized and which require little follow-up servicing. In this case, a knowledgeable sales force may still be necessary to educate customers about new products, but a close relationship between sales personnel and customers

Thus, providing high-powered rewards for easily measured outputs will lead individuals (and their organizations) to neglect tasks that may be important but whose results are difficult to measure.

is not as important. For these components, the customer list should not be owned by an independent sales organization such as DC. The customer list is still very valuable, and DC's bargaining power would mainly undercut AC's engineers' incentive to work hard without a countervailing gain.

RICHES BEYOND MEASURE

A third approach, the *multitask* approach, doesn't view bargaining and

hold-up problems as the central influence on the design of firms. Instead, this approach draws out the implications of a simple, but powerful pair of insights: (1) Most managers and employees are engaged in the production of many outputs — or, viewed differently, engaged in a variety of tasks; (2) Some outputs are easy to measure, and some are hard to measure.²²

For example, the members of any sales force are really engaged in a variety of activities when they sell a component. The most obvious is the sale itself, something that is relatively easily measured. But sales personnel also collect information about customers' needs and problems, and this information can be tremendously valuable to the product's designers and engineers. Unlike booking a sale, diligence and ingenuity in collecting information are hard to measure. Of course, these activities will ultimately be reflected in future sales, but the effects may take a long time to come to fruition, and they will be spread widely.

When Measurement Is Difficult, Low-Powered Incentives Are Best. Consider a compensation scheme such as the one between AC and DC, a straight percentage commission on total sales. When sales personnel are heavily rewarded for the volume of sales, they will predictably allocate their time and attention to selling, and they will neglect the less rewarding task of collecting intelligence to be passed on to AC's designers. Thus, providing high-powered rewards for easily measured outputs will lead individuals (and their organizations) to neglect tasks that may

¹⁹However, David de Meza and Ben Lockwood's article demonstrates that the question of who should be given power over assets may be slightly more complicated than this discussion suggests. In particular, we need to know more about the details of the bargaining environment to make precise predictions.

²⁰Note that Anderson and Schmittlein's evidence that standardized electronic components are sold by independent manufacturers' representatives is inconsistent with this interpretation. However, for the most part, it has proved relatively difficult to devise convincing and powerful empirical tests that distinguish one theory's empirical predictions from another's. See Michael Whinston's paper for one attempt to devise a formal framework for distinguishing between the predictions of the transactions cost and property rights approaches.

²¹In the language of the formal theory, individuals' investments in human capital are complements. One individual's investment in human capital raises the marginal return to investment in human capital for other individuals.

²²The multitask approach has been developed in a series of papers by Bengt Holmstrom and Paul Milgrom, although the insight that measurement problems are central to a theory of the firm figures prominently in Yoram Barzel's work. Milgrom's paper with John Roberts on influence costs — politicking within firms — is another application of the multitask approach.

be important but whose results are difficult to measure.

For some products, collecting intelligence from AC's customers and transmitting it to AC's designers is very important. According to the multitask approach, such products should be sold by the in-house sales division, not DC's sales reps. And the lion's share of the in-house sales staff's compensation should be a fixed salary — a low-powered compensation scheme, in the sense that pay is not closely related to measured performance.²³

Of course, this compensation scheme has inevitable drawbacks, too. Internal sales personnel may allocate their time more appropriately between activities with short- and long-term payoffs, but they may simply work less hard than DC's sales force. AC can partially counter this drawback and also directly reward employees' effort on difficult-to-measure activities by using compensation and promotion schemes tied to subjective performance evaluations by supervisors.²⁴

Measurement Difficulties Also Help Explain Job Design. The multitask view suggests that jobs may also be designed differently depending on whether components are sold by an in-house sales force or by independent manufacturers' representatives. For example, if one salesperson offers the products of numerous producers, each producer will worry that its own product

is being shortchanged. Many firms avoid this problem by using an in-house sales force that sells its firm's products exclusively.

But this only raises another question. Why not insist on an exclusive sales relationship with the independent manufacturers' rep also? Why doesn't each sales organization agree to sell only one manufacturer's product line at a time?²⁵ One reason is that for some types of products, there isn't a lot to gain from an exclusive relationship. When it is easy to measure and reward a salesperson's effort in selling a product, an exclusive relationship isn't necessary. Sales figures will accurately reflect the time and effort that sales personnel have spent in selling each manufacturer's products. For some products AC can simply examine DC's sales of AC's components to make sure that DC is not promoting a competitor's product at AC's expense.

exclusivity on its sales employees. The simplest way to make sure that the employee is not shifting time and attention to promote a competitor's components is to impose exclusivity. While it may be hard to keep close tabs on an employee's allocation of time among various sales activities, it is relatively easy to check whether he or she is selling another firm's products on company time.

CONCLUSION

Recent theories of internal organization offer some general lessons for a financial services firm that's thinking about moving into new product lines by integrating backward or forward or by selling related products. Perhaps the most important general insight is that while engaging in complementary activities may be part of the rationale for expanding a firm's activity mix, the firm must also take serious account of

While engaging in complementary activities may be part of the rationale for expanding a firm's activity mix, the firm must also take serious account of incentives.

On the other hand, the goods handled through in-house employees (and not by independent reps) are sold that way precisely because sales figures are not good measures of sales effort for some types of components. The same measurement problems that cause AC to use internal sales personnel and low-powered incentive schemes for some products also cause AC to impose

incentives. The Gramm-Leach-Bliley Act has introduced a new category of activities, *activities complementary to banking*, that could be provided through financial holding companies. The act requires that bankers seeking to engage in a new, complementary activity must describe the nature of the complementarity in detail. Although regulators will not demand that bankers analyze incentive considerations when new activities are brought into the firm, bankers themselves would be well advised to take their own analyses further than the law demands.

Indeed, the starting point of Oliver Williamson's investigations into

²³Anderson and Schmittlein's characterization of the difference between compensation schemes for in-house sales employees and independent manufacturers' reps is clearly consistent with this theory.

²⁴See Robert Gibbons's and Canice Prendergast's surveys for evidence on the use of subjective evaluation in the workplace. An interesting feature of employing subjective evaluations is that internal politicking to influence supervisors' evaluations becomes a significant problem in organizational design. See Milgrom and Roberts' article on influence activities.

²⁵Actually, this hypothetical arrangement has features in common — notably exclusivity — with traditional retail franchise relationships. See Francine Lafontaine and Margaret Slade's article for an evaluation of the empirical literature on retail franchising.


the transactions cost motivations for vertical integration was the recognition that even highly complementary activities could be profitably carried out by separate firms if the terms of the transactions between the firms are relatively predictable and don't require relationship-specific investments. In these cases, specialized firms may well achieve a high degree of coordination through contracts alone, because individuals' incentives to disagree are small and easily overcome.

While the theoretical work is not yet sufficiently well developed to give bank managers precise guidance, the theories do yield some important insights. One of the central insights of the property rights view is that bargaining problems don't just disappear when transactions are brought within a single firm. The key is to assign ownership rights over assets to those whose effort produces the most value for the firm. For example, many bank holding companies have discovered that purchasing an investment banking subsidiary doesn't improve coordination between commercial bankers and investment bankers — one of the ostensible benefits of having a single company handle both activities — and may simply multiply tensions.²⁶ A stand-alone investment firm or one with significant autonomy within a holding company structure may be more realistic, because it may allow investment bankers to capture a larger share of the rewards from their customer relationships.²⁷

The multitask approach teaches that for complex products,

²⁶In the business press and management literature, problems like these have often been classified as "cultural conflicts."

organizations face difficult tradeoffs between providing incentives for maximal effort and promoting cooperation and other hard-to-measure activities. An example of these tradeoffs is the continuing tension between the sales-oriented activities of commercial lenders — it is easy to measure and reward lenders according to new accounts gained — and their responsibility to closely monitor credit quality — it is intrinsically more difficult to measure careful monitoring of the credit risk of existing accounts. Generally, the multitask approach states that difficult-to-measure tasks should be handled by employees, with subjective evaluations supplementing otherwise low-powered incentive schemes. The approach also suggests that, to the greatest extent possible, easy-to-measure tasks and difficult-to-measure tasks should be assigned to separate individuals or groups.

While the theories yield many practical insights, economists who have been influential in developing the theory of the firm argue that testing theories against each other in careful empirical studies is the most immediate task at hand.²⁸ Such testing should yield more refined economic insights into existing business practices and hopefully more refined guidance to businesses making practical decisions. 

²⁷Of course, it may be difficult for the front office to precommit not to meddle in the affairs of any of its supposedly autonomous affiliates.

²⁸Michael Whinston's paper attempts to formalize the different predictions of the property rights view and the transactions cost view within a common framework. His conclusion is that the theories have not yet been developed with sufficient precision to distinguish them empirically.



