

Essays on Annuitization and Housing Choice

by

Thomas Davidoff

B.A. Social Studies, Harvard College, 1993

Submitted to the Departments of Economics and Urban Studies and
Planning

in partial fulfillment of the requirements for the degree of

Doctor of Economics and Urban Studies and Planning

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2002

© Massachusetts Institute of Technology 2002. All rights reserved.


Author
Departments of Economics and Urban Studies and Planning
May 1, 2002

Certified by ..
Peter Diamond
Institute Professor of Economics
Thesis Supervisor

Certified by ...
William Wheaton
Professor of Economics
Thesis Supervisor

Accepted by ...
Peter Temin
Elisha Gray II Professor of Economics
Chairman, Department Committee on Graduate Students

Accepted by
Alice Amsden
Professor of Political Economy
Chair, PhD Committee, Department of Urban Studies and Planning



Essays on Annuitization and Housing Choice

by

Thomas Davidoff

Submitted to the Departments of Economics and Urban Studies and Planning
on May 1, 2002, in partial fulfillment of the
requirements for the degree of
Doctor of Economics and Urban Studies and Planning

Abstract

Chapter 1 For most US households, labor income is the most important source of wealth and housing is the most important risky asset. A natural intuition is thus that households whose incomes covary relatively strongly with housing prices should own relatively little housing. Under plausible assumptions on preferences and distributions, this result holds theoretically. Empirically, I find a significant effect: among US households, a one standard deviation increase in income-house price covariance is associated with a decrease of approximately \$25,000 in the value of owner occupied housing. This empirical result implies greater cognizance of the interaction between labor income and asset risk on the part of households than suggested by most analyses of stock market behavior. The analysis also suggests that many homeowners enter financial markets in a riskier position than typically thought, and reinforces the intuitive appeal of proposals for market- or tax-based risk sharing in housing prices.

Chapter 2 extends the theory of annuitization with no bequest motive in two directions. First, we derive sufficient conditions, in a more general setting than Yaari (1965), under which complete annuitization is optimal, and weaker conditions under which partial annuitization is better than zero annuitization. Second, we explore how incremental and complete annuitization affect consumer welfare in these more general conditions. When markets are complete, all savings are optimally annuitized as long as there is no bequest motive and annuitized assets have greater returns than conventional assets. Consumers' utility need not satisfy intertemporal additive separability nor the expected utility axioms, and annuities need not be actuarially fair. The result is weakened if annuities markets are incomplete, so that there are some assets which do not exist in annuitized form: as long as trade occurs all at once and consumption is positive in every state of nature, a small degree of annuitization is better than no annuitization. When conventional asset markets are incomplete, if annuities are illiquid, then it is possible that no savings are annuitized. We present numerical calculations of the financial benefit and optimal degree of annuitization for consumers with standard CRRA preferences, and compare these results to results where otherwise identical consumers have utility that depends both on present consumption and a standard of living to which they have grown accustomed. In our specification, the

effect of adding intertemporal dependence hinges on the size of initial standard of living relative to resources.

Chapter 3 addresses the measurement of income sorting and the attribution of observed sorting to different causes. In terms of measurement, I show that a standard decomposition of variance of household income into within jurisdiction and between jurisdiction components understates sorting in the presence of measurement error. Using 1990 US Census data, I find that adjusting for this error approximately doubles the estimated extent of sorting. On average, across all US metropolitan areas (MSAs) I find that approximately ten percent of the variation in household income can be explained by differences across jurisdictions. I attempt further to identify the extent to which the observed sorting may be attributed to a “Tiebout” mechanism, by which income sorting follows from sorting by preferences over governance into differently governed jurisdictions. I find that zip codes are considerably more homogeneous than jurisdictions, and that on average, neighboring zip codes in different jurisdictions are only slightly more different from each other than neighboring zip codes in the same jurisdiction. This result implies that we cannot safely assume that observed sorting on characteristics is driven by differences in government, and also implies that extra-governmental neighborhood characteristics are an important source of sorting.

Thesis Supervisor: Peter Diamond
Title: Institute Professor of Economics

Thesis Supervisor: William Wheaton
Title: Professor of Economics

Acknowledgments

For all that I have, I am grateful to my parents Paul and Linda Davidoff. In addition to being loving and supportive, they taught me to care about the welfare of others. My whole extended family has given me the opportunity to test the tenets of liberal orthodoxy through countless hours of political debate.

Hunter High School and Harvard are great places from which to start an academic career. The friends I accumulated in both places have been a constant source of encouragement and competition.

A lot of what I know about real estate and political economy I learned at Forest City Ratner, where Bruce Ratner was a great mentor. I owe Bobby Willig at Princeton for showing me that economics is a fun and interesting discipline.

At MIT, Peter Diamond, Bill Wheaton, Sendhil Mullainathan and Jim Poterba are great teachers and professional role models who were particularly generous with their time. I am enormously appreciative of all their help, and a lot of this thesis and whatever I produce in the future should be attributed to their teaching, advice and example. The second chapter of this thesis was co-written with Peter Diamond and Jeff Brown, another source of helpful advice and encouragement.

I have enjoyed working and socializing with many of my classmates. This thesis benefitted particularly from the thoughts of Hoyt Bleakley, Lakshmi Iyer, and Mark Lewis.

I am grateful for financial support from the Center for Retirement Research at Boston College pursuant to a grant from the U.S. Social Security Administration funded as part of the Retirement Research Consortium.¹

By far the best part of my MIT experience has been meeting and marrying my wife Dulcy Anderson, who has given all the advice, love, time and support a person could ever ask for. She also brought Kitty Petunia Golightly, to whom I dedicate this thesis, into my life.

¹Nothing in this thesis should be construed as representing the opinions or policy of the Social Security Administration or any agency of the Federal Government, or the Center for Retirement Research.

Contents

1	Labor Income, Housing Prices and Homeownership	13
1.1	Introduction	13
1.2	Housing Choice with Stochastic Labor Income and Prices	16
1.2.1	A Two Period Model: Key Features	16
1.2.2	Homeowner Utility Maximization	19
1.2.3	Renters' Expected Utility	26
1.3	Empirical Estimation of the Effect of Income-Price Covariance on Housing Demand	28
1.3.1	Equations to be estimated	28
1.3.2	Estimating the income-price variance-covariance matrix	31
1.3.3	Cross Sectional Homeownership and Demographics Data	38
1.3.4	Identification and Inference	39
1.4	Results	44
1.4.1	Effect of Covariance on Housing Purchases Combining the Extensive and Intensive Margins	44
1.4.2	Effect of Covariance on Housing Purchases Conditional on Homeownership	47
1.4.3	Effect of Covariance on Tenure Choice	48
1.5	Conclusions	48
2	Annuities and Individual Welfare (Joint with Jeffrey Brown and Peter Diamond)	57
2.1	Introduction	57

2.2	Annuitization When Trade Occurs All At Once	61
2.2.1	Two Periods, No Aggregate Uncertainty	61
2.2.2	Many Future Periods and States, Complete Markets	66
2.2.3	More than Two Periods, Complete Bond Markets, Incomplete Annuities Markets, Trading only Once	68
2.3	Annuitization with Trade in Many Periods	72
2.3.1	More than Two Periods, Complete Liquid Bond Markets, In- complete Illiquid Annuities Markets, Trading More than Once	72
2.3.2	Incomplete Markets With Illiquid Annuities and Future Trade	74
2.4	Valuation of a constant real annuity in special cases	75
2.4.1	Additively Separable Preferences, Constant Real Annuity Con- straint, Actuarially Fair Annuities, No Uncertainty Other Than Length of Life	77
2.4.2	Utility Dependent on Standard of Living	80
2.4.3	Numerically Estimated Magnitudes of Welfare Effects	82
2.5	Conclusion	86
3	Income Sorting: Measurement and Decomposition	97
3.1	Introduction	97
3.2	Measuring Income Sorting	100
3.2.1	Adjustment for Sampling Without Replacement	101
3.2.2	Measurement Error	101
3.3	Data	104
3.4	Results	105
3.5	Decomposing Income Sorting	107
3.6	Conclusions	111

List of Figures

1-1	Effect of labor income - price covariance $Cov(P, y)$ on housing purchases	29
2-1	Case 1	88
2-2	Case 2	89
2-3	Case 3	90
2-4	Case 4	91
2-5	Case 5	92
2-6	Case 6	93
2-7	Case 7	94
2-8	Case 8	95
2-9	Case 9	96
3-1	Income Sorting: Observed R^2 values and maximized values with “perfect sorting,” with and without measurement error	108

List of Tables

1.1	Summary Variance - Covariance Statistics	37
1.2	Summary Housing and Demographic Statistics	38
1.3	Household Investment Income by Housing Tenure	39
1.4	OLS Regressions of Value of Housing Owned on Demographic and Variance-Covariance Characteristics	51
1.5	OLS Regressions of Value of Housing Owned on Demographic and Variance-Covariance Characteristics. Cell Fixed Effects and Income Interactions with MSA and SIC Dummies Included	52
1.6	Regressions of $COV(P,y)$ on Demographic and Other Variance-Covariance Characteristics	53
1.7	IV Regressions of Value of Housing Owned on Demographic and Variance-Covariance Characteristics	54
1.8	Regressions of Housing Value on Demographic and Variance-Covariance Characteristics: Homeowners Only	55
1.9	Linear Probability Regressions of a Dummy Variable for Home Ownership on Demographic and Variance-Covariance Characteristics	56
2.1	Summary of simulations	86
3.1	R^2 Estimates of income sorting	106
3.2	Regressions of Income R^2 for Neighboring Zip Codes on a Variable Indicating Different Jurisdictions	111

Chapter 1

Labor Income, Housing Prices and Homeownership

1.1 Introduction

Households' imperfect ability to trade away risk associated with labor income and housing prices complicates standard portfolio analysis. For many households, housing dominates the portfolio and future labor income is the most important component of wealth. Under these conditions, it is natural to think that risk averse households will use housing purchases to hedge income risk. This paper evaluates the intuitive notion that households whose incomes covary relatively strongly with housing prices will purchase relatively little housing. I consider reduced housing purchases both on the extensive own-rent margin and on the intensive margin of value conditional on ownership.

A parallel literature shows that, under some conditions, investment in stocks decreases in the covariance between stock returns and labor income.¹ Formal empirical studies of investor behavior provide mixed evidence on the effect of income-stock return covariance on portfolio choice.² The existence of large holdings of employer

¹Viceira (2001) shows this for CRRA investors facing jointly normal stock price and wage distributions. Davis and Willen (2000) provide a similar result for CARA investors.

²Heaton and Lucas (2000) show that in a panel of US investors, the fraction of wealth put into

stock in retirement plans suggests (but of course does not prove) that a large fraction of investors fail to recognize the importance of income-return covariance to aggregate portfolio risk.

The existing theoretical examinations of how income-return covariance affects portfolio choice assume that investors choose how much stock to own, but do not own housing. My analysis starts from the opposite assumption, which I consider a much closer approximation of reality. Kennickel, Starr-McCluer and Surette (2000) estimate, based on the 1998 Survey of Consumer Finances, that in 1997, 66 percent of US households owned their own home. By contrast, only 56 percent did any form of saving and just 49 percent held any stock, directly or through mutual funds or retirement plans. Among homeowners, the median home value was \$100,000, whereas the median value of equities among those holding equities was \$25,000.

For the large majority of households, consumption of housing and investment in housing are closely linked: ownership of rental housing is highly concentrated, so that renters typically own no housing, and homeowners typically own as much housing as they consume. Thus, unlike stock purchases, housing purchases affect utility both through the budget constraint and through direct present and future consumption benefits. Future consumption of housing implies that house prices affect welfare through both the numerator and denominator of future real income, a point frequently lost in the housing literature.³

To my knowledge, only three other papers consider housing choice in the context of simultaneously uncertain housing prices and labor income. Campbell and Cocco (2001), Cocco (2000) and Yao and Zhang (2001) solve numerically for optimal lifetime mortgage and housing behavior, estimating a single population covariance matrix for

stocks decreases in the covariance between total income and stock market returns. However, their result appears to occur over a disproportionately old and high wealth sample and decomposing income into wage and entrepreneurial components, the estimated effect of labor income covariance is insignificant. Vissing-Jorgenson (2000) emphasizes her failure to reject the null hypothesis of no relationship between portfolio weight on stocks and correlation between income and stock market returns.

³Sinai and Souleles (2001) emphasize renters' implicit short position in future housing prices. Similarly, young homebuyers may expect to move on to a higher quality of housing, as in Ortalo-Magne and Rady (1998), so that their utility may be decreasing, rather than increasing in future prices.

prices, labor income and interest rates (and zero-covariance stocks in the case of Yao and Zhang), and assuming jointly normal distributions. By contrast, I confine the theoretical analysis to a two period setting, but allow for population heterogeneity in the covariance between labor income and housing prices, and describe analytically conditions under which housing purchases fall with covariance. The assumptions required seem quite reasonable, but the result does not follow directly from the primitive assumption of concave utility.

Empirical evaluation of the relationship between income-price covariance and housing purchases requires data on housing investment, the joint distribution of income and prices, and variables plausibly correlated with both (I use “price” to refer to the price of housing unless otherwise noted. For empirical purposes, I use real housing prices and labor income, deflated by the US Consumer Price Index for non-housing goods). The standard approach to estimating the covariance between labor income and asset prices is to examine the co-movements between prices and the incomes of a panel or repeated cross section of households. I instead estimate the covariance between the mean wages paid by different industries at the metropolitan area (MSA) level, and MSA housing prices, and impute these estimated covariances to a large cross section of households from the 1990 US Census. Thus, for example, I obtain separate covariance estimates for retail workers and construction workers in the Boston MSA, and separate covariance estimates for retail workers in Boston and retail workers in Detroit. This approach lets me estimate covariances with local housing prices, which is important given the heterogeneity in price movements across markets. Further, it is plausible that households form estimates of income-price covariance based not on their personal histories, but rather on the experience of the industry in which they work. The cost of my approach is that job separations, geographic mobility and intraindustry differences in income movements are missed.

The equation of primary interest is a regression of the dollar value of housing owned (which takes on a value of zero for renters) on income-price covariance, expected growth and variance of income and prices, demographic controls, and dummy variables indicating industry (two-digit SIC code) and MSA. I also estimate sepa-

rately the effect of covariance on the intensive margin of purchases conditional on owning, and on the probability of deciding to own rather than rent.

The second section of this paper presents a model of housing choice with uninsurable labor income and uncertain housing prices and lays out sufficient conditions for the result of decreasing housing purchases with increasing income-price covariance. In the case of additive mean variance preferences with no future repurchase of housing, the result holds unambiguously. The third section details the panel data on wages and prices and the cross-sectional microdata on housing investment I use to estimate the effect of covariance on the value of housing owned. As expected, I find generally positive covariances, with correlations larger in the “right” industries, such as stock brokerage in New York and amusement in Orlando. In the fourth section, I present regression results. The estimated effect of covariance on housing owned is consistently significantly negative, and using instrumental variables to overcome measurement error increases the estimated effect dramatically. Combining effects on both the extensive and intensive margins, I find that a one standard deviation increase in income-price covariance is on average associated with a reduction in the value of housing owned of approximately \$25,000. This negative effect operates on both the extensive and intensive margins. The fifth section concludes with a discussion of the consequences of the results for our understanding of households’ financial risk and their awareness of this risk, and for the potential gains to households from public and private sector mechanisms to offset housing risk proposed by Berkovec and Fullerton (1992) and Shiller (1993).

1.2 Housing Choice with Stochastic Labor Income and Prices

1.2.1 A Two Period Model: Key Features

Present housing decisions affect lifetime utility directly through the benefits of consuming more or less housing, and indirectly through the lifetime budget constraint. I

assume that housing investment and consumption are non-separable: renters' housing consumption is free to vary, but investment must equal zero; for owners, the quantity of housing owned must equal the quantity of housing consumed.⁴

How present housing investment affects the budget constraint depends on both present and future housing prices. Purchasing housing involves the sacrifice at the time of purchase of a combination of debt and equity in the amount of the (hedonic) quantity purchased times the present (hedonic) price, HP_1 . Whenever the house is resold, "period 2," absent transaction costs, today's housing investment yields HP_2 .

The date of resale is surely uncertain to homebuyers at the time of purchase. Venti and Wise (2000) show that older households appear generally to cash out only when severe financial shocks such as the need for long term healthcare arise. Presumably, in considering expected resale value, and the riskiness of resale income, homeowners weight the distribution of prices in each future state by the probability that sale will occur in that period. Because of transaction costs, homeownership is unattractive if households expect to move within a short horizon. Rather than imposing a probability distribution of moves over multiple horizons, I will assume that all homeowners know for certain that they will resell at a future date which is fixed before housing purchases are made. In reality, habit formation and transaction costs presumably endogenize the date of resale to the housing choice problem.⁵

As discussed above, housing generally swamps non-housing investment in portfolios. This justifies the simplification that no other risky assets are available. Naturally, I allow homeowners to take on mortgage debt. I assume that mortgage debt is riskless both in the sense that the interest rate is deterministic and in that default is not possible.⁶ With unrestricted mortgage choice, households separately choose how much

⁴These constraints can be relaxed through direct or indirect ownership of rental real estate. However, the fraction of working aged households who own rental real estate either directly or through Real Estate Investment Trust shares is small, and the fraction of renters owning such assets is particularly small. Indirect real estate ownership through pension funds is presumably more widespread, but the exposure to local price risk likely small in most cases.

⁵For example, a household purchasing a three-bedroom house today may find it more inconvenient and psychologically difficult to move to a one-bedroom apartment in retirement than would a similar couple purchasing a two-bedroom apartment today. However, the former couple might find the adjustment more worth the trouble.

⁶In evaluating the effects of income-price covariance empirically, I will consider the possibility that

housing to purchase and how much of the other good to consume in the present. This is only an approximation to reality: mortgage rates are typically lower than consumer loan rates, and households may be constrained by a debt-equity ratio.⁷

I do not incorporate government policy into the analysis: US households are almost universally able to avoid capital gains taxes on housing. The progressivity of taxes and social insurance programs attenuate income risk, so income should be thought of as after tax and transfers. It would be interesting to extend the empirical analysis in this way. Similarly, the deductibility of mortgage interest realistically creates heterogeneity in what I assume to be a constant borrowing rate.

These assumptions allow us to confine the analysis to a two-period setting. In period 1 households earn labor income and purchase or rent housing. For households choosing to purchase, there is a simultaneous decision of how large of a mortgage M to take on. Conditional on owning, the difference between first period income and the equity put into the home ($y_1 + M - HP_1$) goes to consumption of a composite non-housing numeraire good. While utility is defined over housing and numeraire consumption, more intuitive analytical results can be described if choice is considered to occur over housing and mortgage debt, with numeraire consumption implicit. Renters' first period numeraire consumption is equal to first period income less rental payments. Allowing renters to borrow or lend a riskless asset would not affect the analysis.

Period 2 represents the date of both homeowner resale and lease expiration for renters. At this time, households earn stochastic labor income, pay the principal and interest R on any mortgage taken on in the first period, take in the value of their home and allocate wealth optimally between housing and the numeraire good. This gives rise to indirect utility v .⁸ From a first period perspective, utility over first period

households use present interest rates to predict future changes in income and house prices, and the possibility that comovements between labor income and housing prices are driven by comovements between labor income and interest rates.

⁷Alternatively, Laibson, Repetto and Tobacman (2000) suggest that households may use housing purchases as a device to force themselves to save.

⁸This leaves open the frequently observed outcome that households never sell their homes: because I have not included transaction costs, a non-sale corresponds to purchase and resale of the same quantity of housing.

housing and numeraire consumption are deterministic, but future indirect utility is stochastic, depending on the realization of period two income and housing prices.⁹ I assume that first period housing decisions do not affect the realization of second period income, ignoring potential psychological benefits to homeownership.