REFERENCES

- Aghion, Phillippe, and Jean Tirole. "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105, 1994, pp. 371-402.
- Anderson, Erin, and David C. Schmittlein. "Integration of the Sales Force: An Empirical Examination," *Rand Journal of Economics*, 15, Autumn 1984, pp. 385-95.
- Barzel, Yoram. "Measurement and the Organization of Markets," *Journal of Law and Economics*, 25 (April 1982), pp. 27-48.
- Berger, Allen N., J. David Cummins, and Mary A. Weiss. "The Coexistence of Multiple Distribution Systems for Financial Services: The Case of Property Liability Insurance," *Journal of Business*, 70, 1997, pp. 515-47.
- Berlin, Mitchell. "Jack of All Trades? Product Diversification in Nonfinancial Firms," Federal Reserve Bank of Philadelphia *Business Review* (May/June 1999), pp. 15-29.
- Bolton, Patrick, and Michael Whinston. "Incomplete Contracts, Vertical Integration, and Supply Constraints," *Review of Economic Studies*, 60, January 1997, pp. 121-48.
- Carlton, Dennis. "Vertical Integration in Competitive Markets Under Uncertainty," *Journal of Industrial Economics*, 27, March 1979, pp. 189-209.
- Coase, Ronald A. "The Nature of the Firm," *Economica*, 1937, pp. 396-405.
- De Meza, David, and Ben Lockwood. "Does Asset Ownership Always Motivate Managers? Outside Options and the Property Rights Theory of the Firm," *Quarterly Journal of Economics*, May 1998, pp. 361-86.
- Gibbons, Robert. "Incentives in Organizations," *Journal of Economic Perspectives*, 12, Fall 1998, pp. 115-32.
- Grossman, Sanford J., and Oliver D. Hart. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94, 1986, pp. 691-719.
- Hart, Oliver. *Firms, Contracts, and Financial Structure*. Clarendon Lectures in Economics, Oxford: Oxford University Press, 1995.
- Hart, Oliver, and John Moore. "Property Rights and the Theory of the Firm," *Journal of Political Economy*, 1990, pp. 1119-58.
- Holmstrom, Bengt. "The Firm as a Sub-Economy," *Journal of Law, Economics, and Organization*, 15, 1999, pp. 74-102.
- Holmstrom, Bengt, and Paul Milgrom. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, and Organization*, 7, 1991, pp. 24-52.
- Holmstrom, Bengt, and John Roberts. "The Boundaries of the Firm Revisited," *Journal of Economic Perspectives*, 12, Fall 1998, pp. 73-94.
- Klein, Benjamin, R. A. Crawford, and Armen A. Alchian. "Vertical Integration, Appropriable Rents, and the Competitive Contracting Practice," *Journal of Law and Economics*, 21, October 1978, pp. 297-326.
- Knez, Marc, and Duncan Simester. "Direct and Indirect Bargaining Costs and the Scope of the Firm," Working Paper, University of Chicago, June 2000.
- Lafontaine, Francine, and Margaret E. Slade. "Retail Contracting: Theory and Practice," *Journal of Industrial Economics*, 45, March 1997, pp. 1-25.
- Milgrom, Paul, and John Roberts. "An Economic Approach to Influence Activities in Organizations," *American Journal of Sociology*, 94, Supplement, pp. S154-79.
- Prendergast, Canice. "The Provision of Incentives in Firms," *Journal of Economic Literature*, 38, March 1999, pp. 7-63.
- Rajan, Raghuram G., and Luigi Zingales. "Power in a Theory of the Firm," *Quarterly Journal of Economics*, 113, May 1998, pp. 387-432.
- Tirole, Jean. *The Theory of Industrial Organization*. Cambridge, MA: MIT Press, 1992.
- Whinston, Michael. "On the Transactions Costs Determinants of Vertical Integration," Working Paper, Northwestern University, 1997.
- Williamson, Oliver. *The Economic Institutions of Capitalism*. New York: Free Press, 1985.

The Gains from International Risk-Sharing

BY KEITH SILL

Do residents of different countries trade financial assets to insure themselves against country-specific risks? In this article, Keith Sill examines the degree of such risk-sharing and whether there could be further gains from increased risk-sharing across borders.

Our economy has become increasingly global. We import and export more than ever before. Yet, three facts about international financial transactions, when taken together, pose a puzzle. First, financial capital moves freely across country borders, at least in the case of the developed countries. Returns on similar dollar-denominated assets in different countries are very close to each other — differences in returns have essentially been eliminated because some investors buy and sell assets internationally. Second, residents of most major industrialized countries hold most of their wealth in domestic assets, forgoing the benefits of diversifying their portfolios by including foreign assets. A fundamental tenet of finance

holds that portfolios should be diversified, and presumably, such diversification includes holding foreign stocks and bonds. Third, domestic saving is closely tied to domestic investment. However, if financial capital moves freely across borders, countries that want to invest more than they are saving domestically should be able to borrow from other countries to finance investment, while countries that have excess savings should be able to lend those savings to foreigners. This would mean domestic saving and investment wouldn't necessarily move together, but they are closely linked in the data.

These facts pose a puzzle because we would expect residents of different countries to trade goods, services, and financial assets in such a way as to insure themselves against country-specific risks that affect the amount of goods the country produces (output) and the amount of goods residents can buy (consumption). Insuring against risks is possible because countries' economies do not always move in sync: When one country is in a recession, another may be experiencing

an expansion. So shouldn't the residents of two such countries try to share some of the risk they each face individually, so that people in both countries can be better off?

In this article we will discuss some of the benefits that accrue to residents of a country when economic risk is shared with residents of other countries. We will examine some of the data on the extent of international risk-sharing in developed and developing countries. Those data suggest that the *amount* of international risk-sharing is rather small. This finding leads to another question: Are there significant unexploited gains from risk-sharing? Though the jury is still out, that seems unlikely for financially developed countries. However, for developing countries, it's more likely the case that there are substantial unexploited gains from international risk-sharing.

RISK-SHARING AND THE BENEFITS OF PORTFOLIO DIVERSIFICATION

People prefer to have a relatively steady amount of consumption from year to year rather than wild swings. This preference for smoothing out fluctuations in consumption reveals itself in the pattern of household borrowing and saving. Households often borrow funds or reduce their savings when current resources are low in order to maintain their lifestyle.

Let's take a simple example. Suppose you can choose between consuming \$20,000 of goods a year for the next two years or consuming \$10,000 this year and \$30,000 next year. Most people prefer the first plan, in which the



Keith Sill is a senior economist in the Research Department of the Philadelphia Fed.

amount of consumption is constant over the two years. If a household were stuck with the second plan, it might borrow to increase its consumption today and repay that loan with its higher income in the second year. That way, the household could increase its consumption in the first year by forgoing some consumption in the second year.

What else can a household stuck with the second plan do? One possibility is to find another household with different consumption opportunities and trade with it. For example, another household may get \$30,000 worth of consumption in the first year and \$10,000 in the second. The two households could agree to pool their resources in each year and divide the total down the middle. That way, each household would get \$20,000 worth of consumption in both years.¹ Each household would be better off by agreeing to share the “risk” of fluctuating consumption with the other, thereby lowering, or in this case eliminating, consumption risk.

But the world is actually much more complicated because households face uncertainty about their future resources and typically make plans over long periods — indeed, over their lifetimes. How do households share risk in the real world? It’s much too time consuming and difficult to find other households with which mutually agreeable arrangements can be made and enforced. One alternative is to make such arrangements indirectly through financial markets, for example, by purchasing insurance. In addition, households can purchase stocks and bonds, which represent claims on the assets and revenue streams of businesses (and possibly governments, in the case

of bonds). By purchasing a stock or bond, investors are entitled to a share of those assets or revenues. And it’s all to the better if the payoff from households’ financial assets is high when their income is low. That allows households to smooth out fluctuations in the amount

Households can lower some of the economic risk they face by holding a portfolio of stocks and bonds; however, the portfolio must be diversified.

of goods and services they consume over time. In effect, households can lower some of the economic risk they face by holding a portfolio of stocks and bonds; however, the portfolio must be diversified.

Why is a diversified portfolio so important? Different businesses face different risks. For example, automobile manufacturers face much more business-cycle risk than do electric utility companies, since buying a new car involves discretionary spending on the part of a consumer whereas paying utility bills does not. It’s much more likely that, in a recession, households will forgo buying a new car than forgo the use of heating and telephone services. Similarly, firms that produce agricultural products or build houses are generally more sensitive to adverse weather conditions than are firms that produce steel and plastic. So, an autoworker who wants to diversify some of the income risk he faces may do well to purchase stocks and bonds of firms that typically do well when the auto sector is doing poorly. Likewise, an

agricultural worker may want to purchase a portfolio of assets that offset some of the bad-weather risk he faces.

In short, it is better not to keep all your eggs in one basket — as anyone who had heavily invested in technology stocks in 2000 can tell you. A simple exercise will help us see how valuable diversification can be.

We can measure the risk of holding a stock by the volatility of its return — how much the return varies from month to month. A statistical measure of this risk is the standard deviation, which quantifies the average variability of an asset’s return. The higher the standard deviation, the higher the average volatility of the return and the bigger the swings in the return.

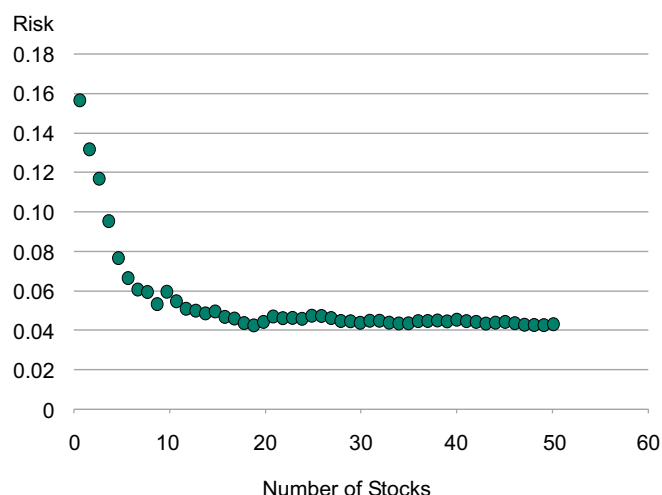
To see how the standard deviation (risk) of a portfolio of stocks declines as the number of stocks in the portfolio increases, let’s look at an experiment first undertaken by economist Eugene Fama in 1976. First, randomly select 50 stocks listed on the New York Stock Exchange (NYSE). Out of this sample, randomly choose one stock, and calculate the standard deviation of its return over a month. That number is a measure of the risk of a portfolio that consists of only that stock. Next, out of the remaining 49 stocks, randomly choose another stock. Combine that stock with the first stock chosen to form a two-stock portfolio. Calculate the standard deviation of that two-stock portfolio to quantify the portfolio’s risk. The risk of the two-stock portfolio should, on average, be lower than that of the one-stock portfolio. Next, randomly choose a third stock out of the remaining 48, add it to the two-stock portfolio to form a three-stock portfolio, and calculate its standard deviation. Proceed along these lines until the portfolio contains all 50 stocks.

We have calculated standard deviations in this manner using average monthly returns from 1995 to 1999

¹For our purposes here, we’ve ignored discounting.

FIGURE 1

Portfolio Diversification and Risk 1995-1999



(Figure 1). The figure shows that the risk associated with the portfolio declines fairly smoothly as the number of stocks in the portfolio increases. Notice, though, that once the portfolio has 10 to 15 stocks, adding more doesn't seem to decrease the risk of the portfolio much further. This remaining risk — the part not affected by holding more NYSE stocks — is called market risk.

Is there a way to lower portfolio risk even further? Only if we can lower the market risk. One way to lower market risk is to hold stocks not traded on the NYSE, in particular, stocks that trade on foreign markets, which represent claims on foreign assets.² Foreign economies generally do not move one-for-one with the U.S. economy. When the U.S. economy is in a recession, foreign economies might be

²Note, though, that many firms whose stocks are listed on the NYSE have extensive foreign operations. So, to some extent, a diversified portfolio of NYSE stocks already embodies some elements of international diversification.

in expansions and vice-versa. More generally, residents of other countries face different types of domestic risks — risks particular to those countries. Because every resident of a given country faces these same risks, they cannot be shared by trading with other domestic residents, that is, they cannot be diversified away by internal trade. So by trading financial assets with residents of other countries, households can share some of their domestic risk, thereby lowering the risk associated with fluctuations in consumption.

In a world where there is a lot of international risk-sharing, fluctuations in consumption across countries should be very similar, and investors' portfolios should include both domestic and foreign securities. Let's review some of the empirical evidence on these issues.

HOW MUCH RISK-SHARING IS THERE ACROSS COUNTRIES?

For international risk-sharing to occur, people must have the opportunity to trade in goods, services, and financial

capital across countries. If the costs of investing in foreign assets are too high — for example, if investors face barriers such as high transaction costs, tax and tariff payments, and certain types of capital controls — domestic investors will not find it profitable to do so. But when barriers to international capital flows are small, financial capital is mobile across countries, and international financial markets are said to be “open.” If international financial markets are open, we might expect that (1) investors would hold portfolios that are diversified internationally, (2) risk-sharing would allow fluctuations in consumption to be smoothed out relative to fluctuations in income, and (3) domestic saving would not be too closely tied to domestic investment, since residents of countries would be free to borrow from and lend to each other. Indeed, simple economic models of risk-sharing suggest that international economic data should confirm these predictions. However, as we will see, the data show there is less-than-perfect international risk-sharing. One reason, at least in the developed countries, might be that the benefits of undertaking further measures to reduce risk may not outweigh the costs.

International Financial Flows. It is difficult to get good, direct data on how freely financial capital flows across borders. Instead, we must look at indirect evidence on cross-border flows. If financial markets are open, dollar-denominated returns on nearly identical assets in different countries should be nearly the same. For example, the interest rate on large-dollar certificates of deposit sold in New York and

³Eurodollars are U.S. dollars deposited in foreign banks outside the United States or in foreign branches of U.S. banks. There is no exchange-rate risk in making these interest-rate comparisons across countries because the assets are denominated in a common currency.

the rate on London eurodollar deposits of the same maturity should be nearly identical.³ Many empirical studies have examined these and other onshore-offshore interest rate differentials. Generally, the studies have found very close links between onshore and offshore money markets for financially developed countries such as the U.S., the U.K., France, Italy, Germany, and Japan.

Figure 2 plots the difference between the interest rate on a eurodollar deposit with a three-month maturity and the interest rate on a three-month U.S. certificate of deposit. Note how the interest rate differential between the two series has declined over time and is now, on average, very close to zero. This demonstrates that financial capital now flows quite freely between the U.S. and the London financial market.

Similarly small differentials are found when economists analyze many of the world's developed financial markets.⁴ However, financial markets in many other countries, especially less developed ones, are not as open. Thus, it is costly for their residents to share economic risks with investors in other countries.


Cross-Border Portfolio

Diversification. In financially developed countries, where capital is mobile, one might assume that international risk-sharing would be present and investors' portfolios would be diversified internationally. However, economic research indicates that investor portfolios are not highly diversified internationally, especially those of investors in the United States and Japan. A 1991 study

⁴Partial surveys of these studies can be found in the article by Jeffrey Frankel and the 1986 article by Maurice Obstfeld. Obstfeld's 1994 study discusses more recent evidence on the financial openness of France, Italy, Germany, Japan, the United Kingdom, the U.S., Spain, Portugal, Ireland, and Greece.

FIGURE 2

Onshore-Offshore Interest Differential



by Kenneth French and James Poterba found that, at the end of 1989, the share of foreign equities in total equity holdings was 4 percent for residents of the United States, 2 percent for residents of Japan, and 18 percent for residents of the United Kingdom. A 1994 study by Linda Tesar and Ingrid Werner estimated that more than 96 percent of U.S. wealth was invested in U.S. equity in 1991. They also found that the fraction of the total U.S. stock market held by Germany, Canada, Japan, and the U.K. was below 12 percent in 1991, suggesting that residents of those countries were not as internationally diversified as we might expect. Marianne Baxter and Urban Jermann confirmed that, in 1991, over 95 percent of equities held by U.S. investors were those of U.S. corporations.

In a 1998 paper, Linda Tesar and Ingrid Werner provided more recent evidence on the lack of portfolio diversification internationally (Figure 3). The fact that equity holdings are disproportionately invested in domestic equities is called the *home equity bias*. The figure shows that home equity bias is smaller for the United Kingdom and


Germany than it is for Canada and the U.S. Home bias is greatest for Japan. Note that in each of the countries, the home equity bias has been getting smaller over time (that is, the percent of wealth invested in foreign assets is increasing). Nevertheless, even though this bias appears to have lessened over time, there is still not as much diversification as simple financial models suggest there should be.⁵

However, there are some problems with these measures of home equity bias. In particular, many large domestic firms have overseas operations, and some firms cross-list their securities on more than one market. Thus, to some extent, holding a well-diversified portfolio of U.S. stocks gives investors

⁵Simple models of portfolio choice suggest that optimal investment strategies involve holding a fraction of wealth in a risk-free asset and the remainder in the global market portfolio. On average, from 1980 to 1988, U.S. equities represented about 45 percent of total global market capitalization (Tesar 1995). Thus, the simple model predicts that U.S. residents should have a substantial fraction of their equity portfolios allocated to foreign stocks.

FIGURE 3

Home Equity Bias 1987-1996*



*From Linda Tesar and Ingrid Werner, "The Internationalization of Securities Markets since the 1987 Crash," *Brookings-Wharton Papers on Financial Services*, Washington, DC, Brookings Institution, 1998, reprinted with permission.

exposure to international developments. Nevertheless, the consensus among researchers studying this issue is that residents of many countries show a home bias in their portfolio holdings.

Fluctuations in Consumption Across Countries. Let's focus on financially open countries where there are few barriers to international capital flows. Suppose further that asset markets are *complete*, which means that households can purchase securities that insure them against all possible risks to their consumption of goods and services.⁶ If markets are complete and financial markets are open, people could help insure themselves against the

economic risks that lead to fluctuating consumption by purchasing assets from foreigners.