I assume that no individual's demand affects housing prices, so that present and future prices are taken as exogenous. Although with imperfectly elastic housing supply, we would expect positive correlations between income and prices on average, as a matter of interpretation it should be emphasized that covariances are not restricted to be positive, nor need they be destructive of welfare. Indeed, renters may benefit from positive real income-price correlations and homeowners from negative correlations. This may partially explain the persistent home ownership of the elderly observed by Venti and Wise (2000) and Sinai and Souleles (2001).

In period one, households calculate expected lifetime utility under optimal behavior conditional on renting and on owning and choose the regime with the greater expected level. I consider first the utility maximization problem conditional on deciding to purchase a home. The objects of interest here are the effect of an increase in income-price covariance $Cov(P, y)$ on optimal housing purchases H and on expected utility conditional on owning. I consider expected utility conditional on renting later.

1.2.2 Homeowner Utility Maximization

Combining the assumptions above with intertemporal additivity of utility, conditional on deciding to own, expected utility is given by:

$$U(H, M, \Theta, Z) = u(y_1 + M - HP_1, H, Z) + Ev(W_2, P_2, Z, \Theta) \quad (1.1)$$

u is a concave utility function of first period numeraire and housing consumption.

⁹The chief cost of such a two period description relative to a many period model is that I am assuming a resolution of uncertainty upon resale. Incorporating continuing uncertainty might change the effect of prices on indirect utility, but the potential for added insight strikes me as small relative to the cost in complexity. Realistically, renters' and homeowners' investment horizons are different: lease expiration is typically one year, whereas homeowners rarely sell so quickly. I take this issue up in the empirical section below.

v is an indirect utility function, concave and increasing in second period numeraire wealth, W_2 , and nonincreasing in second period relative housing price P_2 . Z denotes household characteristics which shape preferences, and Θ is the set of relevant moments of the joint distribution of period two income and housing prices.

Defining the mortgage rate as R , second period wealth is:

$$W_2|_{own} = y_2 + HP_2 - RM. \quad (1.2)$$

The concavity assumptions imply that expected utility is maximized when housing purchases and mortgage debt satisfy the first order conditions:

$$0 = U_H = -P_1u_1 + u_2 + E(P_2v_1), \quad (1.3)$$

$$0 = U_M = u_1 - REv_1. \quad (1.4)$$

Effect of increasing covariance on conditional housing purchases

Expected second period utility will, in general, depend on all the moments of the joint distribution of future housing prices and income. If we consider a change in a particular parameter of the joint distribution θ , holding characteristics Z and the rest of the moments Θ constant, then we can think of the other moments as fixed parameters of the utility function. We can thus rewrite expected utility (1.1) conditional on Z and all of Θ except for θ as

$$U(H, M, \theta).$$

Noting the optimality conditions:

$$0 = U_{HM}dM + U_{HH}dH,$$

$$0 = U_{HM}dH + U_{MM}dM,$$

total differentiation of the first order conditions (1.3) and (1.4) gives us two equations in two unknowns, which can be solved jointly for the change in optimal housing purchases associated with a small increase in the parameter θ . These total derivatives are given by:

$$0 = U_{M\theta} + U_{MM} \frac{dM}{d\theta} + U_{MH} \frac{dH}{d\theta}, \quad (1.5)$$

$$0 = U_{H\theta} + U_{HM} \frac{dM}{d\theta} + U_{HH} \frac{dH}{d\theta}. \quad (1.6)$$

Combining conditions (1.5) and (1.6), and rearranging gives the result:

$$\frac{dH}{d\theta} (U_{MM} U_{HH} - U_{MH}^2) = -U_{H\theta} U_{MM} + U_{MH} U_{M\theta} \quad (1.7)$$

The term multiplying the derivative of interest $\frac{dH}{d\theta}$ must be positive by concavity of u and v (see, for example, Mas-Collel, Whinston and Green (1995), Appendix D). The second derivative U_{MM} similarly must be negative, so dividing equation (1.7) by $-U_{MM}$ we have the relation:

$$\text{sign}\left(\frac{dH}{d\theta}\right) = \text{sign}\left(U_{H\theta} - \frac{U_{MH}}{U_{MM}} U_{M\theta}\right). \quad (1.8)$$

Intuitively, a parameter shift tends to reduce the quantity of housing if the shift reduces the marginal benefit of housing purchases. This effect is modified by changes in mortgage debt if changes in housing investment affect the marginal benefit of mortgage debt. An induced increase (decrease) in the marginal benefit of mortgage debt tends to increase (decrease) housing purchases if increased housing investment makes mortgage debt relatively attractive. The opposite implications arise if mortgage debt becomes less attractive with housing purchase.

In our case, the distributional parameter of interest θ is the covariance between income and prices, $Cov(P, y)$. Equation (1.8) implies the following result:

Result 1 *The following are sufficient conditions for housing purchases conditional*

on ownership to decrease in the covariance between labor income and housing prices, holding all other relevant moments of the joint distribution constant:

- The second derivative $U_{HCov(P,y)}$ is negative and $U_{MCov(P,y)}$ is zero
- $U_{HCov(P,y)}$ and $U_{MCov(P,y)}$ are negative and U_{HM} is positive.

Additively Separable Mean Variance Utility

The first condition for housing purchases to decrease in covariance in Result 1 is satisfied under a pair of assumptions shared by Berkovec and Fullerton (1992) and Flavin and Yamashita (2001). These papers specialize the homeowners' maximization problem by assuming that (i) housing is purchased only once, so that expected indirect utility Ev in equation (1.1) depends only on the distribution of future wealth; and (ii) expected indirect utility depends only and additively on the mean and variance of second period wealth:

$$Ev = a(EW_2) + b(Var(W_2));$$

$$a' > 0, b' < 0.$$

With wealth given by (1.2), and the borrowing rate R fixed between purchase and sale of housing, the variance of future wealth is given by:

$$Var(W_2) = Var(y_2) + 2HCov(P, y) + H^2Var(P_2). \quad (1.9)$$

In this case, an increase in covariance (holding expected income and prices constant) has no direct effect on the first period utility or on the value of expected second period wealth. Equation (1.8) thus reduces to:

$$\text{sign}\left(\frac{dH}{dCov(P, y)}\right) = \text{sign}\left(b' \frac{\partial^2 Var(W_2)}{\partial H \partial Cov(P, y)}\right) = \text{sign}(2b') < 0.$$

Hence, in this setting, optimal housing purchases conditional on owning are decreasing in covariance, matching intuition. We can also see that for constant variance

and mean growth in income and prices, for any positive level of housing, the variance of wealth is increasing in the covariance term. Thus, expected utility falls for any level of housing, and by implication, expected utility conditional on owning must fall.

Both mean-variance utility and the absence of future housing purchases are highly restrictive assumptions. Normality of prices and income are rejected empirically. Quadratic utility, required to guarantee mean variance preferences absent knowledge of the distribution of wealth¹⁰ implies counterintuitively increasing absolute risk aversion. More importantly, indirect utility will take the price of housing as a separate argument unless homeowners are certain that when they dispose of their home, they and the heirs they care about will be dead or in a place with uncorrelated housing prices.¹¹

General Indirect Utility

Without the mean-variance and no future purchase of housing assumptions, an increase in the covariance between income and prices will affect the net marginal benefit of both housing and mortgage debt by changing the riskiness of future real wealth. To obtain a clear prediction on the effect on housing purchases, we must appeal to the second condition of Result 1. That is, we need to show that the marginal benefit of housing is decreasing in covariance, that the marginal benefit of housing is not decreasing in mortgage debt and that the marginal benefit of mortgage debt falls with increasing covariance. The result of monotonically decreasing housing purchases in income-price covariance can survive failure of the latter two conditions, but cannot be guaranteed without them.

Using the first order conditions (1.3) and (1.4), noting that income-price covariance has no effect on deterministic first period utility and using properties of expectations, we obtain:

¹⁰see Rothschild and Stiglitz (1970).

¹¹This assumption may not be too restrictive with respect to the representative stockholder.

$$U_{HCov(P,y)} = EP_2 \frac{\partial Ev_1}{\partial Cov(P,y)} + \frac{\partial Cov(P_2, v_1)}{\partial Cov(P,y)}, \quad (1.10)$$

$$U_{MH} = u_{12} - P_1 u_{11} - RE(P_2 v_{11}), \quad (1.11)$$

$$U_{MCov(P,y)} = -R \frac{\partial Ev_1}{\partial Cov(P,y)}. \quad (1.12)$$

To a first order, the effect of an increase in income-price covariance $Cov(P, y)$ on the covariance between price and marginal utility of wealth, the second term on the right hand side of (1.10), is equal to the second derivative of indirect utility with respect to wealth, v_{11} .¹² By concavity, this term is negative. This effect is intuitive: marginal utility is decreasing in labor income, so if prices are high when income is high, larger realizations of price are associated with smaller realizations of income.

The effect of a change in income-price covariance on the product of expected housing prices and expected marginal utility $E(P_2 v_1)$ is more difficult to sign.¹³ Expected marginal utility is typically judged to be increasing in risk¹⁴, so we can infer that the effect of income-price covariance on expected marginal utility will be positive if in-

¹²To see this result, define a vector valued function of price and income, conditional on first period decisions and parameters as follows:

$$X = \begin{bmatrix} P_2 \\ y_2 \end{bmatrix}, F(X) = \begin{bmatrix} P_2 \\ v_1(y_2 - MR + P_2 H, P_2) \end{bmatrix}.$$

To a first order approximation,

$$\begin{bmatrix} Var(P_2) & Cov(P_2, v_1) \\ Cov(P_2, v_1) & Var(v_1) \end{bmatrix} \approx \begin{bmatrix} \frac{\partial P_2}{\partial P_2} & \frac{\partial P_2}{\partial y_2} \\ \frac{\partial v_1}{\partial P_2} & \frac{\partial v_1}{\partial y_2} \end{bmatrix} \begin{bmatrix} \sigma_h^2 & Cov(P, y) \\ Cov(P, y) & \sigma_y^2 \end{bmatrix} \begin{bmatrix} \frac{\partial P_2}{\partial P_2} & \frac{\partial v_1}{\partial P_2} \\ \frac{\partial P_2}{\partial y_2} & \frac{\partial v_1}{\partial y_2} \end{bmatrix};$$

so

$$Cov(P_2, v_1) \approx \frac{\partial v_1}{\partial P_2} \sigma_h^2 + \frac{\partial v_1}{\partial y_2} Cov(P, y).$$

¹³One might consider that differences in covariance are generated by purely idiosyncratic shocks to income and housing prices, uncorrelated with the underlying distribution, so that covariance increases through the covariance in the idiosyncratic shocks. In this case, a Taylor expansion of marginal utility would show no first order effect of an increase in covariance, and a second order effect proportional to the expectation over the distribution of price and income shocks of $Hv_{111} + v_{112}$. This is qualitatively the same result as discussed in the text. More likely, differences across groups in covariances arise from differences in income and prices that are correlated with the existing distributions, calling a Taylor expansion approach into question.

¹⁴This is a precondition for precautionary savings. Venti and Wise (2000) present evidence suggestive of precautionary motives with respect to home equity among the elderly.

creasing covariance increases risk. The effect of covariance on future risk is ambiguous in general. Absent future housing purchases, an increase in covariance, conditional on positive housing purchases implies increased variance (and presumably increased risk¹⁵) of future consumption, as noted above. However, non-housing consumption decreases with housing prices if households become net purchasers of housing in the future (either through purchase of a higher quality home, or a sufficiently long planned stay in sufficiently high quality rental housing). In this event, with additive separability in utility between housing and other consumption, an increase in income-price covariance acts to smooth marginal utility, thereby acting as a form of insurance. In the likely case that households trade up in quality in some cases, and down in others, or with strong complementarities between the two forms of consumption, the effect of covariance on risk becomes yet more difficult to evaluate.

Thus, an increase in income-price covariance has a first order negative effect on the benefit of purchasing housing by decreasing the covariance between expected prices and marginal utility. Precautionary saving motives work in the opposite direction for homeowners that can be considered to have a long position in local housing prices. A negative sign on the term $U_{HCov(P,y)}$ hence seems probable, but cannot be deduced from concavity alone.

Turning to equation (1.12), we expect intuitively that U_{HM} should be positive: purchasing more housing should make mortgage debt more attractive. However, there is the complication that increasing mortgage debt implicitly increases first period consumption. The cross partial u_{12} is difficult to sign: larger homes require more maintenance and afford more room for durable goods, but more comfortable homes might substitute for other goods. By concavity, $-P_1u_{11}$ is positive. Similarly, since second period prices are never negative, and v_{11} is everywhere negative, the term reflecting the cross effect on second period wealth $-RE(P_2v_{11})$ must be positive. Hence U_{HM} is positive unless there are very large negative cross-consumption effects in the first period.

¹⁵Without knowing the distribution of income and price shocks, we cannot be certain that an increase in variance increases risk, as emphasized in Rothschild and Stiglitz (1970).

The term $R \frac{\partial EV_1}{\partial Cov(P,y)}$, the negative of the cross partial $U_{MCov(P,y)}$, reflects precautionary motives. Again, for homeowners certainly in a long position in housing, this term is most likely positive, and $U_{MCov(P,y)}$ thus negative, if marginal utility is convex.

Summarizing, the second set of conditions in Result 1 for housing purchases to decrease monotonically in income-price covariance, conditional on ownership, seem likely to be met, but require assumptions on the parameters and functional form of both utility and the income and price distributions to be certain.

In the absence of future housing purchases, the increase in variance of future consumption engendered by increased income-price covariance most likely reduces expected utility conditional on home ownership, holding means constant. With large expected future housing needs, increasing covariance may act as a form of insurance, so that expected utility conditional on owning may be increasing. In this case, however, renters, who most likely hold a relatively shorter position in housing prices, presumably also benefit from the increase so that homeownership most likely remains relatively unattractive.

1.2.3 Renters' Expected Utility

We can think of renters as postponing the purchase of housing to a future date. At the end of a lease (typically one year, but frequently less - see US Census Bureau (1995)), renters may either rent housing again, or purchase housing. In either event, upon lease termination, present renters will have to pay for local housing services for the remainder of their stay in the same housing market. I assume that the price of housing upon lease termination is a sufficient statistic for the present value of these payments.¹⁶ Renters' utility is thus given by:

$$EU = u(y_1 - H \times Rent_1, H, Z, \Theta) + Ev(y_2, P_2, Z, \Theta),$$

It is natural to assume that renters' utility increases in the covariance between

¹⁶Empirically, I find a high correlation between (i) the covariance between income and prices and (ii) the covariance between income and rents.

income and prices: when prices are high, renters face diminished utility, and thus presumably value income relatively more in such states of nature. This intuition works most clearly with the additive mean variance utility over numeraire consumption discussed above, and with second period housing needs fixed at some level \bar{H} and hence irrelevant to the maximization. Renters' second period wealth then has mean and variance given by:

$$EW_2 = Ey_2 - \bar{H}EP_2$$

$$Var(W_2) = Var(y_2) - 2\bar{H}Cov(P, y) + \bar{H}^2Var(P_2).$$

In this case, the mean of second period wealth is not changed by an increase in covariance, but variance of wealth falls. Hence, renting is relatively more attractive with an increase in covariance.

Fixed future housing needs are a peculiar assumption in the context of housing choice. The fact that renters may substitute away from housing in high price future states clouds the insurance value of renting. For example, with additive log utility over consumption and housing, covariance between income and prices can be shown to have no effect on expected utility conditional on renting, holding the rest of the income-price distribution constant.¹⁷ With such preferences, however, utility conditional on ownership most likely falls, since real wealth presumably becomes riskier with increasing covariance between labor income and housing resale income. The important, and plausible, condition is that expected utility conditional on renting increases monotonically in covariance relative to expected utility conditional on owning.

¹⁷In this case, expected utility conditional on renting is

$$\log(y_1 - H_1 \times Rent) + \log(H_1) + E(\log(y_2 - H_2P_2) + \log(H_2)).$$

This implies constant expenditures on housing and expected second period utility is $E(\log(y_2) - \log(P_2))$, which does not change with covariance. It is interesting to note that households are risk seeking in the price level, and risk neutral with respect to log prices in this case.

1.3 Empirical Estimation of the Effect of Income-Price Covariance on Housing Demand

1.3.1 Equations to be estimated

The theoretical discussion suggests that we should observe empirically a relationship between covariance and housing purchases as depicted in Figure 1-1. The covariance between income and prices, $Cov(P,y)$ is measured on the horizontal axis; optimal housing purchases conditional on covariance (H^* , represented by the thick line) and maximized utility conditional on owning ($U|Own$) or renting ($U|Rent$) are measured vertically. Housing purchases are shown to decrease in covariance (the intensive margin) to a critical level Cov^* , above which a combination of the risk conditional on owning and the low consumption of housing implies greater expected utility conditional on renting (the extensive margin). Naturally, such a relationship will be conditional on covariates. This pattern of homeownership appears quite plausible theoretically, but because the results are not unambiguous, we cannot interpret an empirical test of the model as a test of rational investor behavior. Rather, we are jointly testing that the model presented has some application to risk as perceived by households and that households act on this risk.

Such a figure suggests estimation of the effects of increasing covariance on the extensive margin, the intensive margin, and the combined effect on both margins:


$$OWN = F(Cov(P, y), Z, \epsilon), \quad (1.13)$$

$$VALUE|OWN = b_0 + b_1Cov(P, y) + b_2Z + \epsilon, \quad (1.14)$$

$$VALUE = \beta_0 + \beta_1Cov(P, y) + \beta_2Z + \epsilon. \quad (1.15)$$

Here Z is a set of observable characteristics potentially correlated with demand for housing consumption or investment, $Cov(P, y)$ is the covariance between income and price levels, OWN indicates home ownership and ϵ represents idiosyncratic household

Figure 1-1: Effect of labor income - price covariance $Cov(P, y)$ on housing purchases



tastes for housing consumption and investment. In equation (1.15), $VALUE$ is a variable which takes on the value of a household's home if it is owner-occupied, or zero if the household rents. In the conditional regression (1.14), $VALUE|OWN$ is the value of a homeowner's house.

Whether we confine analysis to the intensive margin, or consider the intensive and extensive margins simultaneously, we cannot interpret the estimated coefficient on covariance as the effect of covariance on desired investment. Desired investment in housing is unobservable. What we do observe is optimal investment subject to the constraints that owner occupiers' investment must equal consumption¹⁸ and that renters may not own any housing. It is possible that some households would desire

¹⁸As noted above, failure to observe rental holdings impacts only a small fraction of non-landlord working age heads of households.

negative investment in housing without the constraints, and that the zero value assigned to renters is a form of censoring, but confining the sample to owner occupiers does not overcome the failure to observe desired investment.

Estimating the effect of covariance on either the intensive or extensive margin alone involves technical and data problems. A probit with the instrumental variables and fixed effects required for identification is computationally infeasible. I thus present a linear probability model in the extensive margin equation (1.13). In the absence of fixed effects and instruments, unreported probit and linear probability estimates are virtually identical. On the intensive margin (equation (1.14)), there is concern that high covariance households who choose to own may have unobservably large taste for owner occupied housing ϵ .¹⁹

I will thus focus on estimating the combined equation (1.15), but present estimates of the effect of covariance on the intensive and extensive margins separately. I regard as unrealistic the assumptions required to convert an estimate of b_1 or β_1 into a parameter of risk aversion, or cognizance of risk. However, assuming Z includes any demand characteristics plausibly correlated with income-price covariance, a significantly negative estimated value of β_1 implies that portfolio considerations do enter housing demand, and that, on average, households consider covariance between income and prices to augment the riskiness of homeownership. The estimated coefficient on covariance also informs whether joint income-price risk creates a distortion worthy of intervention.

Theoretically, both short and long term covariances should be considered as factors in housing purchases, as potential renters must consider the joint distribution of income and rents between lease signing and termination, typically one year. By contrast, homeowners have much smaller annual moving probabilities. However, since we do not know the actual horizon which homeowners use to consider risk, and because short and long term covariances are highly correlated, I will restrict myself to

¹⁹A Heckman sample selection approach is indicated, but requires either normality of errors or an exogenous shifter of housing tenure choice uncorrelated with demand conditional on ownership. A theoretically plausible candidate, mean length of time in the same residence by MSA-SIC cell, appears to meet neither criterion.

consideration of a single covariance, choosing five years as a reasonably long horizon which does not sacrifice too many observations with limited panel price and income data.²⁰

1.3.2 Estimating the income-price variance-covariance matrix

The covariance $Cov(P, y)$ must be estimated. To do so, I assume that real log labor income and housing prices follow AR(1) processes:

$$y_t = y_{t-5}(1 + g_y + \epsilon_{yt}) \Rightarrow \ln(y_t) \approx \ln(y_{t-5}) + g_y + \epsilon_{yt};$$

$$P_t = P_{t-5}(1 + g_h + \epsilon_{ht}), \Rightarrow \ln(P_t) \approx \ln(P_{t-5}) + g_h + \epsilon_{ht};$$

$$E\epsilon_{it} = 0, \quad i = y, h$$

$$E\epsilon_{it}^2 = \sigma_i^2 \quad i = y, h$$

$$E\epsilon_{ht}\epsilon_{yt} = \sigma_{hy}.$$

$$E\epsilon_{yt}\epsilon_{yt-x} = E\epsilon_{ht}\epsilon_{ht-x} = E\epsilon_{yt}\epsilon_{ht-x} = 0; \quad x \neq 0$$

Here, g_y and g_h are mean growth rates of income and housing prices and ϵ_{yt} and ϵ_{ht} are deviations from mean growth in year t .

The covariance between the level of income and prices between years t and $t+5$ is thus the product of period t income and price times the estimated covariance of percent (approximately log) changes. If we observed housing choice and income in period 1, we would calculate:

$$Cov(P, y) = y_1 P_1 \sigma_{hy},$$

²⁰Estimates based on one and three year horizons give results consistent with those reported below. The Chicago Title Company reports that approximately 55 percent of all homebuyers are repeat buyers. Given that the Census Bureau reports that approximately 17 percent of working age individuals move each year, with a constant hazard rate, approximately half of all present homebuyers will have moved within five years.

$$Var(y) = y_1^2 \sigma_y^2.$$

The interaction of income and log change covariances and variances to create level covariance and variance measures is notationally unpleasant. This problem cannot be overcome by estimating equations (1.13) through (1.15) in log form. The equation of primary focus (1.15) includes a dependent variable that takes on zero values, and covariance itself may be zero or negative. In the tables of results, I label $Cov(P, y)$ by $COV(P,y)$ and σ_{hy} by $COV(\ln P, \ln y)$.

In the model, P_1 is an hedonic price. For estimation purposes, I make the strong assumption that this price is equal (and normalized to \$1) for all households in all areas, so that the observed dollar value of housing is assumed equal to the hedonic quantity H . By including metropolitan area fixed effects (and interactions with income) in my regressions, I overcome some of the attendant problems. Because housing demand is likely nonlinear in price, this is an imperfect fix. However, the alternative of estimating different hedonic prices for separate housing markets would be highly suspect. With the assumption on hedonic prices, my estimate of the covariance between income and price levels is thus the product of income and the covariance of income and price shocks:

$$Cov(\hat{P}, y) = \hat{y} \sigma_{hy}.$$

Similarly,

$$Var(\hat{P}) = \sigma_P^2.$$

The standard approach to estimating household level income-return covariances is to compare asset price changes to changes in individual households' incomes, with household level data coming from panel or repeated cross sectional sources.²¹ Instead, I use wage data by industry (2 digit SIC code) and region (MSA). This data comes from the Bureau of Labor Statistics' Covered Employment Series, and covers the years 1975 to 1999. This deviation from standard practice is motivated by two

²¹As in Heaton and Lucas (2000), Vissing-Jorgenson (2000), Cocco (2000), Campbell and Cocco (2001).

considerations. First, good quality, regionally disaggregated household income data is not readily available for long time series, and housing price changes vary dramatically across regions²², rendering the covariance between income and national housing prices difficult to interpret. Second, it is not clear how individuals form expectations about the joint movements of their income and asset prices. It seems no less reasonable that they would consider the experience of the industry in which they work, than that they would consider their personal earnings history. A similar argument is implicit in Davis and Willen (2000). I use mean wages rather than aggregate wages to avoid overstated changes in wage prospects in industries with small numbers of employees. This compromise means that I fail (arguably rightly) to observe zero wages for the unemployed within any MSA-SIC cell.

The effect of variance of log prices and income and the covariance term on the variance of lifetime wealth will depend on individuals' expected length of stay in their industry and MSA and the sensitivity of their own wages to industry shocks. There is no clear prediction on the relative effects across ages. Younger households may have longer expected stays in an industry, but older, more senior workers' pay may be more sensitive to industry performance.²³

To estimate housing price changes, I use the Office of Federal Housing Agency Oversight's repeat sales Conventional Mortgage Housing Price Index, which provides indices for 148 MSAs for the years 1975 to 2001.²⁴

I deflate both income and house prices by the US consumer price index for all non-housing goods, so that variances, covariances and growth rates are in real terms. Local price indices are available, but less reliable than the national index.

The income and price innovations ϵ_y and ϵ_h are somewhat predictable based on lagged changes and current economic conditions, particularly interest rates. In some

²²The housing price data reveals a range from 6 percent (in Bellingham, WA) to 34 percent (in Charleston, SC) in deviation of percentage change in CPI deflated housing prices from 1995 to 2000 relative to mean changes over the period 1976 to 2000.

²³Splitting the sample, I find that the effect of covariance on younger households' purchases is slightly greater in magnitude than the same effect for older households.

²⁴The repeat sale methodology is meant to yield an index of prices for units of comparable quality: renovations and depreciation are not observed, and may bias the index.

specifications, I take out the components of the ϵ 's that are predictable based on interest rates, and I obtain almost identical results, as I do in unreported specifications removing the estimated effect of lagged shocks. In the former case, I regress log changes in price and income on lagged interest rates, and then estimate the covariance of the residuals. In this setting, I do not consider it appropriate to differentiate between economy-wide shocks and industry-specific shocks: this would be appropriate only if households held risk-minimizing positions in some regional index. I do, however, estimate the covariance between the wages of each MSA-SIC cell and aggregate MSA wages. This statistic may relate to the economic integration of local industry with their local economies, which might correlate with geographic integration and hence employee housing prices.

I estimate expected log price and income growth, the variance of ϵ_y , the variance of ϵ_h and the covariance σ_{hy} as means within MSA-SIC cells:

$$\hat{g}_y = \frac{\sum_{t=6}^T (\ln(y_t) - \ln(y_{t-5}))}{T - 5}$$

$$\hat{\sigma}_h^2 = \frac{\sum_{t=6}^T ((\ln(P_t) - \ln(P_{t-5}) - \hat{g}_h)^2)}{T - 6}$$

$$\hat{\sigma}_{hy} = \frac{\sum_{t=6}^T (\ln(P_t) - \ln(P_{t-5}) - \hat{g}_h)(\ln(y_t) - \ln(y_{t-5}) - \hat{g}_y)}{T - 6}.$$

Because I have overlapping five year changes, I estimate standard errors for the covariance estimates following Newey and West (1987), assuming a five year lag.

Variance-Covariance Results

Aggregate variance and covariance statistics are reported in Table 1.1. These statistics arise from a merge of the variance-covariance estimates with income, industry and MSA data from the 1990 US Census one percent metropolitan sample. The census population I consider consists of household heads with positive labor income, no retirement income and identifiable MSA and SIC categories for which time series

data was available both from the OFHEO house price and BLS wage series. These limitations leave me with just over 300,000 observations in 6,241 MSA-SIC cells.²⁵

The mean variance of log income growth ($\text{VAR}(\ln y)$) is approximately 0.9 percent, relative to mean growth ($\text{GROW}(y)$) of approximately 4.0 percent. Mean variance of housing prices ($\text{VAR}(\ln P)$) is 3.7 percent, around a mean five year growth $\text{GROW}(P)$ 4.7 percent. The mean log covariance ($\text{COV}(\ln P, \ln y)$) is 0.5 percent, associated with a mean correlation CORR of 0.32. There is considerable variation in the magnitude of covariance, and approximately one quarter of household heads work in industries with negative income-price covariances.

The variance of house prices must be multiplied by the value of housing squared to obtain the contribution of price variance to the variance of wealth. Assuming a fairly small \$100,000 house, the standard deviation of wealth attributable to housing wealth alone, at the mean variance level of 3.7 percent is approximately \$19,000. However, the contribution of price variance to income variance is overstated when repurchase is ignored.²⁶ The mean level of income variance suffers from measurement error in both income and estimated variance of income, and from the skewness of the income distribution, which bias the mean upward. On the other hand, the level of variance of income five years from today is just a fraction of the variance of lifetime income, because income is variable in intervening years, and because the level of income in five years informs all future earnings. Further, household incomes only imperfectly track industry incomes: idiosyncratic shocks to individual income will increase the variance attributable to income. The mean standard deviation of income is 2,729. With a reasonably large multiplier on annual income, the contribution to variance is thus potentially greater than that of housing price variance. The mean income-price correlation of 0.32 indicates that covariance contributes meaningfully to financial risk and this correlation is presumably biased to zero by measurement error.

$\text{COV}(\ln y_M, y)$ is the interaction of household head income and the covariance be-

²⁵To avoid nearly singular matrices, I also eliminate industries which never have more than 200 observations in any MSA.

²⁶Further, to the extent that homeowners can “time the market,” variance in prices may not be entirely bad.

tween log growth between cell wages and mean wages in the workers' MSA. $BETA(S,y)$ is the covariance between stock market returns (from CRSP's value weighted index) and cell income divided by the variance of stock market returns. $BETA(R,y)$ is the analogous measure for nominal interest rates. Notably, the stock beta measure is on average positive and significantly different from zero for more than half of the cells observed. This stands in contrast to the results of Davis and Willen (2000). The failure of that paper to find significant occupation-stock market covariances may be due to small samples or short horizon, and hence noisy (see Griliches and Hausman (1986)) estimation.

In stark contrast to the existing literature on housing and risk, I find similarly significant and typically positive Betas calculated for stocks and housing prices $BETA(S,P)$, with a mean of 0.07 and approximately half significantly different from zero. The conventional view (as in Flavin and Yamashita (2001)) that stock returns and house price increases are uncorrelated may again be premised on noisy short horizon estimation. In entering the stock market, workers must thus consider not only background income and price risk and stock market risk, but also considerable covariance between existing sources of wealth and stock market returns.

The results for particular industries (SICs) and MSA-SIC cells largely accord with intuition. The largest income price correlation at the national level (taking averages over regional cells' covariances with regional prices) belongs to the real estate industry, with a mean correlation of 0.61. Other large correlation industries are auto repair services and parking; automotive dealers; engineering, accounting research, management and related services; and building construction general contractors. Nationally, only two industries have negative mean covariances, mining and electronic and electrical components.

Several individual cells have incomes that are perfectly positively or negatively correlated with local housing prices: these may be statistical flukes and somewhat less informative. Among the perfect positive correlations are hotels and lodging in Oklahoma City and apparel manufacturing in Denver. Perfectly negative correlations include amusement and recreational services in Houston and lumber and wood

Table 1.1: Summary Variance - Covariance Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
COV(lnP,lny)	302,657	.0050281	.009061	-.082303	.2894776
COV(P,y)	302,657	169.6814	452.7648	-16,091.56	17,656.74
COV(lnP,lny) R	301,957	.0027901	.0054759	-.088588	.0676396
COV(P,y) R5	301,957	88.25521	267.844	-17,320.38	7240.162
COV(lny _{M,y})	302,657	124.4635	316.5523	-25,080.73	20,443.64
VAR(lnP)	302,657	.0371997	.0327821	8.75e-07	.2918917
INC*VAR(lnP)	302,657	1,276.875	1,954.107	.0006126	32131.4
VAR(lny)	302,657	.0087331	.0133751	1.51e-07	.8390149
VAR(y)	302,657	1.72e+07	8.76e+07	.0007541	1.52e+10
VAR(lnP)VAR(lny)	302,657	.0003491	.0007369	4.57e-12	.0956494
VAR(P)VAR(y)	302,657	775,312.5	5,069,265	.0000585	6.17e+08
BETA(S,y)	302,657	.142363	1.148891	-121.3615	9.453909
INC*BETA(S,y)	302,657	4526.118	74743.58	-2.34e+07	1646847
BETA(R,y)	301,957	-.0097335	.065485	-3.001811	16.75981
INC*BETA(R,y)	301,957	-299.5116	2,598.474	-195,117.7	653,632.8
BETA(S,P)	308678	.0721336	.2083547	-.9612067	1.03583
CORR	302,657	.3232355	.4332918	-1	1
INC*CORR	302,657	10,487.66	21,097.78	-195,875	197,111
GROW(y)	305,889	.0398326	.0561761	-.6964799	.73247775
GROW(P)	305,889	.0468219	0.0717214	-.3926537	.5175165

Notes: The level of observation is household heads in the 1990 US Census IPUMS 1 percent sample. Log covariances and Betas are calculated at the cell (MSA-SIC) level. COV(P,y) is equal to household head wage and salary income from the US times the cell level log covariance COV(lnP,lny). CORR is the correlation coefficient between log cell mean income and MSA housing prices. Variances and covariances of price and log price are identical by the assumption that the hedonic price is equal (and normalized to one) across MSAs).

manufacturing in Harrisburg. The MSA-SIC cell that partly inspired this study, stock brokers in New York City, have the very large correlation of 0.54. Amusement and recreational workers in Orlando also have a predictably large correlation at 0.64. Interestingly, the regionally identified oil industries in Houston do not have disproportionately large income-price correlations. Because covariances are estimated at a five year horizon, high frequency cyclical movements should largely disappear. This is appropriate in assessing the risks faced by homeowners with relatively long anticipated stays.

Table 1.2: Summary Housing and Demographic Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
INC	308,494	32,449.71	28,683.77	1	197,927
INC ²	308,494	1.88e+09	4.88e+09	1	3.92e+10
AGE	308,494	40.71609	11.65216	16	90
INC*AGE	308,494	1,378,805	1,414,118	40	1.61e+07
INC*AGE ²	308,494	6.24e+07	7.67e+07	1600	1.35e+09
MALE	308,494	.7219038	.4480618	0	1
INC*MALE	308,494	26,630.86	31,141.76	0	197,927
BLACK	308,494	.1018691	.3024766	0	1
INC*BLACK	308,494	2,223.031	8,477.105	0	197,927
FAMSIZE	308,494	2.749	1.528479	1	17
INC*FAMSIZE	308,494	93,347.6	109,810.7	2	1,942,160
OWN	308,494	.6270722	.4835839	0	1
VALUE	308,494	84,937.65	101,794.5	0	400,000
VALUE OWN	193,448	135,451.2	98,399.65	50,00	400,000
VALUE/INC	308,494	5.767175	786.1982	0	350,000
EDUC	308,494	11.37584	2.756154	1	17

Notes: Source: 1990 US Census Microdata merged with covariance estimates from BLS wage data and OFHEO house price indices. For any variable X, INC*X is equal to reported 1989 income times X. EDUC is a continuous measure of educational attainment.

1.3.3 Cross Sectional Homeownership and Demographics Data

In addition to the imputed variance - covariance and income variables, I observe, at the household level, housing tenure, the estimated value of owner occupied homes and standard demographic correlates with housing demand from Census microdata. These variables, describing the household heads' wage income, age, sex, race, family size and education are summarized in Table 1.2.

OWN is a variable that indicates homeownership; approximately 63 percent of the sampled household heads are homeowners. The variable VALUE, introduced above, is equal to the dollar value of owned housing, or zero for renters. Because this variable includes zeros for renters, there are more observations and a much lower mean than the conditional VALUE|OWN, which is the dollar value of housing for owner-occupiers only. VALUE/INC is equal to VALUE divided by reported 1989 wage income (INC).

While the level of asset wealth is not identified in the Census, total investment

Table 1.3: Household Investment Income by Housing Tenure

	OWN	RENT
All	193,448	115,046
Have Investent Income	97,512	28,267
Mean Investment Income	2,470	534
Median Investment Income	10	0
Mean Home Value / Monthly Rent	135,451	470
Median Home Value / Monthly Rent	112,500	437

Notes: Data comes from 1990 US Census microdata (1 % sample). Values are for household heads with identifiable MSA-SIC cells and positive labor income.

income is. Table 1.3 suggests that the assumption that only housing and debt are held is not a bad approximation for the households in question, although it must be noted that Kennickel et al. (2000) show that a majority of stock ownership is in the form of retirement plans. The almost complete absence of asset income among renters is particularly striking, and suggests that housing is, indeed, the dominant asset for US households. While approximately two-thirds of households own housing, less than one-third have any investment income. Mean investment income is approximately two percent of mean home value among owners, and accounts for approximately one and one-half mean month's rent for renters. Investment income is highly skewed in the population: median investment income is just \$30 for owners and zero for renters. There is some evidence of precautionary savings in the sample population: an unreported regression of investment income on characteristics and on industry variance of mean wages yields a significant and positive relationship.

1.3.4 Identification and Inference

A natural concern in estimating equation (1.15) is that the covariance between income and prices may be correlated with other demand factors. I control directly for the variables commonly thought to influence housing consumption and investment decisions described above and summarized in Table 1.2 and for dummies indicating marital status, both directly and interacted with income. These are similar to the regressors used in, for example, Ioannides and Rosenthal (1994), but I include income interactions and exclude some within-MSA geographic controls which are likely

endogenous.

Both housing investment and income-price covariance can be expected to be correlated with the mean growth and variance of income and prices, (and their interaction), discussed above and summarized in Table 1.1. I control for these variables as well as for MSA-SIC cell fixed effects, and interactions between income and a set of MSA and SIC dummy variables. The cell fixed effects remove the effect of the covariance between log income and house price shocks ($\text{COV}(\ln P, \ln y)$), and the variance of income shocks. However, the covariance between prices and income is equal to the product of income and these shocks, and the variance of income is equal to the variance of the shocks times income squared, as shown above. Thus, income variance and income-price covariance levels are only partly captured by the MSA and SIC interactions with income.

Some of the effect of income-price covariance on housing investment may stem from correlation with unobserved higher moments of the income distribution interacted with other moments of the price distribution. Such concerns can be allayed somewhat by instrumenting for covariance with the interaction of income and the correlation between income and prices, which removes the scaling by variances. This approach is also attractive given measurement error in covariance. Income-price covariance may also be correlated with the covariance between wages and mortgage rates or stock market returns discussed above, in that all of these variables indicate cyclical earnings. Households may be able to save around predictable price and income shocks, and thus I control for $\text{BETA}(R, y)$, $\text{BETA}(S, y)$ and $\text{COV}(y_M, y)$, discussed above and summarized in Table 1.1.

We might be concerned about selection: households who wish to purchase a large quantity of housing might choose occupations with low wage-price covariances. However, either direction of causality is consistent with household-level belief that increasing covariance increases homeowners' exposure to risk, and action upon that belief, and these are the objects of present interest.

Measurement error presents a major challenge. Income and cell level covariance of log income and price shocks are estimated with both conceptual and observational

error. Conceptual error arises in income because Census reported income is not equal to income in the year of housing purchase. Observationally, reported income is known to be a noisy measure of present income.²⁷

There is considerable conceptual error in the covariance estimates when applied to household heads. Individuals' income changes do not track industry mean wage changes both because of occupational mobility and because of different wage structures within industries. In general, this conceptual error cannot be considered noise with a mean of zero. Coefficient estimates must be taken as the effect of industry level covariance, rather than true covariance on individual housing choice. Thus, even if we assumed a utility function, recovering any underlying parameters would require a model of how individual income is linked to industry mean wages. Also, estimating covariances for all households based on the period 1975 to 1999 misses the fact that different covariances will apply to different households depending on the year of their purchase and their resale horizon. It is encouraging in this regard that covariances estimated in the first half of the time series are highly correlated with covariances using the second half. Finally, household heads' labor may not be the dominant source of some households' income.