If residents of countries could insure themselves against the economic risks they face, the data on domestic consumption and output would show that consumption growth rates are more highly correlated across countries than are output growth rates. Why? Because people could use financial assets to, in effect, pool their incomes and then divide up the shared proceeds, much as in the simple example cited earlier. In effect, countries that might otherwise have low consumption and output in a year borrow from those that have high

consumption and output in a year.⁷

How do the data stack up against this hypothesis? (See Table.) To simplify things, the table shows comparisons between consumption and output in individual countries and world consumption and output, which is used as a common benchmark, for the period 1973 to 1992. The table reports the *correlation* of individual-country consumption with world consumption and country output with world output.⁸ If there is a lot of international risk-sharing and financial markets are nearly complete, the correlation between an individual country's consumption and world consumption should be much higher than the correlation between that country's output growth and world output growth.⁹

The table shows that consumption growth rates are not very highly correlated and that output growth rates are more highly correlated than consumption growth rates — not at all what the simple model of international risk-sharing predicts. Thus, despite the fact that, at least among developed countries, capital markets are well integrated and the barriers to trade are generally low, there does not seem to be as much international risk-sharing as we might expect.

⁸Correlation is a statistical measure of co-movement. The closer the correlation is to one, the more closely two series move together. When the correlation is positive, the series move together over time. That is, when one series is high, the other series tends to be high, and when one series is low, the other tends to be low. When the correlation is negative, the series move opposite to each other. When the correlation is zero, the series do not move together at all.

⁹In a world with complete risk-sharing, certain economic models predict the world supply of consumed goods should be allocated across countries approximately in proportion to their share of total world wealth. Therefore, consumption growth in each country should be identical to the growth rate of world consumption. See the 1995 article by Linda Tesar for details.

⁶In reality, of course, asset markets are not complete. For example, you cannot buy an insurance policy that pays off when you become unemployed. But we put that aside for the moment and concentrate on the idealized case of complete asset markets.

⁷For the industrialized countries, domestic consumption fluctuates less than domestic output, so some risk-sharing is taking place, at least *within* countries. In this section we are concerned with cross-country comparisons, to get some information on the extent of risk-sharing *across* countries.

TABLE

International Consumption and Output Correlations 1973-1992

Country	Consumption Correlation ^a	Output Correlation ^a
Canada	0.56	0.70
France	0.45	0.60
Germany	0.63	0.70
Italy	0.27	0.51
Japan	0.38	0.46
United Kingdom	0.63	0.62
United States	0.52	0.68
OECD average ^{b,c}	0.43	0.52
Developing country average ^b	-0.10	0.05

^aCorrelation between the annual change in the log of a country's real per capita consumption (output) and the annual change in the log of the rest of the world's per capita consumption (output) over 1973-92. The world is defined as the 35 benchmark countries in the Penn World Table Mark 5.6.

^bAverage correlations are population-weighted averages of individual country correlations.

^cExcludes Mexico.

(Source: Maurice Obstfeld and Kenneth Rogoff. *Foundations of International Macroeconomics*, Cambridge, MA: MIT Press, 1996, Table 5.1, pg. 291; reprinted with permission)

Domestic Saving and Investment. Another way that households can smooth their consumption over time is by drawing down savings when income is low and increasing savings when income is high. For a country as a whole, domestic saving and investment provide a means whereby residents can insure themselves against some of the economic risks they face.

¹⁰Dissaving occurs when consumption exceeds income.

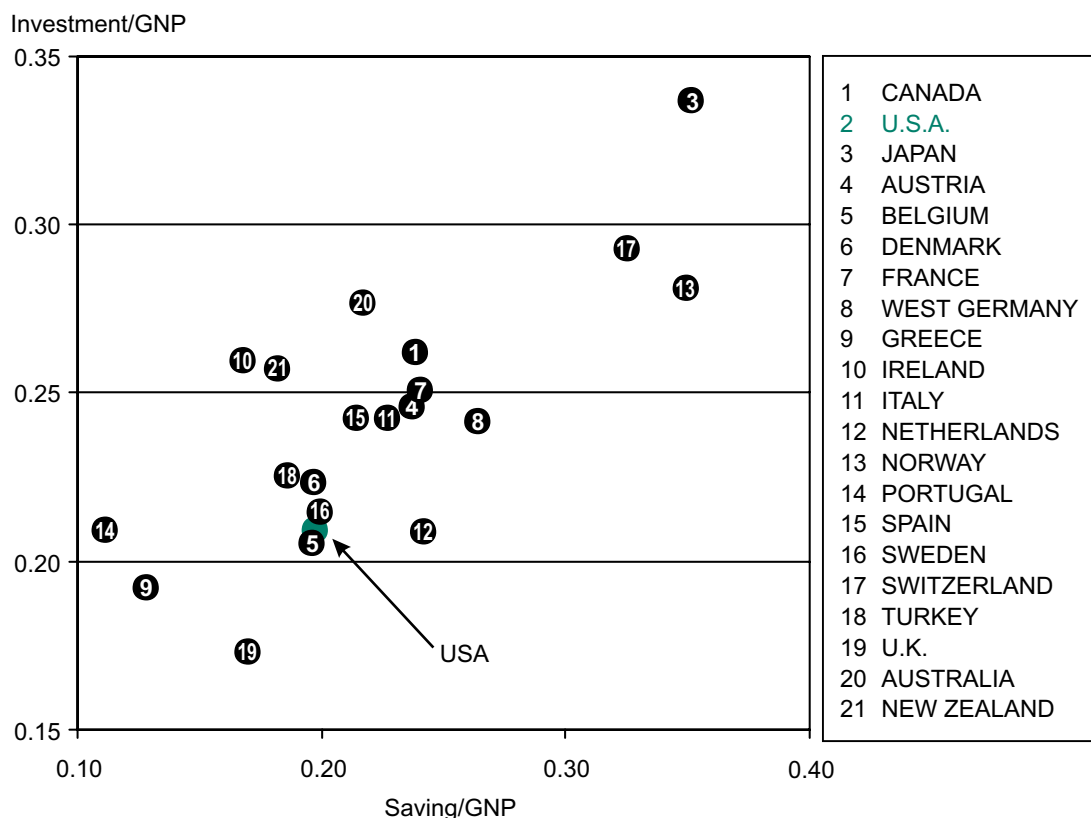
If a country could not engage in trade with any other country, domestic saving would have to equal domestic investment. Under such an economy, the only avenue available to residents for smoothing consumption would be saving and dissaving.¹⁰ But access to international financial markets breaks the link between domestic saving and domestic investment, since savings can be imported to finance domestic investment or exported to find the highest return. Thus, domestic saving and investment need not move together over time if financial markets are open and capital is mobile.

In a very influential 1980 article, Martin Feldstein and Charles Horioka provided evidence that showed saving and investment were highly correlated across a wide sample of countries — not what we would expect if capital is internationally mobile (Figure 4). The figure, which provides some updated evidence on the Feldstein-Horioka puzzle, plots the 10-year average of savings and investment for a sample of 22 countries from 1981 through 1990. Each point represents the average saving rate and average investment rate for an individual country. If saving and investment were unrelated across countries, the points would be evenly scattered about the diagram, with no discernable pattern visible. Instead, the figure shows a clear positive relationship between saving and investment for these countries: When saving is high, investment is high; when saving is low, investment is low. This positive relationship between saving and investment is not what we would have expected in a world of open financial markets and mobile capital.

There is a vast literature in economics on the Feldstein-Horioka puzzle. Many studies have verified that saving and investment tend to be positively correlated both over time within a single country and at a point in time across countries. There is some evidence that the saving-investment correlation may be getting weaker over time as financial markets become more and more integrated. But the data also show that investment-savings correlations tend to be lower in developing countries than in industrialized countries, which is contrary to what we would expect if this correlation were a strong measure of capital's mobility internationally. Developing countries tend to have a low degree of international capital mobility, which suggests that domestic saving and investment should be highly correlated.

FIGURE 4

Saving and Investment Rates for Selected Industrialized Countries 1981-1990



Source: Author's calculations

ARE THERE UNEXPLOITED GAINS FROM INTERNATIONAL RISK-SHARING?

Economic theory suggests that there may be substantial gains from international risk-sharing. But the empirical evidence we reviewed above suggests that there isn't as much international risk-sharing going on as we might expect. Does that mean there are unexploited gains from international risk-sharing? In the case of the developed countries, many economists think not. For these countries, the gains from further international risk-sharing may be very small. But for developing countries,

which often do not have open financial markets, the gains from international risk-sharing may be substantial.

So why do the data *not* bear out the predictions of *simple* models of international risk-sharing? Because the world is complicated. In particular, our simple models do not fully account for three important factors.

For one thing, financial markets are not really complete.¹¹ You cannot buy insurance against all future events, such as becoming unemployed. The less complete financial markets are, the less correlation of consumption there will be across countries.

Another problem is that not all goods are traded. For example, there is

¹¹The fact that consumption is not fully insured against country-specific shocks may be a reflection of incomplete asset markets rather than an inability or unwillingness to trade existing assets internationally. To examine this issue, economists have studied the extent of risk-sharing within countries — such as across states within the U.S. and across prefectures in Japan. Generally, the studies find that there is not complete risk-sharing within countries but that the extent of risk-sharing at the national level exceeds that at the international level. See the studies by Atkeson and Bayoumi (1992), Crucini (1999), and Obstfeld (1994).

no international trade in haircuts, fitness club memberships, Big Macs, or private and government services. Nontraded goods matter because you cannot smooth out the consumption of those goods by trading them with residents of other countries. Over time, consumption of traded goods will move in a similar fashion across countries, but consumption of nontraded goods will show dissimilar movements. The bigger the role nontraded goods play in total domestic consumption, the less cross-country correlation in total consumption we should see.