Observationally, the log income-price shock covariance σ_{hy} is measured with significant error. I estimate a mean standard error of the covariance estimates of 0.02, considerably greater than the mean estimate of covariance of 0.005. The mean ratio of the absolute value of covariance to the standard deviation of the estimate is just 0.36. A particular source of concern is outliers. While 98 percent of observations have estimated σ_{hy} values less than 0.025, the maximal value is approximately 10 times this amount in magnitude, and the minimal value three times.

Interacting noisy covariance measures with income subject to reporting error can be expected to compound the problem of measurement error. If we interact observed σ_{hy} with noisily observed income, denoting by u measurement error in income and by

²⁷See, for example, Bound and Krueger (1991). There is also spousal income to consider - alternative specifications with covariances calculated as within-household weighted averages yield similar results, but cloud the interpretation of cell or SIC fixed effects.

v error in the covariance estimate, we obtain:

$$Cov(\hat{P}, y) = \hat{y}\hat{\sigma}_{hy} = (y + u)(\sigma_{hy} + v).$$

This object has a potentially very small signal to noise ratio, so that we expect considerable attenuation bias in OLS estimation. Assuming that true covariance is uncorrelated with income, and that the errors in each are uncorrelated, the signal to noise ratio of any particular observation should be increasing in the estimated variance of the error $Var(v)$ in measuring σ_{hy} , but weighting by an estimate of this variance will not eliminate attenuation bias.²⁸

A more promising approach to measurement error is to use alternative measures of covariance as instruments. In general, if we can find a $y\sigma_{hy}$ interaction measured with error orthogonal to that in the base estimate, then the IV estimate will not suffer from observational attenuation bias. A more likely outcome is that the instrument's errors will be partially correlated with the error in the original estimate. In this case, IV estimates can be expected to have some, but less bias than OLS estimates. A technique attributed to Wald (1940) is to use the sample rank of a mismeasured variable as an instrument. I present results with such an instrument, along with results using ranks of alternative measures of the covariance. I use the covariance between income and price changes, with shocks purged of components predictable based on interest rates ($COV(P,y)|R$) (the "interest rate adjusted" covariance, as opposed to the "standard" covariance) and the correlation measure $INC*CORR = y \frac{\sigma_{hy}}{\sqrt{(\sigma_y^2)\sqrt{(\sigma_h^2)}}$, both interacted with income as instruments for $COV(P,y)$. In the tables, rX denotes the rank of variable X .

To be explicit, the two stage IV estimator of β_1 in equation 1.15 comes from the

²⁸In the event that the innovation covariance σ_{hy} has a greater signal to noise ratio than the covariance product $y\sigma_{hy}$, it is tempting to divide the relationship $VALUE = \beta_0 + \beta_1(y\sigma_{hy}) + \beta_2Z$ by y , and consider the effect of the innovation covariance on the fraction of income spent on housing. I present such a regression, and find, not surprisingly, that standard errors are so large as to render almost all regressors insignificant. This approach also has the disadvantage of ruling out cell fixed effects.

equations

$$COV\hat{(P, y)} = \gamma_0 + \gamma_1 rC\tilde{O}V + \gamma_2 Z \quad (1.16)$$

$$VALUE = \beta_0 + \beta_1 COV\hat{(P, y)} + \beta_2 Z + \epsilon \quad (1.17)$$

Here, $rC\tilde{O}V$ represents the rank of the possibly alternative measure of $COV(P, y)$. Z includes the demographic and variance-covariance measures discussed above.²⁹

To get an idea of the importance of the IV approach, I estimate an OLS regression of the form

$$EDUC * BETA(S, y) = \eta_0 + \eta_1 COV(P, y) + \eta_2 VAR(y) + \eta_3 INC * VAR(P) + \sum_{m,s} \eta_{ms} CELL_{ms},$$

where $CELL_{ms}$ indicates working in MSA m in industry s .

We expect $COV(P, y)$ to be highly correlated with the product of education and the income-stock price beta, but expect attenuation bias in OLS estimates. Using the rank $rCOV(P, y)$ as an instrument for covariance yields an estimate of η_1 almost 100 times larger than the OLS estimate using uninstrumented $COV(P, y)$. This is a striking increase, but not terribly surprising given the very large measurement error noted above.³⁰

²⁹In undertaking such an approach, we have the opposite of the standard “weak instruments” problem: we do not want instruments to be too highly correlated with estimated covariance. Direct use of alternative measures implies sharing all of the time series error in income, and a large part of the observational error in the covariance of income and price shocks. We can expect the estimated rank of covariance, or the rank of alternative measures to be less strongly correlated both with true covariance and with the error term. An open question is, given a set of instruments measured with error, and only correlations observable, what is the best combination of instruments. With a large sample, we presumably value consistency over efficiency in the first stage, so that a comparatively small first stage R-squared is desirable.

³⁰Relating to the discussion in the footnote above, instrumenting with rank gives a first stage R-squared of .48. Instrumenting with the interest rate adjusted $COV(P, y)|R$, the first stage is stronger, with an R-squared of .72, but the IV estimate is still 10 times smaller than the estimate using rank.

1.4 Results

1.4.1 Effect of Covariance on Housing Purchases Combining the Extensive and Intensive Margins

The object of primary interest is the effect of income-price covariance $\text{COV}(P,y)$ on the value of housing owned VALUE . The additional right hand side control variables labeled Z in equation (1.15) are the demographic and variance-covariance variables discussed above and summarized in Tables 1.1 and 1.2.

Table 1.4 presents OLS estimates of such an equation. Column (1) presents estimates in the absence of the variance-covariance variables. The coefficients are generally as expected, with age and family size exerting significant positive effects on value, and a dummy variable for black household heads having a significant negative effect both in level and interacted with income. Column (2) incorporates the variance-covariance variables, both in levels and interacted with income. Most noteworthy, income-price covariance $\text{COV}(P,y)$ has a significant negative effect. Log covariance $\text{COV}(\ln P, \ln y)$ also has a significantly negative effect. Mean income and price growth predictably increase housing purchases. Notably, $\text{VAR}(y)$ has a significant positive effect, consistent with a precautionary investment demand for housing. Column (3) illustrates the difficulty of dividing the expected relationship through by income. Whereas almost all of the demographic variables have significant effects on the level of housing, only the black indicator, income squared and education have significant effects on VALUE/y . The coefficient of interest in this case is on $\text{COV}(P,y)/y$, which is equal to $\text{COV}(\ln P, \ln y)$. The estimated coefficient is negative but insignificant. In all cases, the standard errors are heteroskedasticity robust and allow for clustering at the MSA-SIC cell level. The heteroskedasticity control is important given the expected nonlinear relationship between housing purchases and covariance, and because purchases are bounded below at zero. Given the attenuation bias expected due to measurement error, I defer interpretation of magnitudes to the IV estimates.

I add MSA-SIC cell fixed effects and interactions of MSA and SIC dummies with

income to the OLS estimates in Table 1.5. The presence of cell fixed effects implies that all log variances and covariances, common to all workers in any particular cell disappear from the analysis. Further, the income \times MSA interaction absorbs income times growth and variance of income. Column (1) reports an OLS regression with no weights, and column (2) weights observations by the inverse of the estimated standard error of $\text{COV}(\ln P, \ln y)$. Column (3) replaces the standard covariance estimate with the interest rate adjusted covariance measure $\text{COV}(P, y)|R$. In all specifications, income price covariance has a negative and significant effect on VALUE, although smaller in magnitude than without fixed effects. The comparable results in columns (1) and (3) show that there is no significant difference in the coefficients on the alternative estimates of covariance.

Tables 1.6 and 1.7 present first and second stage instrumental variables regressions of the form described in equations (1.16) and (1.17). In all cases, we find that $\text{COV}(P, y)$ has a negative effect on the value of owned housing, significantly larger in magnitude than the OLS estimate. In each table, Column (1) uses the rank of the standard covariance measure, $r\text{COV}(P, y)$ as an instrument for the level of $\text{COV}(P, y)$. Column (2) uses the rank of the interest rate adjusted covariance, $r\text{COV}(P, y)|R$, and columns (3) and (4) use the rank of income times the correlation between income and prices, $r\text{INC}*\text{CORR}$. In specification (4), I also instrument for income and the income interaction with income and price variance measures, as well as for the covariance of cell mean wages with MSA total wages. This approach is indicated by the fact that income and all the elements of the variance-covariance matrix are measured with considerable error. There are separate first stage regressions for each variable marked with an 'a', with the rank of each instrumented variable providing a single instrument, so that the system of equations is exactly identified. Because I have almost 150 metropolitan area and 72 SIC interactions with income, it is not feasible to instrument separately for each of these interactions. In specifications (5) and (6) in the second stage Table 1.7, I confine analysis only to the 10 largest MSAs, and instrument for each remaining MSA-income interaction with its rank in column (6). Column (5) is presented as a comparison, and reveals that instrumenting for the

MSA-income instruments reduces the magnitude of the coefficient on $\text{COV}(P,y)$, but not quite significantly.

Given measurement error, the result that instrumenting increases the magnitude of the coefficient estimate on almost every instrumented variable is expected. Most notably, in Table 1.7 the estimated coefficient β_1 on $\text{COV}(P,y)$ increases approximately by ten to twenty times over the OLS estimates in Table 1.5. Given the significant negative coefficient on the interaction of variance of income and variance of prices $\text{VAR}(P)\text{VAR}(y)$, instrumenting with the rank of the interaction between income and variance-scaled correlation, rINC*CORR appears appropriate. This instrument is less likely to be correlated with variances and higher moments of the joint income-price distribution that might be correlated with optimal housing purchases than the level of covariance is, because of the scaling. Estimates using this instrument for $\text{COV}(P,y)$ appear in columns (4) through (6) of Table 1.7. At the cost of lost cross sectional observations, using the ten largest MSAs likely reduces any measurement error surviving the IV strategy, since mean wages are measured with larger samples in the panel data from which variance and covariance measures are obtained. In sum, the estimated coefficient on $\text{COV}(P,y)$ of -54, obtained in column (4) without instrumenting for $\text{MSA} \times \text{income}$ interactions, and in column (6) with the reduced sample, and instrumenting for these interactions seems to be a more plausible estimate than the larger coefficients estimated without instrumenting for other income and variance-covariance variables, as in columns (1) through (3). In all cases, the first stage instrument is highly significant and the R^2 statistic is large.

To interpret this coefficient, multiplying -54 by the standard deviation of $\text{COV}(P,y)$, which is 453, implies that a one standard deviation increase in income-price covariance is associated with a decrease in housing purchases of approximately \$25,000. Alternatively, holding income constant, the standard deviation of log covariance $\text{COV}(\ln P, \ln y)$ is 0.009. Multiplying a one standard deviation increase in this variable by the estimated coefficient implies that housing purchases would decrease by approximately 49 percent of a year's wages. Neither effect is large relative to the very large standard deviation of VALUE , but both clearly have economic significance.

1.4.2 Effect of Covariance on Housing Purchases Conditional on Homeownership

Table 1.8 presents OLS and IV estimates of the effect of income-price covariance on housing purchases conditional on homeownership. The dependent variable $VALUE|OWN$ is equal to the value of owner occupied housing, and the sample is confined to owner-occupiers. Given that we have the “selection” equation reported below: $OWN = F(COV(P, y), Z, \epsilon)$, we expect that holding characteristics and other variance-covariance elements Z constant, households with large covariance values who own will have positive idiosyncratic taste for owner occupied housing ϵ . In this event, absent an instrument correlated with the tenure decision to own or rent housing, but uncorrelated with taste for housing conditional on ownership, we expect estimation results to understate the effect of covariance on conditional housing purchases. It should be emphasized again that even with a sample selection mechanism, the conditional regression fails to estimate the effect of covariance on desired investment in housing. Rather, it estimates the effect on optimal investment subject to the housing investment equals housing consumption constraint. I do not see this as an inherently more interesting effect than the effect on optimal investment subject to either the investment equals consumption constraint *or* the zero investment conditional on renting constraint. The latter effect was estimated in Tables 1.4, 1.5 and 1.7. The conditional coefficient is interesting in that a significant coefficient indicates that consumption is effected directly by the portfolio consideration, not just through the portfolio (as could be consistent with an effect only on the extensive margin).

Column (1) of Table 1.8 presents the OLS estimate of the effect of $COV(P,y)$, demographics and other variance-covariance terms on $VALUE|OWN$. $COV(P,y)$ has a negative, but insignificant effect in the OLS specification. Column (2) presents the result with covariance estimated by the rank of the interaction of income with income-price correlation, $rINC*CORR$. Again, as expected, the IV approach reduces attenuation bias. I obtain an estimated effect of approximately -59 dollars per unit of covariance, statistically indistinguishable from the comparable estimate in column

(4) of Table 1.7. Hence, we expect housing purchases to fall by the same figure of approximately \$25,000 estimated on the combined margin with a one standard deviation increase in covariance.³¹

1.4.3 Effect of Covariance on Tenure Choice

Table 1.9 Presents OLS and IV estimates of the effect of covariance of shocks $COV(\ln P, \ln y)$ on the probability of ownership. I divide covariance by income in this case, because the fraction of income devoted to housing, not the level seems most likely to determine ownership. In the absence of fixed effects, unreported linear probability and probit specifications give approximately identical results, which provides some confidence in the fixed effects case. In specifications (1) and (2), I include MSA and SIC dummy variables, but not interacted with income. Column (1) is an OLS estimate, and column (2) presents IV results where $COV(\ln P, \ln y)$ is instrumented with the rank of the income-price correlation, $rCORR$. The OLS and IV coefficient estimates are both negative, but not either is significantly different from zero.

1.5 Conclusions

Because housing is the most important asset, and labor income the most important source of wealth for most households, we expect intuitively that housing decisions will incorporate the desire to hedge against income risk. Putting some theoretical structure on the question of housing choice with risky prices and income, under what look like reasonable conditions, I find that households optimally purchase less housing on both the intensive and extensive margins as the covariance between housing prices and labor income increases. This theoretical prediction is borne out empirically. I estimate that, on average, an increase of one standard deviation in covariance reduces

³¹Using income interacted with mean length in the same residence by MSA-SIC cell as an “exogenous” selection variable generates an increase in the estimated coefficient on $COV(P, y)$, but the increase is small relative to the increase generated by instrumenting to overcome measurement error. Further, a negative effect of this interaction in the conditional regression is significant, but the positive effect on selection is not.

housing investment by approximately \$25,000. It is clear that this effect operates on the intensive margin. On the extensive margin, covariance has a negative effect on the probability of ownership, but significance cannot be established. An implication is that uninsurable labor income and housing prices, combined with non-diversification of housing investment, act to distort consumption and investment decisions substantially. The results are interesting both because they extend our understanding of household financial risk and because they suggest that households are, on average, aware of these risks and take some measures to reduce risk.

Existing studies of stock market behavior present ambiguous evidence that households act on labor income - asset return covariance. The result of Heaton and Lucas (2000) leaves open the possibility that only relatively astute entrepreneurs take covariance into account. Significantly, I find in separate, unreported regressions that the result of decreasing housing purchases in income - house price covariance extends both to college graduates and household heads with no college experience. This suggests that some degree of financial sophistication extends to the broader public.³²

The theoretical and empirical results are interesting with respect to stock market behavior. Because homeowners are wealthier on average than renters, financial assets are concentrated in the hands of homeowners. On average, the incomes of these homeowners covary positively with housing prices. For homeowners considering the purchase of stock, there is thus background risk from income, from housing returns, and from the typically positive covariance of the two risks. Over long horizons, I find a positive correlation not only between stock market returns and labor income, but also between stock market returns and housing prices. The consequences for risk aversion over stock returns, and the welfare consequences of incremental investment in equities are worthy of further consideration.

Karl Case, Robert Shiller and Allan Weiss have proposed³³ that derivatives mar-

³²On the extensive margin, I find a significant effect only for college graduates, and only an insignificant effect among non-college heads. We might expect the less educated household heads to be more likely to face liquidity constraints that would complicate the covariance effect analysis. Hence, we cannot be certain that the difference on the extensive margin is driven by differences in financial sophistication.

³³As in Shiller (1993).

kets in regional housing prices might offset risk attributable to variability in capital gains on housing investment. Evidently, if households completely insured against house price risk through such markets, there would be no incentive to shift housing consumption or investment with changes in income-price covariance. While the general equilibrium welfare effects of the introduction of such markets are ambiguous, the analysis suggests that such securities, if fairly priced, would have direct benefits for many households. Indeed, given the large average correlation found between income and prices, it appears that households might wish to hold short positions in regional price indices to smooth labor income across states of nature, independent of desire to smooth capital gains.

As a practical matter, most households directly hold few or zero non-housing assets, so that complete insurance against housing risk seems highly unlikely for most of the population. Given this, and in light of the analysis presented here, proposals to remove the exemption of imputed rental services and the virtual exemption of capital gains on housing from taxation warrants further consideration. Berkovec and Fullerton (1992) emphasize the attendant implicit risk sharing in housing prices. Assuming strictly positive nominal price changes (and no offsetting reduction in income taxes), a tax at rate τ on housing capital gains would proportionately reduce the covariance between income and prices for homeowners, and should hence proportionately reduce the substantial consumption and investment distortion estimated above. Again, the general equilibrium welfare consequences are uncertain, but we might expect the presence of income-price covariance to augment the positive effects found by Berkovec and Fullerton. Heterogeneity in income-price covariances across households can be expected to complicate any such analysis.

Table 1.4: OLS Regressions of Value of Housing Owned on Demographic and Variance-Covariance Characteristics

Dep. Var	(1) VALUE	(2) VALUE	(3) VALUE/y
COV (P,y)		-9.362 (2.285)**	0.003 (0.004)
AGE	1,388.571 (133.843)**	1,424.345 (156.572)**	1.808 (1.976)
INC*AGE	0.074 (0.005)**	0.082 (0.006)**	-0.000 (0.000)
AGE ²	3.526 (1.541)*	1.703 (1.817)	-0.011 (0.022)
INC*AGE ²	-0.001 (0.000)**	-0.001 (0.000)**	0.000 (0.000)
EDUC	5,280.167 (64.455)**	5,196.537 (167.329)**	1.648 (0.641)*
INC	0.432 (0.115)**	-0.252 (0.140)	0.000 (0.001)
INC ²	-0.000 (0.000)**	-0.000 (0.000)**	0.000 (0.000)*
FAMSIZE	4,553.462 (189.623)**	3,655.356 (245.919)**	3.613 (3.697)
INC*FAMSIZE	0.009 (0.006)	0.013 (0.006)*	-0.000 (0.000)
MALE	-13,506.482 (716.595)**	-12,820.228 (879.958)**	11.064 (6.516)
INC*MALE	0.136 (0.029)**	0.212 (0.035)**	-0.000 (0.000)
BLACK	-15,517.928 (757.541)**	-13,981.445 (1,154.892)**	-12.128 (5.685)*
INC*BLACK	-0.423 (0.035)**	-0.441 (0.040)**	0.000 (0.000)
VAR(lnP)		219,145.985 (22,283.352)**	96.825 (99.274)
INC*VAR(lnP)		6.960 (0.871)**	-0.001 (0.002)
VAR(lny)		-11,323.797 (42,736.410)	156.053 (174.535)
VAR(y)		0.000 (0.000)**	-0.000 (0.000)
VAR(lnP)VAR(lny)		1,964,898.986 (1,025,995.84)	-3,590.493 (2,243.771)
VAR(P)VAR(y)		-0.001 (0.000)**	0.000 (0.000)
BETA(R,y)		50.920 (5,241.719)	-11.702 (10.047)
INC*BETA(R,y)		-0.091 (0.224)	0.000 (0.000)
BETA(S,y)		1,555.647 (1,917.837)	2.244 (5.028)
INC*BETA(S,y)		-0.104 (0.066)	-0.000 (0.000)
COV(lnP,lny)		-283,631.399 (65,720.745)**	-229.485 (264.028)
COV(lny _M ,lny)		41,981.995 (72,992.684)	830.504 (659.253)
COV(y _M ,y)		6.264 (3.263)	-0.014 (0.011)
GROW(y)		17,766.275 (8,675.355)*	9.303 (56.045)
INC*GROW(y)		0.739 (0.295)*	-0.000 (0.001)
GROW(P)		83,906.355 (9,343.784)**	24.831 (27.364)
INC*GROW(P)		0.442 (0.494)	-0.000 (0.000)
Constant	-87,030.952 (2,833.372)**	-89,877.565 (3,615.726)**	-68.260 (50.938)
Observations	308,494	301,957	301,957
R-squared	0.38	0.41	0.00

Notes: Interactions of income with marital status dummies are included, but not reported. Dependent variable is equal to the dollar value of household's housing unit if the household owner occupies, or zero if the household rents. Robust standard errors, clustered at the cell level in parentheses, * significant at 5%; ** significant at 1%. Neither cell fixed effects, nor interactions of income with MSA or SIC are included. In columns (1) and (2), the dependent variable is the dollar value of housing owned. In column (3), the dependent variable is this value divided by reported 1989 income.

Table 1.5: OLS Regressions of Value of Housing Owned on Demographic and Variance-Covariance Characteristics. Cell Fixed Effects and Income Interactions with MSA and SIC Dummies Included

Dep. Var.	(1) VALUE	(2) VALUE	(3) VALUE
COV(P,y)	-3.963 (1.124)**	-4.467 (1.750)*	
COV(P,y) R			-4.166 (1.592)**
AGE	1,125.836 (130.057)**	886.551 (198.617)**	1,125.851 (130.046)**
INC*AGE	0.108 (0.005)**	0.120 (0.007)**	0.108 (0.005)**
AGE ²	3.957 (1.496)**	7.792 (2.277)**	3.957 (1.496)**
INC*AGE ²	-0.001 (0.000)**	-0.001 (0.000)**	-0.001 (0.000)**
EDUC	5,279.918 (67.389)**	5,545.541 (108.794)**	5,280.872 (67.392)**
INC	-1.543 (0.268)**	-2.168 (0.474)**	-1.574 (0.269)**
INC ²	-0.000 (0.000)**	-0.000 (0.000)**	-0.000 (0.000)**
FAMSIZE	4,270.372 (122.249)**	4,294.929 (195.163)**	4,270.192 (122.256)**
INC*FAMSIZE	0.000 (0.005)	0.000 (0.008)	0.000 (0.005)
MALE	-20,456.749 (676.674)**	-23,334.121 (1,047.859)**	-20,451.827 (676.302)**
INC*MALE	0.502 (0.026)**	0.552 (0.040)**	0.501 (0.026)**
BLACK	-11,750.677 (750.412)**	-12,505.885 (1,206.885)**	-11,752.941 (749.274)**
INC*BLACK	-0.511 (0.033)**	-0.542 (0.053)**	-0.511 (0.033)**
VAR(y)	0.000 (0.000)**	0.000 (0.000)**	0.000 (0.000)**
VAR(P)VAR(y)	-0.001 (0.000)**	-0.002 (0.000)**	-0.001 (0.000)**
INC*BETA(R,y)	-0.114 (0.199)	0.362 (0.347)	-0.084 (0.197)
INC*BETA(S,y)	-0.021 (0.025)	-0.106 (0.076)	-0.016 (0.025)
COV(y _M ,y)	0.841 (1.319)	1.829 (2.524)	0.176 (1.298)
INC*GROW(y)	-0.329 (0.199)	0.124 (0.339)	-0.220 (0.196)
Constant	0.001 (512.510)	0.000 (816.975)	0.001 (512.500)
Observations	301,957	225,503	301,957
R-squared	0.37	0.37	0.37
Weights?	N	Y	N

Notes : Interactions of income with marital status dummies are included, but not reported. Also included are MSA × SIC cell fixed effects, and MSA and SIC × income interactions. Robust standard errors in parentheses, Column (2) weights by the inverse standard error of the log income-price covariance measure. In Column (3), the standard covariance measure is replaced by the interest rate adjusted measure.

Table 1.6: Regressions of COV(P,y) on Demographic and Other Variance-Covariance Characteristics

Dep. Var.	(1) COV(P,y)	(1) COV(P,y)	(3) COV(P,y)	(4) COV(P,y)
AGE	-1.758 (0.558)**	-1.591 (0.558)**	-2.050 (0.547)**	0.822 (1.173)
INC*AGE	0.000 (0.000)**	0.000 (0.000)*	0.000 (0.000)**	-0.000 (0.000)
AGE ²	0.019 (0.006)**	0.017 (0.006)**	0.022 (0.006)**	-0.008 (0.012)
INC*AGE ²	-0.000 (0.000)**	-0.000 (0.000)*	-0.000 (0.000)**	0.000 (0.000)
EDUC	0.146 (0.129)	-0.033 (0.131)	0.005 (0.127)	-0.148 (0.144)
INC	-0.001 (0.005)	0.001 (0.005)	-0.004 (0.005)	-0.000a (0.000)**
INC ²	0.000 (0.000)**	0.000 (0.000)**	0.000 (0.000)**	0.000 (0.000)
FAMSIZE	0.664 (0.256)**	0.683 (0.257)**	0.566 (0.248)*	0.059 (0.258)
INC*FAMSIZE	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
MALE	-14.745 (2.985)**	-12.286 (2.971)**	-15.946 (2.915)**	-6.397 (3.140)*
INC*MALE	0.001 (0.000)**	0.001 (0.000)**	0.001 (0.000)**	0.000 (0.000)*
BLACK	-1.198 (3.974)	-2.336 (3.957)	3.130 (3.920)	-0.670 (4.150)
INC*BLACK	0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)
VAR(y)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.000a (0.000)**
VAR(P)VAR(y)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.001a (0.000)**
INC*BETA(R,y)	-0.008 (0.001)**	-0.009 (0.001)**	-0.008 (0.001)**	-0.000a (0.000)**
INC*BETA(S,y)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001a (0.000)**
COV(y _M ,y)	0.396 (0.061)**	0.419 (0.061)**	0.364 (0.060)**	-0.000a (0.000)**
INC*GROW(y)	-0.029 (0.002)**	-0.031 (0.002)**	-0.028 (0.002)**	-0.022 (0.002)**
rCOV(P,y)	0.002 (0.000)**			
rINC*CORR			0.002 (0.000)**	0.003 (0.000)**
rCOV(P,y) R		0.001 (0.000)**		
Constant	-0.000 (1.750)	-0.000 (1.746)	-0.000 (1.722)	-0.000 (1.799)
Observations	301,957	301,957	301,957	301,957
R-squared	0.71	0.71	0.72	0.70

Notes: Interactions of income with marital status dummies are included, but not reported. Also included are MSA × SIC cell fixed effects, and MSA and SIC × income interactions. Robust standard errors in parentheses. (a) indicates rank instrument used for variable in question. rX is the rank of variable X. In columns (1) through (3), only COV(P,y) is instrumented. In column (4), multiple variables are instrumented.

Table 1.7: IV Regressions of Value of Housing Owned on Demographic and Variance-Covariance Characteristics

Dep. Var.	(1)	(2)	(3)	(4)	(5)	(6)
	VALUE	VALUE	VALUE	VALUE	VALUE	VALUE
COV(P,y)	-85.694 (6.106)**	-89.498 (6.710)**	-55.278 (3.713)**	-53.666 (5.977)**	-63.873 (6.038)**	-54.408 (5.649)**
AGE	1,130.429 (138.047)**	1,130.643 (138.760)**	1,128.720 (133.381)**	-1,252.534 (267.258)**	386.309 (474.508)	608.037 (429.644)
INC*AGE	0.109 (0.005)**	0.109 (0.005)**	0.108 (0.005)**	0.221 (0.011)**	0.175 (0.018)**	0.166 (0.016)**
AGE ²	3.871 (1.582)*	3.867 (1.590)*	3.903 (1.531)*	28.619 (2.978)**	16.047 (5.264)**	13.858 (4.768)**
INC*AGE ²	-0.001 (0.000)**	-0.001 (0.000)**	-0.001 (0.000)**	-0.002 (0.000)**	-0.002 (0.000)**	-0.002 (0.000)**
EDUC	5,256.410 (68.286)**	5,255.315 (68.398)**	5,265.158 (67.756)**	5,401.512 (80.328)**	6,315.823 (142.786)**	6,287.282 (126.646)**
INC	-1.107 (0.565)	-1.087 (0.583)	-1.269 (0.431)**	-16.442a (3.155)**	-13.857a (5.847)*	-11.614a (5.005)*
INC ²	-0.000 (0.000)**	-0.000 (0.000)**	-0.000 (0.000)**	-0.000 (0.000)**	-0.000 (0.000)**	-0.000 (0.000)**
FAMSIZE	4,296.228 (123.871)**	4,297.432 (124.036)**	4,286.606 (122.837)**	4,076.618 (136.408)**	5,909.859 (267.612)**	5,954.974 (234.777)**
INC*FAMSIZE	0.000 (0.005)	0.000 (0.005)	0.000 (0.005)	0.000 (0.007)	-0.153 (0.040)**	0.000 (0.014)
MALE	-20,947.863 (717.316)**	-20,970.720 (721.355)**	-20,765.098 (692.266)**	-22,907.840 (1,025.363)**	-28,197.806 (1,836.791)**	-28,233.492 (1,773.840)**
INC*MALE	0.525 (0.028)**	0.526 (0.029)**	0.516 (0.027)**	0.626 (0.042)**	0.645 (0.065)**	0.650 (0.064)**
BLACK	-12,406.507 (867.644)**	-12,437.030 (876.455)**	-12,162.444 (810.185)**	-10,767.258 (1,252.254)**	-20,086.928 (1,872.549)**	-20,866.056 (1,661.659)**
INC*BLACK	-0.484 (0.038)**	-0.483 (0.039)**	-0.494 (0.036)**	-0.551 (0.056)**	-0.466 (0.071)**	-0.441 (0.063)**
VAR(y)	0.000 (0.000)**	0.000 (0.000)**	0.000 (0.000)**	0.000a (0.000)**	0.001a (0.000)**	0.001a (0.000)**
VAR(P)VAR(y)	-0.001 (0.000)**	-0.001 (0.000)**	-0.001 (0.000)**	0.004a (0.001)**	-0.004a (0.001)**	-0.003a (0.001)*
INC*BETA(R,y)	-0.825 (0.170)**	-0.858 (0.172)**	-0.561 (0.171)**	-18.695a (2.775)**	-43.158a (3.639)**	-38.325a (3.167)**
INC*BETA(S,y)	-0.140 (0.105)	-0.146 (0.109)	-0.096 (0.071)	-4.540a (0.421)**	-5.700a (0.514)**	-6.209a (0.458)**
COV(y _M ,y)	38.324 (7.600)**	40.069 (7.996)**	24.375 (4.735)**	39.488a (7.867)**	105.693a (13.968)**	74.917a (12.812)**
INC*GROW(y)	-2.813 (0.295)**	-2.929 (0.302)**	-1.889 (0.248)**	-0.844 (0.839)	4.391 (0.729)**	4.512 (0.680)**
Constant	0.001 (532.581)	0.001 (534.440)	0.001 (520.576)	0.001 (700.656)	-98,970.825 (9,255.197)**	0.006 (2,107.177)
Observations	301,957	301,957	301,957	301,957	118,420	118,420
R-squared	0.35	0.35	0.37	0.18	0.05	0.27
Instrument for COV(P,y)	rCOV(P,y)	rCOV(P,y) R	rINC*CORR	rINC*CORR	rINC*CORR	rINC*CORR

Notes: Interactions of income with marital status dummies are included, but not reported. Also included are MSA × SIC cell fixed effects, and MSA and SIC × income interactions. Robust standard errors in parentheses. (a) indicates the variable is instrumented. All specifications instrument for COV(P,y). Columns (5) and (6) include only the 10 largest MSAs. In Column (6) MSA × income interactions are instrumented with their ranks.

Table 1.8: Regressions of Housing Value on Demographic and Variance-Covariance Characteristics: Homeowners Only

Dep. Var	(1) VALUE OWN	(2) VALUE OWN
COV(P,y)	-0.856 (0.932)	-58.684 (4.203)**
AGE	485.548 (171.067)**	483.688 (179.049)**
INC*AGE	0.063 (0.005)**	0.064 (0.005)**
AGE ²	1.597 (1.866)	1.643 (1.947)
INC*AGE ²	-0.001 (0.000)**	-0.001 (0.000)**
EDUC	6,019.023 (73.684)**	6,001.346 (74.493)**
INC	-0.924 (0.242)**	-0.770 (0.526)
INC ²	-0.000 (0.000)**	-0.000 (0.000)**
FAMSIZE	1,430.031 (124.375)**	1,457.953 (125.671)**
INC*FAMSIZE	-0.008 (0.005)	-0.008 (0.006)
MALE	-18,504.329 (820.104)**	-18,867.443 (860.551)**
INC*MALE	0.413 (0.025)**	0.430 (0.027)**
BLACK	-16,412.992 (1,075.276)**	-16,946.121 (1,227.063)**
INC*BLACK	-0.327 (0.037)**	-0.308 (0.043)**
VAR(y)	0.000 (0.000)**	0.000 (0.000)*
VAR(P)VAR(y)	-0.000 (0.000)**	-0.000 (0.000)*
INC*BETA(R,y)	-0.131 (0.193)	-0.626 (0.160)**
INC*BETA(S,y)	0.007 (0.025)	-0.103 (0.090)
COV(y _M ,y)	-1.455 (1.187)	24.334 (5.698)**
INC*GROW(y)	0.118 (0.174)	-1.599 (0.243)**
Constant	-6,705.237 (588.352)**	-6,786.954 (607.538)**
Observations	189,166	189,166
R-squared	0.30	0.28
Instrument for COV(P,y)	None	rINC*CORR

Notes: Interactions of income with marital status dummies are included, but not reported. Also included are MSA × SIC cell fixed effects, and MSA and SIC × income interactions. Robust standard errors in parentheses.

Table 1.9: Linear Probability Regressions of a Dummy Variable for Home Ownership on Demographic and Variance-Covariance Characteristics

Dep. Var	(1)	(2)
	OWN	OWN
COV(lnP,lny)	-0.214 (0.172)	-0.110 (0.258)
age	0.035 (0.001)**	0.035 (0.001)**
INC*AGE	-0.000 (0.000)**	-0.000 (0.000)**
AGE ²	-0.000 (0.000)**	-0.000 (0.000)**
INC*AGE ²	0.000 (0.000)**	0.000 (0.000)**
EDUC	0.010 (0.000)**	0.010 (0.000)**
INC	0.000 (0.000)**	0.000 (0.000)**
INC ²	-0.000 (0.000)**	-0.000 (0.000)**
FAMSIZE	0.021 (0.001)**	0.021 (0.001)**
INC*FAMSIZE	-0.000 (0.000)	-0.000 (0.000)
MALE	-0.030 (0.004)**	-0.030 (0.004)**
INC*MALE	-0.000 (0.000)	-0.000 (0.000)
BLACK	-0.134 (0.005)**	-0.134 (0.005)**
INC*BLACK	0.000 (0.000)**	0.000 (0.000)**
VAR(lnP)	-0.074 (0.097)	-0.085 (0.098)
VAR(lny)	0.038 (0.087)	0.033 (0.088)
VAR(lnP)VAR(lny)	0.699 (1.824)	0.689 (1.807)
BETA(R,y)	-0.000 (0.011)	0.000 (0.011)
BETA(S,y)	-0.006 (0.003)*	-0.006 (0.003)*
COV(lny _M ,lny)	0.090 (0.155)	0.038 (0.191)
GROW(y)	-0.018 (0.026)	-0.016 (0.026)
GROW(P)	-0.003 (0.025)	-0.003 (0.025)
Constant	-0.525 (0.031)**	-0.462 (0.033)**
Observations	301,957	301,957
R-squared	0.32	0.32
Instrument for COV(lnP,lny)	None	rCORR

Notes: Interactions of income with marital status dummies are included, but not reported. Also included are MSA and SIC fixed effects. No income \times MSA or SIC interactions are included. Robust standard errors, clustered at the cell level, in parentheses.

Chapter 2

Annuities and Individual Welfare (Joint with Jeffrey Brown and Peter Diamond)

2.1 Introduction

Annuities play a central role in the theory of a life-cycle consumer with an unknown date of death. Yaari (1965) shows that certain consumers should annuitize all savings, and evaluates the welfare gain for such consumers if they move all of their savings from unannuitized to annuitized assets. These consumers satisfy several restrictive assumptions: they are von Neumann-Morgenstern expected utility maximizers with intertemporally separable utility, they face no uncertainty other than time of death, the markets in which they trade are complete, actuarially fair annuities are available, and they have no bequest motive.¹ The subsequent literature has typically relaxed one or two of these assumptions, but has universally retained expected utility and additive separability, the latter dubbed “not a very happy assumption” by Yaari.²

¹It must be noted that Yaari does not suggest that any of these assumptions (other than absence of bequest motives) are required.

²Kotlikoff and Spivak (1981), Bernheim (1991), Hurd (1989) Jouten (2001) and Walliser (2001) consider incomplete markets and bequest motives. Brown (2001b) and Palmon and Spivak (2001) consider actuarially unfair pricing. Milevsky (2001) explores the consequences of uncertainty con-

This paper extends the theory of annuitization with no bequest motive in two directions: first, we derive sufficient conditions in a more general setting than Yaari's under which complete annuitization is optimal, and yet weaker conditions under which partial annuitization is better than zero annuitization. Second, we explore how incremental and complete annuitization affect consumer welfare in these more general conditions.

Section 2.2 considers annuitization when trade takes place all at once. In section 2.2.1, we consider a two period setting with no uncertainty other than individuals' date of death. Here, all savings are placed in annuities as long as there is no bequest motive and annuities pay a greater return, net of transaction costs, than conventional assets in the event that the consumer survives. The intuition behind this result is that consumers who do not care about wealth after death will prefer any asset which pays out more in every period of life than some other asset, regardless of how small the difference in payouts. We are thus able, in section 2.2.2 to extend the result to the Arrow-Debreu case with arbitrarily many future periods with aggregate uncertainty, as long as conventional asset and annuities markets are complete. Neither expected utility nor additive separability are required for these results.

In section 2.2.3, we provide a weaker result for the case where conventional asset markets are complete, but annuities markets are incomplete. In this case, as long as trade occurs all at once and consumers avoid zero consumption in every state of nature, some positive degree of annuitization is better than zero annuitization. All that is required for this result is that zero consumption in any state of nature is always avoided and that the annuities that exist pay greater returns than any conventional asset which pays off in all the same states of nature. Also, if trade in unannuitized assets occurs all at once, we have the result that an annuitized version of any conventional asset always dominates the underlying asset for consumers with no bequest motive. This result holds even if the asset in question does not pay off in every state of nature. An important consequence is that the result that annuities dominate conventional assets extends past riskless bonds to risky securities including

cerning asset returns.

mutual funds and certificates of deposit. For example, suppose the provider of some family of mutual funds doubles the set of available funds by offering a matching, “annuitized” fund which periodically takes the accounts of investors who die and distributes the proceeds proportionally across the accounts of surviving investors. With a large number of investors (and small additional administrative costs), the returns to this annuitized fund would strictly exceed the returns of the underlying funds for investors as long as they live. Since we assume investors do not care about wealth after death, the annuitized funds must dominate the underlying funds as an investment.

When conventional markets are incomplete, so that trade occurs more than once, it is possible that zero wealth is optimally annuitized, but only if annuities are illiquid relative to conventional assets. This possibility is discussed in section 2.3.