Third, transaction costs may prevent portfolios from being as internationally diversified as simple models of risk-sharing predict. Developed countries afford their residents ample opportunity to diversify risk by trading domestic assets. If transaction costs must be incurred in trading goods and assets with foreign countries, it may not be beneficial to domestic residents to try to exploit further gains from risk-sharing, even if the transaction costs are small. A number of economic studies have concluded that there is little in the way of unexploited gains to further risk-sharing among the developed countries.¹²

Linda Tesar's conclusions in her 1995 study are typical. She found that, for developed countries, the gains from international risk-sharing are usually less than one-half of 1 percent of the lifetime consumption of a typical household. Since these gains are so small, even small transaction costs for

international trade are enough to offset the gains. These conclusions are similar to those reached by Maurice Obstfeld and Kenneth Rogoff. They found that small transaction costs can help explain the perceived lack of international risk-sharing.¹³

In addition to these three factors, there are other channels through which households can share risk without using international financial markets: international trade in goods and services and domestic saving and investment.

time. By cutting back on consumption today and increasing saving and investment, more consumption can take place in the future. Similarly, by cutting back on saving and investment today, more consumption can be had today at the expense of future consumption. In this way, residents of a country can shift consumption from the present to the future and vice versa. The ability to adjust saving and investment to smooth out fluctuations in consumption lessens the need to use international financial markets to do the same.

The bigger the role nontraded goods play in total domestic consumption, the less cross-country correlation in total consumption we should see.

Trade in goods and services across countries can, to some extent, offset the need to diversify portfolios and share risk internationally. International trade in goods allows households to import goods when the benefits from consuming them are high, substituting in part for sharing risk by trading in international financial assets. This substitution can happen because a country's *terms of trade* may change in such a way as to offset bad output shocks.¹⁴ If a country's terms of trade are high when domestic output is low, it is relatively inexpensive to import goods and thereby smooth consumption.

As we have seen, domestic investment offers some scope to smooth out fluctuations in consumption as well. Investment allows residents of a country to reallocate their consumption over

Would countries experience large gains if there was more international risk-sharing? The answer to this question depends on the economic model being used. In a typical economic model, the gains to an average consumer of fully eliminating consumption risk are usually very small. For example, an influential study by Robert Lucas suggested that a typical consumer would be willing to pay only about \$80 a year to totally eliminate variability in consumption. Lucas's results are based on some special assumptions that have been challenged, but many of the models that economists have used to analyze gains from international risk-sharing have a structure very similar to the one Lucas used.

Thus, in these models, consumers aren't willing to pay much to eliminate variability in consumption. Models in which consumers care more about lowering consumption fluctuations would demonstrate bigger gains from risk-sharing.

There are conditions, however,

¹²See the papers by Harold Cole and Maurice Obstfeld; Linda Tesar (1995); Enrique Mendoza; David Backus, Patrick Kehoe, and Finn Kydland (1992), and Obstfeld and Rogoff.

¹³In addition, individuals may not like the risks that international stocks entail, since they are likely to be less informed about such stocks.

¹⁴A country's terms of trade can be defined as the price of exports relative to the price of imports.

under which the gains from further international risk-sharing for large developed countries may be substantial. Suppose that domestic investment rises in response to expanded opportunities for diversification, so that domestic purchases of capital equipment rise as investors diversify portfolios internationally. Increased investment in capital goods leads to a greater increase in the capital stock than would otherwise be the case and may lead to faster long-run economic growth. Over time, the economy could experience a large increase in its standard of living because of the diversification opportunities presented by foreign markets.

In addition, although the aggregate gains to the economy of increased risk-sharing may be small, the benefits perceived by individuals within a society may be large. For example, avoiding taxes may motivate some individuals to trade in international financial markets in an effort to shelter income. Or the income variability faced by a typical household may be much larger than the variability of income in the economy as a whole. So some income groups in the economy may gain substantially from international diversification and risk-sharing.

For developing countries, the gains to further risk-sharing may be quite a bit larger than any gains developed countries might realize from increased international risk-sharing. First, developing countries contribute less to world output than developed countries, making it less likely that their domestic output would rise and fall with world output. Thus, more of their country-specific risk can be eliminated by trading assets with residents in the rest of the world. Second, developing countries' output is often much more variable than that of developed countries, which means the potential benefit from risk-sharing is greater, since there is more scope to reduce output volatility.



In sum, the precise magnitude of the gains from further international risk-sharing remains an open question. But it seems that developing countries are the most likely to be its strongest beneficiaries.

SUMMARY

Residents of a country might be better off if they could share some economic risks with residents of other

countries, who may face different economic risks. Sharing these risks allows residents of both countries to potentially smooth out the fluctuations in consumption they might otherwise face. However, the empirical evidence on the correlations between international consumption, the link between saving and investment, and portfolio diversification suggests that the extent of international risk-sharing is not as

great as we might at first suspect. Perhaps, the residents of developed countries have already shared about as much risk as it is worthwhile for them to do so. The costs of undertaking further measures to reduce risk may not be worthwhile.

For developing countries the benefits of international risk-sharing are likely to be substantial. Output and consumption in these countries tend to be very volatile from year to year. If

these countries can gain increased access to world financial markets, they would be able to substantially smooth fluctuations in their aggregate consumption. 

REFERENCES

Atkeson, Andrew, and T. Bayoumi. "Do Private Capital Markets Insure Regional Risk? Evidence from the United States and Europe," University of Chicago Working Paper, 1992.

Backus, David, Patrick Kehoe, and Finn Kydland. "International Real Business Cycles," *Journal of Political Economy* 100 (1992), pp. 745-75.

Baxter, Marianne, and Urban Jermann. "The International Diversification Puzzle Is Worse Than You Think," *American Economic Review*, 87 (1997), pp. 170-91.

Cole, Harold, and Maurice Obstfeld. "Commodity Trade and International Risk Sharing," *Journal of Monetary Economics*, 28 (1991), pp. 3-24.

Crucini, Mario. "On International and National Dimensions of Risk Sharing," *Review of Economics and Statistics*, 81 (1999), pp. 73-84.

Fama, Eugene. *Foundations of Finance*. New York: Basic Books, 1976.

Feldstein, Martin, and Charles Horioka. "Domestic Savings and International Capital Flows," *Economic Journal*, 90 (1980), pp. 314-29.

Frankel, Jeffrey. "Quantifying International Capital Mobility in the 1980s," in *On Exchange Rates*. Cambridge, MA: MIT Press, 1993.

French, Kenneth, and James Poterba. "Investor Diversification and International Equity Markets," *American Economic Review*, 81 (1991), pp. 222-26.

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

Mendoza, Enrique. "The Terms of Trade, the Real Exchange Rate, and Economic Fluctuations," *International Economic Review*, 36 (1995), pp. 101-37.

Obstfeld, Maurice. "Capital Mobility in the World Economy: Theory and Measurement," *Carnegie-Rochester Conference Series on Public Policy*, 24 (1986), pp. 55-103.

Obstfeld, Maurice. "International Capital Mobility in the 1990s," in Peter B. Kenen, ed., *Understanding Interdependence: The Macroeconomics of the Open Economy*. Princeton, NJ: Princeton University Press, 1994.

Obstfeld, Maurice, and Kenneth Rogoff. "The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?" NBER Working Paper 7777, 2000.

Tesar, Linda. "Evaluating the Gains from International Risksharing," *Carnegie-Rochester Conference Series on Public Policy*, 42 (1995), pp. 95-143.

Tesar, Linda, and Ingrid Werner. "International Equity Transactions and U.S. Portfolio Choice," NBER Working Paper 4611, 1994.

Tesar, Linda, and Ingrid Werner. "The Internationalization of Securities Markets Since the 1987 Crash," *Brookings-Wharton Papers on Financial Services*, Washington: Brookings Institution, 1998.

How Responsive Is the Demand for Residential Land to Changes in Its Price?

BY RICHARD VOITH

Public policies from zoning to income-tax deductions for mortgage interest affect the price of residential land. In this article, Richard Voith's estimates help measure the effect of public policies on land consumption in the United States.

The dominant trend in metropolitan development in the 20th century was the increasing use of land per capita. As households moved to the suburbs, both houses and residential lots increased dramatically in size. These changes were direct consequences of higher incomes and lower transportation costs: people could afford to spend more on housing, and traveling longer distances between home and work became more feasible.¹ As commuting distances lengthened, the supply of land deemed to be acceptable for residential development increased greatly, and this greater supply meant that residential land became more affordable.