Given these results, the empirically observed dearth of private supply of annuitized assets is even more puzzling than previous analysis suggests.³ Taking this shortage of annuitized assets as given, it is interesting to consider the welfare consequences of moving savings from conventional assets, which allow complete choice over payout trajectories, to a single annuitized asset, which imposes a particular, and possibly unattractive payout trajectory. One common approach to measuring the benefits of annuitization is to compare expenditures with zero annuitization to expenditures with complete annuitization holding utility constant (as in Yaari (1965), Kotlikoff and Spivak (1981), Mitchell, Brown, Poterba and Warshawsky (1999), Brown (2001b)). This is an equivalent variation, the standard measure of the welfare consequences of policies with economic impact.⁴ We may also be interested in the effect on expenditures of a differential increase in annuitization: Mitchell et al. (1999) estimate a sort of marginal valuation of annuities by determining the difference between the change

³While Poterba (1997) notes that variable annuities are becoming increasingly popular, these insurance products are tax-favored savings devices that include an option to annuitize, an option which rarely appears to be activated.

⁴Mitchell et al. (1999) calculate a measure of “wealth equivalence,” which is the ratio of annuitized to non-annuitized wealth that achieves the no annuitization utility level. Other studies often calculate the ratio of non-annuitized to annuitized wealth required to achieve the full annuitization utility level, and this is sometimes referred to as the “annuity equivalent wealth.”

in expenditures from the move from zero to complete annuitization and from the move from half annuitization to complete annuitization. Brown (2001a) calculates a similar move from partial to complete annuitization for individuals in the Health and Retirement Survey.

Results on the sign of welfare effects follow from those on optimal choice of annuitization: completeness of markets is sufficient to guarantee that incremental (and hence complete) annuitization is welfare improving. With incomplete markets, the welfare effects of incremental annuity purchases are ambiguous (although we know that a move from zero to a small degree of annuitization is welfare increasing). We also analyze the demand characteristics that determine the size of welfare benefits of complete and incremental annuitization.

In section 2.4, following the literature, we apply the welfare analysis to the provision of fixed real annuities using two particular utility functions. We start with the case of a 65-year old single male with CRRA utility, and provide theoretical results on the optimal degree of annuitization as well as numerical estimates of the welfare consequences. We also relate the gain from increased annuitization to parameters of the utility function. Consistent with previous research, we find that annuitization is equivalent to a large increase in retirement wealth. We also find that the optimal fraction of savings placed in a constant real annuities is very large: 100 percent when the rate of time preference equals the interest rate (this result generalizes to all additively separable preferences), and 72 percent when the rate of time preference is large relative to the interest rate. We then examine how the results change when this same 65-year old single male has preferences that are intertemporally dependent. Specifically, we consider a utility function in which instantaneous utility is a function of the ratio of present consumption to a standard of living, which is itself derived from past consumption. We find that the value of annuitization increases with the introduction of a standard of living effect when the initial (age 65) standard of living is small relative to wealth. Conversely, for cases where the initial standard of living effect is large relative to initial retirement wealth, the welfare effects of annuitization are much smaller, and sometimes even negative. In what we regard as an

extremely bad case for annuitization however, approximately 60 percent of wealth is optimally annuitized. We also calculate the welfare benefits associated with a move from zero annuitization to complete annuitization when consumers are free to choose the trajectory of annuity payments.

To the extent that the population exhibits heterogeneity in the importance of standard of living effects, the results suggest that there may be important heterogeneity in the value placed on annuitization. However, the large fraction of savings placed in annuities for all preferences reinforces the empirical puzzle of low annuitization rates and suggests that absent bequest motives and markets for private annuitization, the current fairly large mandatory annuitization program may be social welfare improving just on the grounds of increased annuitization.

2.2 Annuitization When Trade Occurs All At Once

Analysis of intertemporal consumer choice is greatly simplified if resource allocation decisions are made all at once. Consumers will be willing to commit to a fixed plan of expenditures at the start of time under either of two conditions. The first condition, standard in the complete market Arrow-Debreu model is that, at the start of time, consumers are able to trade goods across time and all states of nature. Alternatively, first period asset trade obviates future trade across states of nature if consumers live for only two periods. In this case, absent a bequest motive, there is no reason to do anything with available wealth in the second period other than consume it.⁵

2.2.1 Two Periods, No Aggregate Uncertainty

Yaari considers annuitization in a continuous time setting where consumers are uncertain only about the time at which they will die. Some important results can be seen more simply by dividing time into two discrete periods: the present, period 1, when the consumer is definitely alive, and period 2, when the consumer is alive with

⁵With multiple consumption goods, trade would generally be optimal, but here we consider a single good.

probability $1 - q$. We maintain the assumption that there is no bequest motive, and for the moment assume that only survival to period 2 is uncertain. In this case, lifetime utility is defined over first period consumption c_1 and planned consumption in the event that the consumer is alive in period 2, c_2 . By writing

$$U = U(c_1, c_2)$$

we allow for the possibility that the effect of second-period consumption on utility depends on the level of first period consumption. This formulation does not require that preferences satisfy the axioms for U to be an expected value.

We approach both optimal decisions and the welfare evaluation of the availability of annuities by taking a dual approach. That is, we analyze consumer choice in terms of minimizing expenditures subject to attaining a minimal level of utility. We measure expenditures in units of first period consumption. Assume that there is a bond available which returns R_B units of consumption in period 2, whether the consumer is alive or not, in exchange for each unit of the consumption good in period 1. Assume in addition the availability of an annuity which returns R_A in period 2 if the consumer is alive and nothing if the consumer is not alive. Whereas the bond requires the supplier to pay R_B whether or not the saver is alive, the annuity pays out only if the saver is alive. If the annuity were actuarially fair, then we would have $R_A = \frac{R_B}{1-q}$. Adverse selection and transaction costs may drive returns below this level (the relevant definition of R_A for our purposes is net of transaction costs and inclusive of the effects of adverse selection). However, because any consumer will have a positive probability of dying between now and any future period, thereby relieving borrowers' obligation, we regard the following as a weak assumption ⁶:

Assumption 1 $R_A > R_B$

Denoting by A savings in the form of annuities, and by B savings in the form of

⁶That $R_B < R_A < \frac{R_B}{1-q}$ is supported empirically by Mitchell et al. (1999). If the first inequality were violated, annuities would be dominated by bonds.

bonds, if there is no income in period 2 (e.g. retirees), then

$$c_2 = R_A A + R_B B, \quad (2.1)$$

and expenditures for lifetime consumption are

$$E = c_1 + A + B. \quad (2.2)$$

The expenditure minimization problem can thus be defined as a choice over first period consumption and bond and annuity holdings:

$$\begin{aligned} \min_{c_1, A, B} \quad & c_1 + A + B \\ \text{s.t.} \quad & U(c_1, R_A A + R_B B) \geq \bar{U} \end{aligned} \quad (2.3)$$

By Assumption 1, purchasing annuities and selling bonds in equal numbers would cost nothing and yield positive consumption when alive in period 2 but leave a debt if dead. However, such an arbitrage would imply that lenders would be faced with losses in the event that such a trader failed to live to period 2. The standard Arrow-Debreu assumption is that planned consumption is in the consumption possibility space. For someone who is dead, this would require that the consumer not be in debt. In this simple setting the restriction is therefore that

$$B \geq 0.$$

This setup yields two important results. The first considers improving an arbitrary allocation while the second refers to the optimal plan.

Result 2 *(i) If $B > 0$, then (i) annuitization can be increased while reducing expenditures and holding the consumption vector constant. (ii) The solution to problem (2.3) sets $B = 0$.*

Proof. For (i) a sale of $\frac{R_A}{R_B}$ of the bond and purchase of 1 annuity works by Assumption 1 and definition of c_2 . For (ii), by (i), a solution with $B > 0$ fails to

minimize expenditures. Solutions with the inequality reversed are not permitted. ■

In this two period setting, Part (ii) of Result 2 is an extension of Yaari's result of complete annuitization to conditions of intertemporal dependence in utility, preferences that may not satisfy expected utility axioms and actuarially unfair annuities. All that is required is that there is no bequest motive, and that the payout of annuities dominates that of conventional assets.

Part (i) of Result 2 implies that the introduction of annuities reduces expenditures for constant utility, thereby generating increased welfare (a positive equivalent variation or a negative compensating variation). We might be interested in two related calculations: the reduction in expenditures associated with allowing consumers to annuitize a larger fraction of their savings (particularly from a level of zero), and the benefit associated with allowing consumers to annuitize all of their savings. That is, we want to know the effect on the expenditure minimization problem of loosening or removing an additional constraint on problem (2.3). To examine this issue, we restate the expenditure minimization problem with a constraint on the availability of annuities as:

$$\min_{c_1, A, B} : c_1 + A + B \tag{2.4}$$

$$s.t. : U(c_1, R_A A + R_B B) \geq \bar{U}$$

$$A \leq \bar{A}. \tag{2.5}$$

$$B \geq 0 \tag{2.6}$$

Given the linearity in both the objective function and the budget constraint, the solution to this problem must satisfy either $A = \bar{A}$ or $B = 0$ (or as a knife-edge, both).

We know that utility maximizing consumers will take advantage of an opportunity to annuitize as long as second-period consumption is positive. This is ensured by the

plausible condition that zero consumption is extremely bad:

Assumption 2

$$\lim_{c_t \rightarrow 0} \frac{\partial U}{\partial c_t} = \infty \text{ for } t = 1, 2$$

We can see from the optimization that allowing consumers previously unable to annuitize any wealth to place a small amount of their savings into annuities (increasing \bar{A} from zero) leaves second period consumption unchanged (since the cost of the marginal second-period consumption is unchanged, and so too, therefore, is the level of consumption in both periods). By Result 2, in this case, a small increase in \bar{A} generates a very small substitution of the annuity for the bond proportional to the prices

$$\begin{aligned} \frac{dA}{d\bar{A}} &= 1 \\ \frac{dB}{d\bar{A}} &= -\frac{R_A}{R_B}, \end{aligned}$$

leaving consumption unchanged: $dc_2 = R_A - R_A \frac{R_B}{R_B} = 0$.

The effect on expenditures is equal to $1 - \frac{R_A}{R_B} < 0$. This is the welfare gain from increasing the limit on available annuities for an optimizing consumer with positive bond holdings.

If constraint (2.5) is removed altogether, the price of second period consumption in units of first period consumption falls from $\frac{1}{R_B}$ to $\frac{1}{R_A}$. With a change in the cost of marginal second-period consumption, its level will adjust. Thus the cost savings is made up of two parts. One part is the savings while financing the same consumption bundle as when there is no annuitization and the second is the savings from adapting the consumption bundle to the change in prices. We can measure the welfare gain in gong from no annuities to unlimited annuities by integrating the derivative of the expenditure function between the two prices:

$$E|_{\bar{A}=0} - E|_{\bar{A}=\infty} = - \int_{R_A^{-1}}^{R_B^{-1}} c_2(p_2) dp_2, \quad (2.7)$$

where c_2 is compensated demand arising from minimization of expenditures equal to $c_1 + c_2 p_2$ subject to the utility constraint without a distinction between asset types.

Equation (2.7) implies that consumers who save more (have larger second-period consumption) benefit more from the ability to annuitize completely:

Result 3 *The benefit of allowing complete annuitization (rather than no annuitization) is greater for consumer i than for consumer j if consumer i 's compensated demand for second period consumption (equivalently, compensated savings) exceeds consumer j 's for any price of second period consumption.*

2.2.2 Many Future Periods and States, Complete Markets

The result of the optimality of complete annuitization survives subdivision of the aggregated future defined by c_2 into many future periods and states. A particularly simple subdivision would be to add a third period, so that survival to period 2 occurs with probability $1 - q_2 = 1 - m_2$ and to period 3 with probability $(1 - q_2)(1 - q_3) = 1 - m_3$. In this case, bonds and annuities which pay out separately in period 2 with rates R_{B2} and R_{A2} , and period 3 with rates R_{B3} and R_{A3} are sufficient to obviate trade in periods 2 or 3. That is, defining bonds and annuities purchased in period 1 with the appropriate subscript⁷,

$$E = c_1 + A_2 + A_3 + B_2 + B_3$$

$$c_2 = R_{B2}B_2 + R_{A2}A_2,$$

$$c_3 = R_{B3}B_3 + R_{A3}A_3.$$

If Assumption 1 is modified to hold period by period, Result 2 extends trivially. Note that we have set up what we will call “Arrow bonds” by combining two states of nature that differ in no other way except whether this consumer is alive. “Arrow

⁷Implicitly, we are assuming that if markets reopened, the relative prices would be the same as are available in the initial trading period

annuities” which also recognize whether this consumer is alive are the true Arrow securities of standard theory.

In order to take the next logical step, we can continue to treat c_1 as a scalar, and interpret c_2 , B_2 , and A_2 as vectors with entries corresponding to arbitrarily many (possibly infinity) future periods ($t \leq T$), as well as arbitrarily many states of nature ($\omega \leq \Omega$) in each period. R_{A_2} and R_{B_2} are then $T\Omega \times T\Omega$ matrices with columns corresponding to annuities (bonds) and rows corresponding to payouts by state of nature. Thus, the assumption of no aggregate uncertainty can be dropped. Multiple states within each period might refer to uncertainty about aggregate issues such as output, or individual specific issues beyond mortality such as health.⁸ In order to extend the analysis, we need to assume that the consumer is sufficiently “small” that for each state of nature where the consumer is alive, there exists a state where the consumer is dead and the allocation is otherwise identical. Completeness of markets still allows construction of Arrow bonds which represent the combination of two Arrow securities. We assume that the bonds and annuities have payouts $R_{B_{t\omega}}$ and $R_{A_{t\omega}}$.

Annuities with payoffs in only one event state are contrary to our conventional perception of (and name for) annuities as paying out in every year until death. However, with complete markets, separate annuities with payouts in each year can be combined to create such securities. It is clear that the analysis of the two-period model extends to this setting, provided we maintain the standard Arrow-Debreu assumptions that do not allow an individual to die in debt. In addition to the description of the optimum, the formula for the gain from allowing more annuitization holds for state-by-state increases in the level of allowed level of annuitization. Moreover, by choosing any particular price path from the prices inherent in bonds to the prices inherent in annuities, we can measure the gain in going from no annuitization to full annuitization. This parallels the evaluation of the price changes brought about by a lumpy investment (see Diamond and McFadden (1974)).

Thus we have extended the Yaari result of complete annuitization to conditions

⁸For a discussion of annuity payments that are partially dependent on health status, see Warshawsky, Spillman and Murtaugh (Forthcoming).

of aggregate uncertainty, actuarially unfair (but positive) annuity premiums and intertemporally dependent utility that need not satisfy the expected utility axioms. Moreover, we have the result that increasing the extent of available annuitization increases welfare if there is positive holding of Arrow bonds.⁹

2.2.3 More than Two Periods, Complete Bond Markets, Incomplete Annuities Markets, Trading only Once

Annuities are frequently assumed to require a particular time path of payouts, thereby combining in a single security a combination of Arrow securities. Such annuities, for example with a constant real payout, have been evaluated in the literature. The United States Social Security system works this way (ignoring the role of the earnings test). Private annuities have generally been fixed in nominal terms rather than real. Variable annuities make the payoff depend on returns on some given portfolio, and combine Arrow securities in that way. CREF annuities also vary the payoff with mortality experience for the class of investors, which is also a combination of Arrow securities (see Poterba (1997)). To consider such lifetime annuities in this setup, we continue to assume a double set of states of nature, differing only in whether the particular consumer we are analyzing is alive. We continue to assume a complete set of Arrow bonds and consider the effect of the availability of particular types of annuities. We also need to consider whether the return from annuities and bonds can be reinvested (markets are open) or must be consumed (markets are closed) In general, we will lose the result that complete annuitization is optimal. Nevertheless, we will get a similar result for real annuities provided that optimal consumption is rising over time and markets are open. In addition we will examine a sufficient conditions for the result that the optimal holding of annuities is not zero.

To illustrate these points, we consider a three-period model with no aggregate

⁹The generalization of Result 3 to this case requires the very strong condition that after the present, consumption for agent i exceeds that of agent j state of nature by state of nature. That i 's consumption grows at a greater rate than j 's is not sufficient: allowing complete annuitization may yield reduction in many different prices by increasing any of many ratios $\frac{A_{t\omega}}{A_{t\omega} + B_{t\omega}}$. In general, these price changes are non-monotonic in time past period 1.

uncertainty and a complete set of bonds. Then we will show how the results generalize. If there are no annuities, then the expenditure minimization problem is:

$$\begin{aligned} \min_{c_1, A, B} & : c_1 + B_2 + B_3 & (2.8) \\ \text{s.t.} & : U(c_1, R_{B_2}B_2, R_{B_3}B_3) \geq \bar{U} \end{aligned}$$

That is, we have:

$$\begin{aligned} c_2 &= R_{B_2}B_2, \\ c_3 &= R_{B_3}B_3. \end{aligned}$$

With the assumption of infinite marginal utility at zero consumption, all three of c_1 , B_2 , and B_3 are positive. Now assume that there is a single available annuity, A , that pays given amounts in the two periods. Assume further that there is no opportunity for trade after the initial contracting. The minimization problem is now

$$\begin{aligned} \min_{c_1, A, B} & : c_1 + B_2 + B_3 + A & (2.9) \\ \text{s.t.} & : U(c_1, R_{B_2}B_2 + R_{A_2}A, R_{B_3}B_3 + R_{A_3}A) \geq \bar{U} \end{aligned}$$

That is, we have:

$$\begin{aligned} c_2 &= R_{B_2}B_2 + R_{A_2}A, \\ c_3 &= R_{B_3}B_3 + R_{A_3}A. \end{aligned}$$

Before proceeding, we must revise the assumption that $R_{At\omega} > R_{Bt\omega} : \forall t\omega$. A more appropriate formulation for the return on a complex security that combines Arrow securities to exceed bond returns is that for any quantity of the payout stream provided by the annuity, the cost is less if bought with the annuity than if bought through bonds. Define by ℓ a row vector of ones with length $T\Omega$, let the set of bonds be represented by a $T\Omega \times 1$ vector with elements corresponding to the columns of the $T\Omega \times T\Omega$ matrix of returns R_B , and let R_A be a $T\Omega \times 1$ vector of annuity payouts multiplying the scalar A to define state-by-state payouts.

Assumption 3 For any annuitized asset A and any collection of conventional assets B , $R_A A = R_B B \Rightarrow A < \ell B$.

For example, if there is an annuity that pays R_{A2} per unit of annuity in the second period and R_{A3} per unit of annuity in the third period, then we would have $1 < \left(\frac{R_{A2}}{R_{B2}} + \frac{R_{A3}}{R_{B3}}\right)$. By linearity of expenditures, this implies that any consumption vector that may be purchased strictly through annuities is less expensive when financed strictly through annuities than when purchased by a set of bonds with matching payoffs.¹⁰

Given the return assumption and the presence of positive consumption in all periods, it is clear that the cost goes down from the introduction of the first small amount of annuity, which can always be done without changing consumption. Thus we can also conclude that the optimum (including the constraint of not dying in debt) always includes some annuity purchase. It is also clear that full annuitization may not be optimal if the implied consumption pattern with complete annuitization is worth changing by purchasing a bond. That is, optimizing first period consumption given full annuitization, we would have the first order condition

$$U_1(c_1, R_{A2}A, R_{A3}A) = R_{A2}U_2(c_1, R_{A2}A, R_{A3}A) + R_{A3}U_3(c_1, R_{A2}A, R_{A3}A)$$

Purchasing a bond would be worthwhile if we satisfy either of the conditions:

$$U_1(c_1, R_{A2}A, R_{A3}A) < R_{B2}U_2(c_1, R_{A2}A, R_{A3}A) \quad (2.10)$$

or

$$U_1(c_1, R_{A2}A, R_{A3}A) < R_{B3}U_3(c_1, R_{A2}A, R_{A3}A) \quad (2.11)$$

By our return assumption, we can not satisfy both of these conditions, but we might satisfy one of them. That is, the optimum will involve holding some of the annuitized asset and may involve some bonds, but not all of them.

¹⁰This assumption leaves open the possibility considered below that both bond and annuity markets are incomplete and some consumption plans can be financed only through annuities.

It is clear that these results generalize to a setting with complete Arrow bonds and some compound Arrow annuities with many periods and many states of nature. We show below that expenditure minimization requires that there must be positive purchases of at least one annuity.

Lemma 1 *Consider an asset A^* with finite, non-negative payouts R_{A^*} . Any consumption plan $[c_1 \ c_2]'$ with positive consumption in every state of nature can be financed by a combination of first period consumption, a positive holding of A^* , and another strictly non-negative consumption plan.*

Proof. Define $\bar{R}_{A^*} = [\frac{1}{R_{A^*21}}, \dots, \frac{1}{R_{A^*t\omega}}, \dots, \frac{1}{R_{A^*T\Omega}}]'$, and define the scalar $\alpha = \min(c_2 \cdot \bar{R}_{A^*})$. Now $c_2 = R_{A^*}\alpha + Z$, where Z is weakly positive. ■

We now have a weaker version of Result 2:

Result 4 *If Assumption 2 holds, and there exist annuities with non-negative payouts which satisfy Assumption 3, then (i) when no annuities are held, a small increase in annuitization reduces expenditures, holding utility constant. Also, then (ii) expenditure minimization implies $A > 0$.*

Proof. Suppose that the optimal plan (c_1, A, B) features $A = 0$. Then there are two possibilities: first, consumption might be zero in some future state of nature. By Assumption 2 this implies infinitely negative utility and fails to satisfy the utility constraint. If consumption is positive in every state of nature, then consumption is a linear combination of all strictly positive linear combinations of the Arrow bonds. But then since some strictly positive consumption plan can be financed by annuities, by Assumption 3 and Lemma 1, expenditures can be reduced holding consumption constant by a trade of some linear combination of the bonds for some combination of annuities with strictly positive payouts. This contradicts optimality of the proposed solution. ■

Part (i) of Result 4 states that if consumers are willing to commit to lifetime expenditures all at once, then starting from a position of zero annuitization, a small purchase of *any* annuity increases welfare. This applies to any annuity with returns

in excess of the underlying asset, no matter how distasteful the payout stream. Part (ii) is the corollary that optimal annuity holdings are always positive. Lemma 1 shows that up to some point, annuity purchases do not distort consumption, so that their only effect is to reduce expenditures, as in the case where annuities markets are complete. When a large fraction of savings is annuitized, if the supply of annuitized assets fails to match demand, annuitization distorts consumption and conventional assets may be preferred. From the proof of Result 4, it follows that the annuitized version of any conventional asset that might be part of an optimal portfolio dominates the underlying asset.

2.3 Annuitization with Trade in Many Periods

2.3.1 More than Two Periods, Complete Liquid Bond Markets, Incomplete Illiquid Annuities Markets, Trading More than Once

The setup so far has not allowed a second period of trade. Now assume that trade in bonds is allowed after the first period, with bond prices consistent with the returns that were present for trade before the first period. To begin we assume that there is not an annuity available at the second trading time and that the consumer can save the second period annuity payout, but can not sell the remaining portion of the annuity. Since there would be no further trade without an annuity purchase at the start, the optimum without any annuity is unchanged. Utility at the optimum, assuming some annuity purchase, is at least as large as it was without the further trading opportunity. Thus we conclude that the result that some annuity purchase is optimal (Result 4) carries over to the setting with complete bond markets at the start and further trading opportunities in bonds that involve no change in the terms of bond transactions.

Returning to the three period example with no aggregate uncertainty, a sufficient condition for complete annuitization at the start, even if one of the inequalities

(2.10) or (2.11) is violated, is that the consumption stream associated with full annuity purchase at the first trading point was such that saving (rather than dissaving) was attractive. To examine this issue, we now set up the expenditure minimization problem with retrading, denoting saving at the end of the first period by Z .

$$\min_{c_1, A, B} : c_1 + B_2 + B_3 + A \quad (2.12)$$

$$s.t. : U(c_1, R_{B_2}B_2 + R_{A_2}A - Z, R_{B_3}B_3 + R_{A_3}A + (R_{B_3}/R_{B_2})Z) \geq \bar{U}.$$

The restriction of not dying in debt is the nonnegativity of consumption if A is set equal to zero ¹¹

$$B_2, B_3, Z \geq 0$$

$$R_{B_2}B_2 \geq 0$$

$$R_{B_3}B_3 + (R_{B_3}/R_{B_2})Z \geq 0$$

The assumption that dissaving would not be attractive given full annuitization is

$$R_{B_2}U_2(c_1, R_{A_2}A, R_{A_3}A) \leq R_{B_3}U_3(c_1, R_{A_2}A, R_{A_3}A) \quad (2.13)$$

This condition can be readily satisfied for preferences satisfying a suitable relationship between (implicit) utility discount rates and interest rates. The result extends with many future periods, as long as trade is allowed in each. However, once we introduce uncertainty, the sufficient condition along these lines would need to hold in every state of nature.

Absent uncertainty, the presence of future opportunities to purchase annuities would increase the set of circumstances where full initial annuitization is optimal.

¹¹ B_3 can be negative if Z is positive. However, a budget-neutral reduction in Z and increase in B_3 , holding A constant, then yields equivalent consumption, so there is no restriction in disallowing negative B_3 . If B_3 is non-negative, then Z must be zero as long as B_2 is positive, or else constant consumption with reduced expenditures could be obtained at a lower price by reducing B_2 and increasing A . That is, there are no savings out of bonds.

With stochastic availability of more favorable annuity returns, annuitization may become relatively unattractive, as in Milevsky (2001).

2.3.2 Incomplete Markets With Illiquid Annuities and Future Trade

The saving condition (2.13) required for complete annuitization becomes implausible once it is recognized that people receive information over time about life expectancy and that such states (like individual life and death) are not distinguished by existing Arrow bonds. A similar issue would arise with other noninsurable events.

Indeed, with incomplete markets, even absent annuities, future trade typically occurs. In this case, the value of an asset is not equivalent to the sum of scheduled payouts multiplied by expected marginal utility in the period of payouts, as with complete markets. Now, assets' values also depend on their trading value in all future periods. If annuitized assets were as liquid as their underlying conventional asset, then it would remain the case that the annuitized version always dominates the underlying version. As a practical matter, the payment of survivorship premia for annuitized assets will require a minimal holding period and a possible penalty for resale, so that the annuitized asset is not as liquid as the underlying asset. As a consequence, it is possible that zero annuitization may be optimal, even under the assumption that annuities provide the cheapest means of purchasing any cashflow they can reproduce (this is Assumption 3, above).¹²

To see this possibility, returning to the three period case with future trade and no aggregate uncertainty, suppose that in period 1, a consumer expects to survive to period 2 with probability $1 - m_2$ and to period 3 with probability $1 - m_3$. However, the consumer knows that in period 2, the conditional probability of survival to period 3 will be updated to zero with probability α or $\frac{1}{1-\alpha} \frac{1-m_3}{1-m_2}$ with probability $1 - \alpha$. In the likely case that neither the annuity nor the bonds distinguish between the two health

¹²The result of partial annuitization survives if there are incomplete markets with no future trade, which technically occurs if people live for only two periods.

conditions, the consumer will sell whatever bonds pay off in period 3 on obtaining bad health news, but will be unable to cash out the illiquid third period annuity claim.

Without annuitization, consumption in period two is thus given by $R_{B2}B_2$ if there is good health news, and $R_{B2}B_2 + \frac{R_{B2}}{R_{B3}}R_{B3}B_3$ if the health news is bad. Assuming the consumer is an expected utility maximizer, zero annuity purchase is thus optimal as long as:

$$\alpha R_{A2}U_2(c_1, R_{B2}(B_2+B_3)) + (1-\alpha)(R_{A2}U_2(c_1, R_{B2}B_2, R_{B3}B_3) + R_{A3}U_3(c_1, R_{B2}B_2, R_{B3}B_3)) \leq R_{B2}(\alpha U_2(c_1, R_{B2}(B_2 + B_3)) + (1 - \alpha)U_2(c_1, R_{B2}B_2, R_{B3}B_3)).$$

In this case, the superior return condition for annuities is $R_{A2} + R_{A3} > R_{B2} + R_{B3}$. The zero annuitization condition above can be consistent with this relationship if the annuities' payouts are sufficiently graded towards future payouts relative to the bonds. Hence, with incomplete markets, zero annuitization, partial annuitization, and complete annuitization are all consistent with utility maximization without further assumptions.

2.4 Valuation of a constant real annuity in special cases

Here, we consider a world with $T - 1$ future periods and no uncertainty except individual mortality, so that future consumption conditional on survival can be described by a vector with one element for each period up to T where death occurs for certain: $c_2 = [c_2, c_3 \dots c_T]'$. We consider the welfare consequences of a policy whereby consumers are forced to purchase some quantity of an actuarially fair annuity which pays out a constant sum R_A in every future period, when annuities are otherwise not available. That is, consumers minimize expenditures in a world with "Arrow" bonds and no annuity products, having already made an irreversible expenditure of A units and a commitment to take in $R_A A$ per period, where A is the amount of required

annuity spending. We assume that no annuities are available after the first period, but that future bond trades are allowed. By completeness of bond markets, we can consider the set of bonds to be described by $T - 1$ securities, each of which pays out at a rate of $(1 + r)^{t-1}$ at date t only.

With a constant real interest rate of r , without the annuity, expenditures are given by

$$E(c, 0) = c_1 + \sum_{t=2}^T c_t R_{Bt}^{-1} = c_1 + \sum_{t=2}^T c_t (1 + r)^{1-t}. \quad (2.14)$$

With annuities, the cost of a consumption plan is equal to the cost of annuitized consumption plus the difference between annuitized consumption and actual consumption in every period:

$$E(c, A) = c_1 + A + \sum_{t=1}^T (c_t - R_A A) (1 + r)^{1-t},$$

where R_A is the per-period annuity payout. For $t > 1$, if consumption is less than the annuity payout, the difference can be used to purchase consumption at later dates, with the relative prices given by bond returns. If consumption is greater than the annuity payout, then a bond with maturity at date t must be purchased.

If $1 - m_t$ is the probability of survival to period t , then actuarial fairness implies that the cost per unit of the annuity is equal to the survival-adjusted present discounted value of bond purchases yielding the same unit per period:

$$\begin{aligned} 1 &= \sum_{t=2}^T (1 - m_t) \frac{R_A}{(1 + r)^{t-1}} \\ \Rightarrow R_A &= \frac{1}{\sum_{t=2}^T (1 - m_t) (1 + r)^{1-t}}. \end{aligned} \quad (2.15)$$

Assumption 3 applies as long as there is a positive probability of death by the end of T periods ($1 - m_T < 1$) because the cost of consuming any plan $R_A A$ per period past period 1 with annuities is $\frac{A}{\sum_{t=2}^T (1 - m_t) (1 + r)^{1-t}}$ which is less than $\frac{A}{\sum_{t=2}^T (1 + r)^{1-t}}$, the cost of purchasing A per period with conventional securities.

As discussed above, a small increase in A from zero has no effect on consumption,

so that the CV from incremental annuitization from 0 to a small number ϵ is equal to the difference between $E(c, 0)$ and $E(c, \epsilon)$:

$$\frac{dE}{dA}|_{A=0} = 1 - \sum_{t=2}^T (R_A(1+r)^{1-t}) < 0. \quad (2.16)$$

The inequality follows from equations (2.14) and (2.15) as long as $m_T > 0$.

Larger increases in annuitization are more difficult to sign because they may constrain consumption. Below, we consider the effects for particular utility functions.

2.4.1 Additively Separable Preferences, Constant Real Annuity Constraint, Actuarially Fair Annuities, No Uncertainty Other Than Length of Life

Optimal Choice of Annuitization

Here, we assume that utility is given by:

$$U(c_1, c_2) = \sum_{t=1}^T \delta^{t-1} (1 - m_t) u(c_t), \quad (2.17)$$

Where $u' > 0$, $u'' < 0$; $\lim_{c_t \rightarrow 0} u' = \infty$, and δ is the rate of time preference.

Because Assumptions 2 (infinite disutility from zero consumption in any future period) and 3 (any consumption plan that can be financed by annuities alone is financed most cheaply by annuities alone) are met:

Claim 1 *The solution to the expenditure minimization problem features $A > 0$.*

Proof. Follows immediately from Result 4. ■

By the no bankruptcy constraint, consumers may undo annuitization by saving if annuitization renders consumption too weighted towards early periods, but not by borrowing if annuitization renders consumption too weighted to later periods. The liquidity constraint given a constant real annuity requires that expenditures on consumption up to any date τ must be less than total planned expenditures less

expenditures committed to future annuity payments. This constraint can be written as:

$$\sum_{t=1}^{\tau} c_t(1+r)^{1-t} \leq c_1 + A + \sum_{t=2}^T B_t - R_A A \sum_{t=\tau+1}^T (1+r)^{\tau-t} \forall \tau. \quad (2.18)$$

This induces one constraint for every period in which consumption is bound from above by the required annuity. Annuities are costly in optimization terms because they contribute to these constraints.

The expenditure minimization problem becomes:

$$\min_{c_1, A, B} c_1 + A + B \quad (2.19)$$

$$s.t. U(c_1, c_2(A, B)) \geq \bar{U}$$

s.t. equation (2.18) is satisfied.

Under these circumstances, consumers whose discount rates are no greater than the interest rate annuitize fully:

Claim 2 *If optimal consumption is weakly increasing, then complete initial annuitization is optimal.*

Proof. (1) *if consumption is increasing, equation (2.15) implies that $\exists t : R_{Bt} < R_A$. Hence, if net bond holdings are greater than zero, expenditures can be reduced and utility increased by an additional purchase of ϵ units of A and sale of $\epsilon \frac{R_A}{R_{B2}} > \epsilon$ units of B_2 . This trade does not violate (2.18) for ϵ sufficiently small because by increasing consumption, $\exists t > 2 : B_t > 0$. (2) If consumption is constant, then satisfaction of Assumption 3 implies that expenditure minimization occurs when $A = E$. ■*

Claim 3 *If the optimal level of annuitization A is less than savings, so that there are positive expenditures on bonds, an increase in δ yields an increase in optimal A relative to savings.*

Proof. *A sufficient condition for this result is that for any s such that $B_s > 0$, an increase in δ renders marginal annuity purchases strictly preferable to the bond*

purchase. At an optimum, there is positive annuitization by Claim 1. It cannot be optimal to hold bonds with maturity at date T . Hence, there is a date $S < T$ which is the latest date for which bonds are held. For any date $s \leq S$ at which bonds are held, the consumer must be indifferent to a trade of a small unit of the bond maturing at s and an equally small unit of the annuity. Further, by construction of the utility function, with an increase in δ , B_S must become relatively more attractive than any B_s for $s < S$, so we need only to show the result for B_S relative to the annuity. The condition for indifference between B_S and the annuity at an optimum is:

$$\delta^{S-1}(1+r)^{S-1}(1-m_S)u'(c_S) = R_A \sum_{t=2}^T \delta^{t-1}u'(c_t)(1-m_t).$$

The annuity places no constraints on consumption up to period S by positivity of B_S . Hence, the consumer is indifferent between a marginal dollar received at any period $s < S$ and receiving $(1+r)^{S-s}$ dollars at period S , and the indifference equation can be rewritten:

$$\delta^{S-1}(1+r)^{S-1}(1-m_S)u'(c_S)(1 - R_A \sum_{t=2}^S (1+r)^{S-t}) = R_A \sum_{t=S+1}^T \delta^{t-1}u'(c_t)(1-m_t).$$

With an increase in δ , the left hand side (LHS) increases by:

$$\frac{\partial LHS}{\partial \delta} = \frac{1}{\delta}(S-1)LHS.$$

By contrast, the right hand side increases by

$$\frac{\partial RHS}{\partial \delta} = \frac{1}{\delta}R_A \sum_{t=S+1}^S (t-1)\delta^{t-1}u'(c_t)(1-m_t) > \frac{1}{\delta}(S-1)RHS.$$

Hence, the marginal utility of annuity purchase exceeds the marginal utility of any bond, so the annuity's share of savings must increase. ■

Claim 4 If $\delta(1+r) \geq 1$, complete initial annuitization is optimal.

Proof. By Claim 3, it is sufficient to show that this is true for $\delta(1+r) = 1$. For complete annuitization to be suboptimal, it must be the case that there exists some t

for which purchasing a bond with maturity at date t provides greater marginal utility than purchase of the real annuity, or:

$$\begin{aligned} \exists t > 1 : \delta^{t-1}(1+r)^{t-1}u'(R_AA)(1-m_t) &> \frac{\sum_{t=2}^T \delta^{t-1}(1-m_t)u'(R_AA)}{\sum_{t=2}^T (1-m_t)(1+r)^{1-t}} \\ \Rightarrow \delta^{t-1}(1+r)^{t-1}(1-m_t) &> \frac{\sum_{t=2}^T \delta^{t-1}(1-m_t)}{\sum_{t=2}^T (1-m_t)(1+r)^{1-t}} \end{aligned}$$

If $\delta(1+r) = 1$, then this is impossible, because the left hand side is less than or equal to one (by non-negative mortality) and the right hand side equals one in the simplified equation. ■

The “dual” to Claim 4 follows from the proof:

Claim 5 *If $\delta(1+r) \geq 1$, then any increase in annuitization in the range $A \in [0, E-c_1]$ is welfare enhancing.*

For more impatient consumers, we solve for the optimal fraction of savings put into annuities numerically. Results are detailed below.

Beyond the results we have above, making statements about the size of EV for a move from complete annuitization to zero annuitization is difficult, because in general, this calculation must take into the period-by-period positive wealth constraints summarized in equation (2.18). That said, a plausible conjecture, based on Claim 3 is that valuation will increase in the patience parameter δ , which should push consumption later in life. Further, in cases where optimal consumption is decreasing over time, increased smoothing should increase valuation. Hence, for $\delta(1+r) \leq 1$, we should expect valuation to increase with any parameter of risk aversion, because the desire for decreasing consumption, which makes the μ constraints bind, would then be tempered by a desire for consumption smoothing. We confirm these intuitions below with numerical examples.

2.4.2 Utility Dependent on Standard of Living

Additive separability of utility does not sit well with intuition. For example, life in a studio apartment with no car is surely more tolerable for someone used to living

in a studio apartment without a car than for someone who was forced by a negative income shock to abandon a four bedroom house and a Lexus for a studio apartment and no car. In this section, we revisit the analysis above with an extreme, and hence illustrative, example of intertemporal dependence in the utility function, taken from Diamond and Mirrlees (2000). The intuition behind this formulation is that it is not the level of present consumption, but the level relative to past consumption that matters. We consider the ratio of present to past consumption, but the difference has also been considered in the literature. In choosing how to allocate resources across periods, consumers with such utility trade off immediate gratification from consumption not only against a lifetime budget constraint, but also against the effects of consumption early in life on the standard of living later in life.

$$U(c_1, c_2) = \sum_{t=1}^T \delta^{t-1} (1 - m_t) u\left(\frac{c_t}{s_t}\right), \quad (2.20)$$

where

$$s_t = \frac{s_{t-1} + \alpha c_{t-1}}{1 + \alpha}.$$

Note that if $\alpha = 0$, so that individuals have no control over their standard of living, we are in the additively separable case. A positive value of α indicates that past consumption makes individuals less satisfied with a given level of present consumption.

In the absence of the positive wealth constraints (2.18), the marginal utility of consumption in any period incorporates two effects not present in the additively separable case: (1) the effect of the present standard of living on present marginal utility and (2) the effect of present consumption on future periods' utility through subsequent standards of living. Under this specification, the marginal benefit of present consumption is given by:

$$\frac{\partial U}{\partial c_t} = \frac{1}{s} u'\left(\frac{c_t}{s_t}\right) - \sum_{k>t} \frac{\alpha}{(1 + \alpha)^{k-t}} \frac{c_k}{s_k^2} u'\left(\frac{c_k}{s_k}\right).$$

We note that if $\lim_{\frac{c_t}{s_t} \rightarrow 0} u'(c_t) = \infty$, then Assumption 2 holds, and Claim 1 applies for finite s_1 .