Not surprisingly, since rising incomes and lower transportation costs

are common throughout developed countries, the trend toward increasing land consumption per capita is an international phenomenon. This trend, however, has been more pronounced in the United States than in other developed countries. Many observers suggest that rising incomes and lower transportation costs had a strong impact on the rate of suburbanization and increasing residential land consumption in the United States because Americans have strong preferences for the open space associated with low-density metropolitan development.² Another

¹In the traditional monocentric model of urban economies where everyone works in the center of the region, increasing income does not necessarily result in larger residential lots because the increased desirability of a larger lot is also associated with a longer commute. As income rises, the cost of commuting in terms of time increases. The two forces associated with rising income — the desire to buy more land and the increased cost of commuting — tend to offset one another. With the rise of suburban employment, however, households have the opportunity to increase the size of their lot without necessarily increasing their commuting costs.

²See the article by Peter Mieszkowski and Edwin Mills and Witold Rybczynski's book.

factor contributing to the pattern of low-density development may be the low price of land in the United States.³ Abundant vacant land means that the supply of land that can be used for residential development can be greatly increased through investment in transportation. Such investment helps keep the price of residential land low.⁴ In turn, these low prices encourage households to buy larger lots. However, public policies — including taxation, transportation, and zoning regulations — have also affected after-tax residential land prices.⁵ Policy choices, therefore, may have played an important role in the patterns of U.S. metropolitan development as well as in the rapid increase in per capita land use.

While the pace of decentralization has continued unabated in the United States, concerns about road

³U.S. and Australian metropolitan areas tend to be much less dense than those in Canada, Europe, and Asia. While international comparisons of residential land prices are not readily available, prices per square foot of office space, which should reflect land value as well, clearly indicate that U.S. metropolitan prices are relatively low. Only two U.S. cities, San Francisco and New York, are in the top 20 highest international office rents. (CB Richard Ellis Global Research and Consulting)

⁴Transportation investments may increase the value of land in areas that benefit from the investment, but these investments serve to increase the overall supply of land suitable for houses and, therefore, help keep the average price of residential land low.

⁵It's likely that public policies in Europe and Asia have affected land prices as well. The prevalence of "greenbelts" and other land-use restrictions reduce the supply of land available for development, which tends to raise the price of land.



Dick Voith is a senior vice president and principal at Econsult Corporation, Philadelphia. When he wrote this article, he was

an economic advisor in the Research Department of the Philadelphia Fed.



congestion, loss of farmland and green space, and the character of community in our decentralized metropolitan areas have increased. These concerns have led some people to ask whether current patterns of low-density growth need to be reexamined and to suggest a whole range of policies that would alter current growth patterns.⁶ To predict whether public policies are likely to have an impact on the amount of residential land households use and to evaluate the costs and benefits of policies that may affect residential land use, we must have a thorough understanding of the nature of demand for residential land.

One key aspect of this demand is how responsive it is to changes in price.⁷ In other words, if the price of

land increases, will consumers adjust their demand for land downward? And, if so, by how much?⁸ If households have strong preferences for residential land and therefore consumers change their land consumption very little in response to a large increase in price, policies that raise the price of land would have very little impact on patterns of land use.

Furthermore, attempts to change land-use patterns would be very costly from a social point of view: Any policy that managed to reduce land consumption would generate hardships for households. Households would find it difficult to derive as much benefit from spending their money on other

more land do households want when their income rises? Glaeser, Kahn, and Rappaport provide estimates that suggest that the demand for residential land is fairly unresponsive to changes in income. They estimate that a 1 percent rise in income results in an increase in expenditures on land of only 0.4 percent.

⁸Economists call this relationship “the elasticity of demand.” This elasticity measures how a change in price affects the quantity of a good demanded. Specifically, it is the percentage change in the quantity of a good demanded resulting from a 1 percent change in its price. The price elasticity for a product is typically negative, that is, a rise in price results in a decline in quantity demanded.

goods, such as more exotic vacations or more expensive clothing, as they did from their large yard.

On the other hand, if consumers readily adjust the quantity of land they consume in response to changes in price, policies that modestly change the price of land could have a large impact on land-use patterns. If consumers significantly adjust their land consumption to changes in its price, it means that there are other goods almost equally as attractive. Therefore, when land prices rise, households simply choose to have a smaller yard and have more money available for other uses. In this case, public policies that affect the price of land may have a large impact on land-use patterns, and these changes may have a relatively small impact on households’ satisfaction. Thus, a key piece of information needed for understanding the forces affecting metropolitan development is how responsive households are to changes in the price of land.

TWO VIEWS OF THE DEMAND FOR RESIDENTIAL LAND

In the 1960s, two researchers in urban economics, Richard Muth and William Alonso, offered different perspectives on the nature of the demand for residential land. Muth viewed residential land as an input to the production of a house. In his view, the demand for residential land was based on the price of land relative to the price of other materials and labor needed to create residential housing. If land were inexpensive, builders would use more land and less lumber, steel, construction labor, and so forth when constructing a house. That is, builders would tend to favor single-floor structures covering more land. If land were expensive, builders would construct taller houses so that more houses could be put on less land.

Muth’s fundamental insight was to apply a well-developed

⁶These policies range from impact fees on new development, to land conservatories, to urban growth boundaries — like the one in Portland, Oregon — which circumscribe the areas in which development is allowed to occur. In the Third District, New Jersey’s land-use plan limits development in the Pine Barrens. The state of Pennsylvania has recently enacted new legislation designed to encourage more cooperative local planning, increase regional land-use planning, and conserve open space.

⁷Another key aspect of the demand for residential land is how responsive it is to changes in income. In other words, how much

microeconomic theory that allows researchers to compute the demand for any input used in the production of any good if they know how responsive demand for the final good is to changes in price, how easy it is to find a substitute for the input, and how large a share of the total of all materials the input is. Thus, for housing, if we knew how responsive housing consumption is to changes in housing prices, how easy it is to substitute capital for land (think of building up, rather than out), and how big a share of total input costs land is, on average, we could compute how responsive changes in demand for land would be to changes in its price.

Using this microeconomic theory and armed with estimates of the variables outlined above, Muth (1964,1971) concluded that the demand for residential land was not very sensitive to changes in its price. Muth estimated that a 1 percent increase in the price of residential land would reduce the amount of land used by 0.75 percent, or a price elasticity of -0.75 . His approach allowed him to estimate the price elasticity of demand for residential land without addressing the issue of consumers' direct demand for land.

Unlike Muth, Alonso focused on the fact that households probably valued residential land for other reasons, not just simply as an input to the production of housing. In his view, consumers' demand for residential land was like that for any other durable good.

Households may want land not only because they need a place on which to build a house but also because they want to plant a garden, create a play area for children, or ensure privacy. The amount of land that consumers want, therefore, will depend not only on how much land costs relative to other materials needed to build a house but also on the consumer's income and tastes; the attributes of the land itself, such as its location in the metropolitan

area; and the price of land relative to the price of other consumer goods. The elasticity of demand for land, in Alonso's view, may be very different from that derived when viewing land as an input to the production of housing. Demand for land that is going to be used for a garden may be more sensitive to changes in price, for example, than

technique, which is called hedonic analysis, is a key tool economic researchers use when analyzing housing markets.

The second problem arises from the fact that the price of residential land is not independent of the kind of house on the lot. Economic theorists have shown that the price of a component of a bundled good, like land in

If consumers readily adjust the quantity of land they consume in response to changes in price, policies that modestly change the price of land could have a large impact on land-use patterns.

demand for land on which a house will be built.⁹ Unfortunately, little work has been done that directly estimates consumers' demand for residential land.¹⁰

CHALLENGES IN ESTIMATING THE DEMAND FOR RESIDENTIAL LAND

Estimating the demand for residential land is a very difficult econometric problem for two reasons. First, we don't usually have direct prices for residential land because most residential land that is sold is bundled with a house: we only see the price of the house and the land together. Thus, researchers must use statistical techniques to break down the sale price into prices for the unit's individual components: bathrooms, bedrooms, and other housing traits, including land. This

housing, will depend on the quantity of the other components in the bundled good. Because the price of residential land depends on the bundle of housing traits, it differs across houses.¹¹ Since each housing bundle is associated with a quantity of land and an implicit price of land, consumers have a range of choices for both price and quantity of land when buying a house. This choice results in what economists call a selection problem: People who have strong preferences for land tend to buy houses with more land and, on the margin, are willing to pay a higher price for land bundled with the house. On the other hand, people without a strong desire for land will buy houses on small lots with relatively inexpensive land.

To make this issue more concrete, consider two houses in the same neighborhood: one with a quarter acre of land and one with two acres. Suppose further there are two similar consumers, but one has a strong preference for large lots. In this situation

⁹While we do not attempt to make separate elasticity estimates for land used for housing and for land used for other purposes, our estimated prices do allow for the fact that the market price per acre of land may depend on the parcel size. That is, larger parcels for a single house might have lower prices per acre.

¹⁰Paul Cheshire and Stephen Sheppard, who studied a variety of British cities, report price elasticities of demand for land ranging from -0.6 to -1.6 .

¹¹Land prices will also differ depending on the land's location within the metropolitan area and on attributes of the land such as whether it has a good view or is near the seashore. In my statistical work, one of the factors determining a parcel's land price is its distance from the center of the metropolitan area.

the consumer with the strong preference for more land would be willing to pay more for the house with the large lot than will the other consumer. Thus, he would both choose to consume more land and be willing to pay a higher price for the house to bid it away from the other consumer.¹²

What we observe in market transactions is how different households choose among housing and land bundles. We cannot tell directly how much of the observed differences in lot size across households is a result of an individual household's adjusting how much land it uses in response to price differences and how much is a consequence of observing a different household with different tastes. Ignoring the selection issue results in biased estimates of quantity's responsiveness to price. To correctly estimate demand, we would like to observe how the same household reacts to a change in price.

Researchers Timothy Bartik and Dennis Epple independently suggested an approach to dealing with the selection problem inherent in estimating demand for residential land. Their method — which applies to estimating demand for components of any bundled good, not just land — requires data that satisfy a number of criteria that are difficult to satisfy; therefore, their method has seldom been used. We implemented their procedure to evaluate how consumers adjust their residential land consumption in response

to changes in land prices. (The box on page 39 describes the Bartik-Epple procedure and my application of it to the estimation of the demand for residential land.)

NEW ESTIMATES OF THE ELASTICITY OF DEMAND FOR RESIDENTIAL LAND

In a recent working paper, Joseph Gyourko and I applied the Bartik-Epple procedures to develop new estimates of the price elasticity of demand for residential land. We used a massive data set on housing sales in Montgomery County, Pennsylvania, over

Our model indicates that the quantity of residential land that households choose is highly sensitive to the price of land.

26 years. Our data on almost 100,000 sales of single-family detached houses from 1972 through 1997 included not only the sale price, date, and detailed information on the characteristics of each house sold but also the amount of land in the parcel.¹³ In addition to information on the house and its lot, all parcels were geocoded so that we had detailed information about the parcel's location within the county, its proximity to employment centers and transporta-

tion, and the characteristics of the neighborhood.

Using these data, we constructed statistical models of housing prices, and these models predicted the value of the property based on characteristics of the property and its neighborhood. The models also yielded estimates of the contribution of each housing trait — bathrooms, central air conditioning, square footage of the lot, and so forth — to the value of the property. From these estimates, we can derive the implied price of residential land for each parcel sold in each year of the sample. (See *House Prices and Residential Land Prices in Montgomery County, 1972-1997*.) This is the first stage of the Bartik-Epple procedure.

The second stage uses the estimated land prices for each parcel and the observed quantities of land associated with each parcel to determine the relationship between prices of land and quantity of land consumed. The second-stage statistical model, the details of which are discussed in *The Bartik-Epple Approach to Estimating the Demand for Bundled Goods*, on page 39, provides an estimate of the relationship between the price and quantity of land that is free of potential biases associated with selection problems discussed earlier.

Our model indicates that the quantity of residential land that households choose is highly sensitive to the price of land. The elasticity of demand is around -1.6, which indicates that a 10 percent increase in price would reduce land consumption 16 percent. Estimates of the elasticity of demand that do not control for the selection problems identified by Bartik and Epple show significantly greater responsiveness of residential land consumption to land prices. Even though our estimates are about 50 percent lower than estimates that do not take selection issues into account, they are still substantially higher than those suggested by Muth, who found that a 10 percent increase in

¹²The discussion of selection ignores the issue of location. Basic urban theory suggests that houses located near the center will have high land prices because commuting costs are low, and because of these high prices, residential lots will tend to be small. Although the selection mechanism described above still occurs, prices of small, centrally located residential lots are likely to be high because of premiums for central locations, but not as high as they might have been if people with strong preferences for land had chosen to live there.

¹³Housing traits include the unit's age, square footage, square footage of the lot, and the presence of central air conditioning, fireplace, pool, and garage. Neighborhood characteristics include the population density of the unit's census tract, percent of the tract with single-family housing units, travel time from the census tract to the Philadelphia central business district, and the presence of commuter rail service in the neighborhood.

House Prices and Residential Land Prices in Montgomery County, PA 1972-1997


T

he pattern of real prices for houses in Montgomery County from 1972 to 1997 differs markedly from the pattern of residential land prices in our statistical models (Figures 1 and 2). Real housing

prices slowly trended downward from 1972 to 1982, rose sharply from 1982 to 1989, then trended downward again in the 1990s. Overall, mean real house prices stood at \$118,500 at the start of the sample and rose to \$155,100 in 1997, an increase of about 31 percent, or 1.2 percent per year.

Our statistical models break down housing prices into prices for the houses' component traits for each year of the sample, an approach that allows us to compute estimates of the price of land for every house sold during the sample period. By averaging the estimated lot prices for all houses sold in each year, we can show the pattern of land prices for houses sold over the sample period (Figure 2). In sharp contrast to the average price of housing, the average price of land fluctuates considerably from year to year but shows no significant trend over time. Average price per square foot of land stood at \$1.03 at the beginning of the sample period and \$1.09 at the end of the period in 1997. Although the difference in prices from the beginning to the end of the sample was less than 6 percent, there were large fluctuations. Low point for the price of land was \$0.72 in 1979, less than half its peak level of \$1.51 in 1988.


FIGURE 1
House Sale Prices
in Montgomery County, PA



The year-to-year variance in the estimated average value of residential land over the sample period is not surprising, since it reflects changes in overall supply and demand in the regional economy.^a The wide variance in land prices is not unexpected because the value of land reflects the value of the location.^b Casual observation of the land-price time series shows that land prices fall substantially at the beginning of recessions and their associated depressed housing markets and rise markedly just after the economic upturn begins. Land prices peaked in Montgomery County in the late 1980s. Since then, prices have trended downward roughly 30 percent in real terms, although the last year of data show a marked upturn in price that the popular press suggests has continued into the new millennium.

Based on the average lot size and the price per square foot shown in Figure 2, the cost of land hovered around 15 percent of mean house value in most years between 1972 and 1997. For example, in 1972, the mean lot size was 19,856 square feet, the price per square foot was \$1.03, and the average value of a house \$120,300. Therefore, land constituted 17 percent of mean house value that year. This percentage reached its low in 1992, when the price of \$0.89 per square foot implied that land was only 10 percent of mean house value.

FIGURE 2
Implied Lot Prices per Square Foot



^aIn addition, the average value of residential land is affected by the location of residential sales in a given year. If, for example, sales were concentrated in areas with low land prices, this would be reflected in a lower countywide average price for that year.

^bThis stands in contrast to the value of housing structures, which is tied to construction costs in the long run. Because the value of land is not tied to construction costs, the price of land fluctuates more to equilibrate supply and demand.

land prices would reduce land consumption less than 10 percent. Our findings provide evidence that Alonso was correct in arguing that the demand for land is based on more than its use as an input to the construction of houses.¹⁴

IMPLICATIONS FOR POLICY AND METROPOLITAN GROWTH PATTERNS

In recent years, rapid rates of low-density suburban growth have convinced most observers, as well as developers and consumers, that single-family detached houses on large lots primarily reflect an American preference for open space and personal privacy. But these same patterns of low-density development have also been associated with a rapid increase in automobile travel and congestion, as well as concerns about the loss of open space, quality of development, loss of community, and decline in older cities and towns.

This juxtaposition of preference and concerns has resulted in a quandary: If Americans strongly prefer very low-density development, addressing the concerns raised is likely to be very costly. Efforts to force people into denser communities through public policies that raise land prices or through land-use regulation would exact a high price in terms of households' welfare and, therefore, would likely be politically unpopular.

On the other hand, higher levels of transportation investment to address the increased demand for automobile travel associated with less dense living patterns could relieve congestion but, in addition to the expense, would likely put older commu-

nities at an even greater disadvantage, perhaps even accelerating their decline.¹⁵

changes in price should help policymakers assess the likely impact of "smart growth" policies. Smart growth policies

Our estimates suggest that American consumers are very flexible with respect to residential land consumption.

The finding that the quantity of residential land that households choose to own is very responsive to the price of land suggests that the quandary described above may not, in fact, be such a thorny issue: Americans' choice of low-density residential development is as much a reflection of the relatively low price of land as it is of a uniquely strong preference for large residential lots. Our estimates suggest that American consumers are very flexible with respect to residential land consumption: When land prices rise, American consumers readily shift their consumption to other goods with relatively lower prices.

Because households are sensitive to the relative price of land, public policies that affect the price of residential land are likely to have a considerable impact on the density of metropolitan areas. For example, some estimates imply that the federal tax treatment of owner-occupied housing lowers the after-tax cost of housing by 12 percent.¹⁶ On the basis of these estimates, the tax treatment of owner-occupied housing lowers residential density by 16.1 percent.¹⁷

Our estimates of how responsive households' land consumption is to

have focused on ameliorating some of the perceived negative consequences — increased reliance on cars, greater travel distances, road congestion, loss of open space, and loss of a sense of place — associated with rapidly increasing land use per capita.¹⁸ To the extent that

¹⁶A good example of a policy that affects the price of residential land is the federal tax treatment of housing. James Poterba has estimated that federal tax policy lowers the after-tax cost of owner-occupied housing and, by extension, residential land, by 15 percent, assuming the market price of housing is unaffected by tax policy. The value of the tax break, however, is offset partially because the lower after-tax land prices increase demand, which, in turn, drives market prices up. Assuming that 20 percent of the tax benefit is capitalized into land prices (Sinai 1997), the housing tax break effectively lowers housing and residential land prices 12 percent.

¹⁷The tax policy increases land consumption by 19.2 percent (12 percent times 1.6). Remember, our estimation of the elasticity of demand for residential land is -1.6 percent. This increase in land consumption lowers density by 16.1 percent. Note that density equals population divided by land (P/L). Since current tax policy increases the demand for land 19.2 percent, density under that policy is $P/L(1.192)$. This represents a 16.1 percent decline in density from what it would be without the tax policy.

¹⁸Growth in land use per capita can arise from increasing lot sizes within communities or shifts in population from communities with small average lot sizes to communities with larger lot sizes. In the Philadelphia metropolitan area, most of the increase in land consumption per capita in the 1990s has not come from increasing average lot sizes in suburban communities but rather from declining population in the city of Philadelphia — which, on average has very small lot sizes — and increasing population in suburban counties, which have larger average lot sizes.

¹⁴Some caution is appropriate in generalizing the implications of elasticity estimates, since they are based on only one county. Although the estimates are based on a great deal of data for the county, other regions of the country conceivably could have different land-price elasticities.

¹⁵Another approach to the traffic congestion problem is more appropriate pricing of automobile travel. Many economists have suggested that tolls, time-of-day pricing and other user fees that reflect the true social costs of car travel would result in less congestion and more efficient development patterns. These approaches have not been widely accepted in the U.S.

The Bartik-Epple Approach to Estimating The Demand for Bundled Goods

I

n 1987, Timothy Bartik and Dennis Epple independently suggested similar approaches to the problem of estimating the demand functions for goods that are sold bundled together with other goods, like land and housing. They suggested that one could overcome the

econometric problem associated with the fact that consumers choose over a range of both price and quantity of residential land if observations on multiple markets were available that met two essential conditions: (1) the distribution of household preferences were unchanged across different market observations; and (2) there must be forces — changes in incomes and prices — that shift household budget constraints across markets. The first condition ensures that, across markets, we would observe differences in quantity of land consumed that reflect, on average, responses to changes in prices and incomes, rather than differing tastes among households. The second condition enables the researcher to statistically isolate the changes in land consumption in response to a change in price.

Using data on nearly 100,000 housing transactions spanning 26 years in Montgomery County, Pennsylvania, we estimate the demand function for residential land following the Bartik-Epple technique. We treat each year of data as a separate market. Essentially, we assume that the preferences of the population do not systematically change over time. With regard to the observable attributes of the Montgomery County population, there has been very little change (aside from size of population) in the underlying composition of households over the period. During this period, there were many factors affecting the supply and demand for housing that shift household budget constraints, including (1) employment shifts in Philadelphia and the suburbs, which served to shift the demand for housing in different ways throughout the county^a; (2) changes in mortgage rates that affect the cost of financing home purchases and are related to the finance costs builders face during the construction of new houses; and (3) the supply of available land changed over time at different rates throughout the county, which affected the market prices of houses and, hence, consumers' budget constraints.

^aThus, shifts in housing demand associated with changes in city and suburban employment vary across space as well as time. See my 1999 article.

Once our data meet the fundamental requirements of the Bartik-Epple procedure, a two-step procedure is then used to estimate the demand for residential land. First, using hedonic regression techniques, we estimate the relationship between house value and housing traits for each house in each year. Then, using this estimated relationship for each market (year), we compute the implicit price of residential land separately for each house. Second, we estimate a function that describes the relationship between the implicit price of land and the quantity of land consumed.^b The estimation is done in such a way that the changes in quantity reflect those changes associated with shifts in the budget constraint and not changes that reflect differences in tastes across households. This is accomplished using an instrumental variables approach that purges the changes in quantity of land consumed that are due to differences in individual tastes. We use the variables described above as instruments for the quantity of land consumed.^c These variables shift the consumer's budget constraint and, hence, shift the quantity of land consumed without shifts in preferences. The instrument equation yields the predicted quantities of land that differ across households only as a result of differences in households' budget constraints (not preferences). We then estimate the relationship between changes in land prices and changes in the quantity of land consumed.^d These estimates allow the computation of unbiased estimates of the price elasticity of demand for residential land.

^bThe function estimated is an "inverse demand" function because price is on the left-hand side of the equation and quantity consumed is on the right-hand side.

^cSpecifically, the instruments include supply shifters: number of new homes built in the tract each year, fraction of homes in a tract each year that are new, census tract size in square miles, vacant land in the tract available for residential development; demand shifters: suburban employment growth lagged one year, suburban employment growth lagged two years, Philadelphia employment growth lagged one year, Philadelphia employment growth lagged two years, suburban and city employment growth rate lagged one year interacted with municipality dummy variables; and variables that affect supply and demand: fraction of households that moved between 1975 and 1980, fraction of heads of household between the ages of 35 and 54, annual mortgage rate, annual mortgage interest rate interacted with municipality dummy variables, total number of sales in the tract each year, and dummy variables for year of sale.

^dFor the inverse demand function to be identified, the instruments that shift supply (and hence are not included in the demand equation) must be significant. Our supply shifters in the instrument equation are all highly significant.


smart growth policies limit the amount of land available for residential use, they will drive up land prices, imposing additional costs on households.

Both the cost and the effectiveness of the policies in achieving their goals will depend, in part, on how responsive to price households' demand for land is. If, for example, consumers were very unresponsive to the price of land, local policies restricting the availability of residential land would have two effects. Because large price changes would be required to make households reduce their land consumption, land prices would likely rise a great deal in response to constraints on supply. High prices for land would result in some decline in land consumption per household, but only at a relatively large cost in terms of households' living standards. Also, higher prices for land would lead some households to seek new communities with lower land prices.

The bottom line is that if consumers are unresponsive to land prices, policies restricting the quantity of residential land will impose high costs on consumers and will likely induce households to circumvent the intent of the restrictions by moving to communities without restrictions. Ironically, if there are communities without land-use restrictions adjacent to communities that enforce such restrictions, the net result may be additional geographic decentralization with little overall impact on density.

Our estimates, however, suggest that consumers' consumption of land is quite responsive to changes in price. This finding raises the likelihood that smart growth policies will have larger effects on patterns of metropolitan land use, at lower costs to households than previously thought. When demand for residential land is elastic — that is, small increases result in relatively large

adjustments in the amount of land consumed — consumers will substitute consumption of other goods for consumption of land relatively easily. Thus, policies that reduce the supply of land will result in increased prices for land, but these increases will be relatively modest. This suggests that the hardship associated with smart growth policies will be smaller than if demand were inelastic and, furthermore, that public policies designed to increase land prices and reduce households' land consumption will likely be more effective in increasing residential density.

Our findings imply that the low-density patterns of residential development so dominant in the United States may reflect not a unique American taste for large lots and open space but rather the low price of residential land. Moreover, public policies that affect the price of residential land could significantly alter these patterns of development. 

REFERENCES

Alonso, William. *Location and Land Use*. Cambridge, MA: Harvard University Press, 1964.

Bartik, Timothy J. "The Estimation of Demand Parameters in Hedonic Price Models," *Journal of Political Economy*, Vol. 95, 1 (February 1987), pp. 81-88.

CB Richard Ellis. *Global Market Rents: Office Rents & Occupancy Costs Worldwide*. CB Richard Ellis Global Research and Consulting, July 2000.

Cheshire, Paul, and Stephen Sheppard. "Estimating the Demand for Housing, Land, and Neighborhood Characteristics," *Oxford Bulletin of Economics and Statistics*, Vol. 60, 3 (1998), pp. 357-82.

Epplé, Dennis. "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products," *Journal of Political Economy*, Vol. 95, 1 (February 1987), pp. 59-80.

Glaeser, Edward, Mathew Kahn, and Jordan Rappaport. "Why Do the Poor Live in Cities?" Harvard Institute of Economic Research Discussion Paper 1891 (April 2000).

Mieszkowski, Peter, and Edwin Mills. "The Causes of Metropolitan Suburbanization," *Journal of Economic Perspectives*, Vol. 7, 3 (Summer 1993), pp. 135-47.

Muth, Richard F. "The Derived Demand for a Factor of Production and the Industry Supply Curve," *Oxford Economic Papers*, July 1964, pp. 221-34.

Muth, Richard F. "The Derived Demand for Residential Land," *Urban Studies*, 1971, pp. 243-54.

Poterba, James. "House Price Dynamics: The Role of Tax Policy and Demography," *Brookings Papers on Economic Activity*, 1991, 2, pp. 143-203.

Rybczynski, Witold. *City Life*. Touchstone Books, 1996.

Sinai, Todd. "Are Tax Reforms Capitalized into House Prices?" in *The Effect of Tax Reform on the Owner-Occupied Housing Market*. MIT Ph.D. thesis, May 1997.

Voith, Richard. "The Suburban Housing Market: Effects of City and Suburban Employment Growth," *Real Estate Economics*, 1999.