To obtain results, we assume that $u(\frac{c}{s}) = \frac{(\frac{c}{s})^{1-\gamma}}{1-\gamma}$ and that $\gamma \geq 1$. Hence:

$$\frac{\partial U}{\partial c_t} = c_t^{-\gamma} s_t^{\gamma-1} - \sum_{k>t} \frac{\alpha}{(1+\alpha)^{k-t}} c_k^{1-\gamma} s_k^{\gamma-2}.$$

For $\gamma > 1$, effect (1) will tend to push consumption towards later periods relative to the no standard ($\alpha = 0$) case if the standard of living is increasing over time. If the standard of living is decreasing over time, and $\gamma \geq 2$, then this will push consumption to earlier periods. For $\gamma < 2$, the effect is ambiguous.

Effect (2) will unambiguously push consumption towards later periods in life. Hence, the result of complete annuitization when the discount rate is less than the interest rate, Claim 4, continues to hold if s is constant or decreasing over the period of annuitization. This occurs if the initial value of s is small and the required level of utility, \bar{U} , is large. If the initial value s_1 is sufficiently large relative to the expenditures required to attain \bar{U} , then the smoothing implied by risk aversion may undo the result by rendering optimal consumption relatively decreasing over time.

With the constraint that the only annuity available pays out a constant real sum, relative valuations are particularly difficult to calculate with standard of living effects, because the intertemporal effects compound the difficulty of the multiple positive wealth constraints. However, we can conjecture that parameter changes that tend to defer optimal consumption will tend to increase valuation. Hence, simulated valuations should tend to be increasing in δ . Further, large s_1 should yield decreasing valuation, and small s_1 increasing valuation, with both effects magnified by γ .

2.4.3 Numerically Estimated Magnitudes of Welfare Effects

To estimate numerical valuations of annuitization, we specify $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$ for both the additively separable and standard of living exponential discounting, flat yield curve cases considered above. In the separable case, this gives constant relative risk aversion and an intertemporal rate of substitution $\frac{\partial U/\partial c_t}{\partial U/\partial c_k} = (\frac{c_t}{c_k})^{-\gamma}$. In the standard of living case, both risk aversion and intertemporal substitution are complicated by the intertemporal utility linkage.

We calculate five values: first, the CV for a mandatory purchase of a very small quantity of a constant real annuity when $A = 0$. This value is identical and negative for all consumers, as shown above. Second, and of more interest for determining parameter effects, we calculate EV for a move from putting all expenditures in such an annuity to putting nothing in the annuity (EV_R). Third, we determine the optimal fraction of savings A^* placed in the real annuity as opposed to in bonds. Fourth, we calculate the EV of a move from complete annuitization with complete annuity markets to zero annuitization (EV_C). This number will be greater than the EV associated with complete real annuitization, or equal in the knife-edge case where optimal consumption is constant with actuarially fair prices (in the additively separable case, this occurs when $\delta(1+r) = 1$). Finally, we compute an intermediate equivalent variation, EV_O : that associated with a move from annuitizing the optimal fraction A^* of wealth to zero annuitization.

We perform these calculations for a single 65 year old male in 1999 with survival probabilities taken from the US Social Security mortality tables, modified (to ease computation) so that death occurs for sure at age 100. We use a real interest rate r of 0.03, and vary δ . We consider coefficients of relative risk aversion γ of 1 (log utility) and 2, which are on the low end of plausible values. We use the same values for the standard of living case (where these cannot be interpreted as coefficients of relative risk aversion). For the case where the discount rate is 1.03^{-1} and there is no standard of living, our results are very close to those found in the existing literature despite the truncation of life by 10 or 15 years.

We “reverse engineer” the expenditure minimization problem to have minimal expenditures of 100 with complete annuitization, and exclude consumption in the year of retirement (so utility is defined only over the 35 element vector c_2). Changes in parameters will change not only the optimal allocation of savings, but also the level of savings. Comparative static analysis of the effect of parameter changes on the welfare consequences for an individual with constant lifetime wealth would thus require solution of a lifetime problem, not just a retirement consumption problem. Adding consumption at age 65 does not fix this problem.

The graphs depict and summarize results for each of 9 cases. Each graph plots optimal consumption with and without annuitization. A positive EV of annuitization indicates that the level of utility achieved with optimal consumption subject to complete actuarially fair annuitization and real annuity constraint (2.18) when 100 units of the annuity are purchased requires a greater level of expenditures to attain when consumption is financed through conventional bonds with no trajectory constraint. A rough estimate of the magnitude of EV can be obtained by observing the difference in trajectories between the two consumption plans: when optimal consumption is sharply decreasing, the constraints implied by (2.18) bind consumption away from the optimal path; in these cases the price benefit of annuitization is largely offset by the constraints. When optimal consumption is hump shaped, and less steeply decreasing, the constraints impose less costs, so the net benefit to annuitization is greater.

Results

Magnitudes

Pursuant to equation (2.16), CV for all consumers is equal to -125.39 when there is no annuitization. This means that for a very small increase in annuitization dA , total expenditures are equal to $100 - 125.39dA$.

Consistent with past results, the EV for a change from complete to zero annuitization is 44 in the case of log utility with no standard of living effect, and the discount and interest rates equalized (case 1). The positive EV in this case is guaranteed by Claim 4. Consistent with our expectations, and past results, EV increases in risk aversion (to 56 with $\gamma = 2$ (case 7), and decreases when discounting is heavier (to 15 with a discount rate of 0.1 and log utility (case 3)). Note that the value of annuitization is increasing as the trajectory of optimal consumption with no constraints approaches the flat (or upward sloping) optimal annuitized consumption path. This is because the positive wealth constraints have less bite when optimal consumption is nearly flat.

As expected, EV increases with the introduction of the standard of living effect

when the initial standard of living is small relative to wealth. The externality of present consumption on future consumption leads consumption to be upward sloping over some range, with large mortality eventually bringing consumption to lower than initial levels. This occurs around age 80 for $\delta = 1.03^{-1}$ and around 70 for $\delta = 1.1^{-1}$. Comparable to Case 1, Case 2 has a small standard of living and large valuation of 67, approximately 50 percent greater. The initial value of 5 is “small” because with expenditures of 100, a constant real annuity pays out 8.5 per period. Case 4 is comparable to case 3, and here valuation almost doubles. Case 8 is comparable to Case 7, and valuation increases, although not as sharply from 56 to 70.

Assuming rationality and retirement with assets sufficient to sustain the standard of living enjoyed going into retirement, we see that this particular relaxation of additivity has an economically very significant effect on valuation.

For the cases where the initial standard of living is large, there is an economically significant effect on valuation in the other direction. Cases 5, 6 and 9 are comparable to cases 1 and 2, 3 and 4, and 7 and 8, respectively. In this case, the consumption smoothing effect of the externality of early consumption on subsequent utility is overcome by the desire to smooth the ratio of consumption to the standard of living. The standard of living must fall over time because the initial standard of living (50 in cases 5, 6 and 9) are not sustainable throughout retirement.

Perhaps the most striking result is the consistently very large fraction of savings placed in the annuity. The minimal annuitized fraction is 60, and this is for a consumption plan (case 9) with very sharply decreasing consumption. For the familiar additively separable case, the minimal annuitized fraction is 72 percent. Following Claim 3, optimal annuitization increases in patience. Following Claim 4, complete annuitization is optimal in the separable case when $\delta(1+r) = 1$. The last column, showing equivalent variations associated with the optimal level of constant real annuitization can be interpreted as stating that for all of our simulations, annuitizing 60 percent of wealth is equivalent to an increase in wealth of at least 20 percent.

For some patient consumers, the presence of a constant real annuity is as good as the presence of complete annuities. This outcome, however, is a knife-edge. Adding

a savings motive through the standard of living effect, we find a benefit to the move to complete annuities markets, because upward sloping consumption can be financed through initial annuity purchases, rather than future savings and bond purchases. For more impatient consumers, the move to complete annuities markets introduces considerable further gains relative to the introduction of the real annuity.

Table 2.1: Summary of simulations

Case	δ	γ	s_1	α	EV_R	A^*_R	EV_C	EV_O
1	1.03^{-1}	1	1	0	44	100%	44	44
2	1.03^{-1}	1	5	1	67	100%	82	67
3	1.1^{-1}	1	1	0	15	72%	24	20
4	1.1^{-1}	1	5	1	36	99%	37	36
5	1.03^{-1}	1	50	1	36	84%	49	46
6	1.1^{-1}	1	50	1	3	63%	24	21
7	1.03^{-1}	2	1	0	56	100%	56	56
8	1.03^{-1}	2	5	1	70	100%	87	70
9	1.03^{-1}	2	50	1	negative	60%	30	27

2.5 Conclusion

With complete markets, the result of complete annuitization survives the relaxation of several standard, but restrictive assumptions. Utility need not satisfy the von Neumann-Morgenstern axioms and need not be additively separable. Further, annuities must only offer positive net premia over conventional assets; they need not be actuarially fair. We have retained the abstractions of no bequest motive, and no learning about health status or other liquidity concerns. Exploring the consequences of dropping these assumptions in the context of non-separable preferences and unfair annuity pricing will be an important generalization, but obtaining results will require strong assumptions both on annuity returns and on the nature of bequests and liquidity needs.

When annuities are restricted to be constant in real terms, we find that adding a particular form of intertemporal dependence in utility reinforces the benefits of annuitization as long as the standard of living entering retirement is not too large

relative to resources. Even in what we regard to be extremely unfavorable conditions for annuitization, 60 percent of savings are optimally placed in the annuity. For consumers who benefit less from constant real annuitization, there are substantial gains to completion of annuity markets.

Figure 2-1: Case 1




Figure 2-2: Case 2

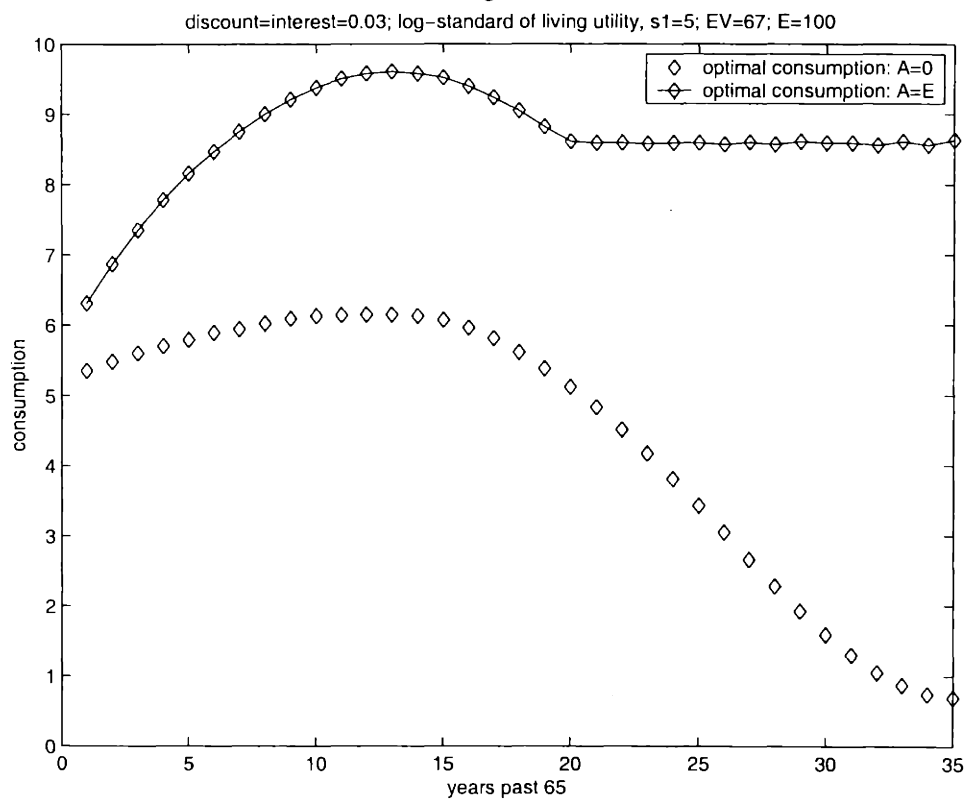


Figure 2-3: Case 3

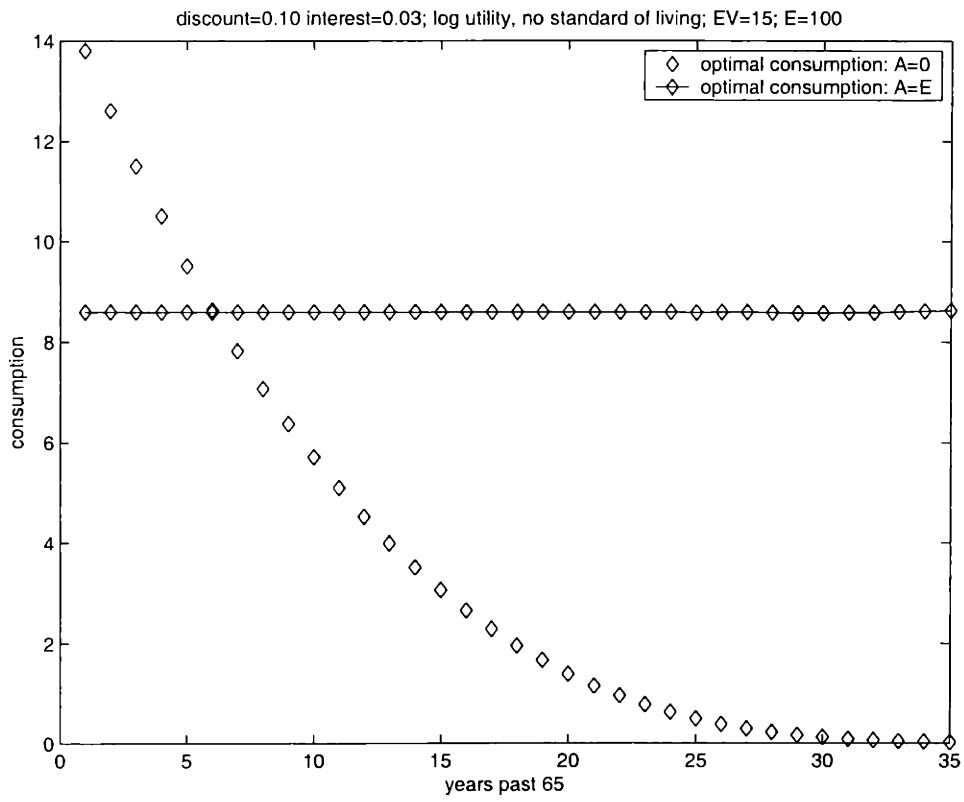


Figure 2-4: Case 4




Figure 2-5: Case 5

discount=interest=0.03; log-standard of living utility, $s_1=50$; $EV=36$; $E=100$




Figure 2-6: Case 6

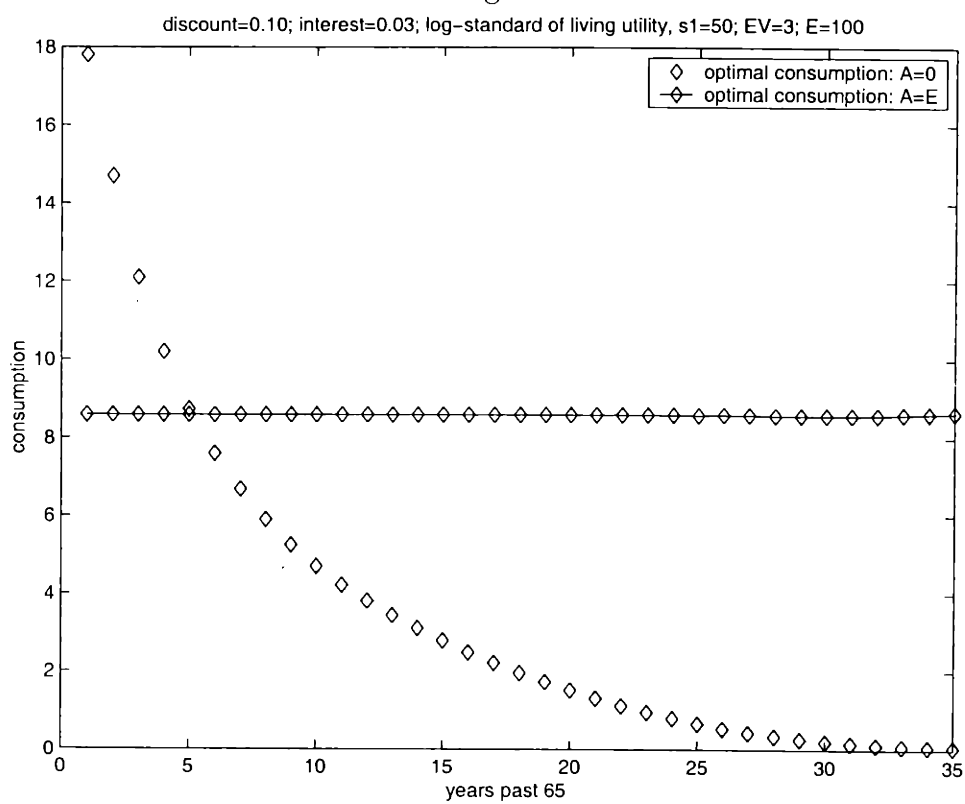


Figure 2-7: Case 7

discount=interest=0.03; gamma = 2, no standard of living; EV=56; E=100




Figure 2-8: Case 8

discount=interest=0.03; gamma = 2, standard of living utility, s1=5; EV=70; E=100





Figure 2-9: Case 9

discount=interest=0.03; gamma = 2, s1=50; EV: negative; E=100



Chapter 3

Income Sorting: Measurement and Decomposition

3.1 Introduction

The causes, extent and consequences of the segregation of demographically heterogeneous populations into relatively homogeneous neighborhoods and jurisdictions are objects of considerable interest among economists and social scientists generally.¹ Segregation on the dimension of income (“income sorting”), at the jurisdictional level, is particularly interesting because, under some conditions, it is an equilibrium condition in the political economy models of jurisdiction choice that follow from Tiebout (1956), as in Epple and Sieg (1999). Tiebout’s notion that households choose jurisdictions based on the package of public goods they offer has also motivated studies of the valuation of public goods based on differences in property values across jurisdictions with different measured public goods quality.

This paper addresses the measurement of income sorting and the attribution of observed sorting to different causes. In terms of measurement, I show that a decomposition of variance into within jurisdiction and between jurisdiction components must be adjusted for the presence of measurement error in income. Using 1990 US Cen-

¹see, for example, Wilson (1987), Benabou (1993), Kremer (1997), Glaeser and Cutler (1997).

sus data, I find that the adjustment approximately doubles the estimated extent of sorting. On average, across all US metropolitan areas (MSAs) I find approximately nine percent of the variation in household income can be explained by differences across jurisdictions. There is a great deal of heterogeneity in sorting: while sorting is statistically significant in almost all MSAs, the fraction of variation attributable to jurisdictions is very close to zero in a large number of MSAs, while the majority of variation can be explained by jurisdictions in others.

That jurisdictional differences do not yield perfect sorting is hardly surprising given the many dimensions of preferences that enter housing choice. Epple and Sieg (1999) observe that some very strong assumptions are required to obtain perfect sorting by income. Among the strongest are that preferences over public goods can be reduced to a single dimension, that the level of public goods provision is equal across locations within jurisdictions and that the housing stock within jurisdictions is determined by the preferences and budget constraints of households of present residents (so that we would only observe fixed housing quality in jurisdictions over time if relative amenity were also unchanging).

An empirical observation that there are significant differences in incomes across jurisdictions, combined with the fact that there are differences in public goods across jurisdictions cannot be interpreted as proof that differences in government drive, or even enable income sorting. Jurisdictions are differentiated not only by government, but often by geographic amenity and housing quality. The well-established difficulties in estimating hedonic values for location and amenity are compounded by the fact that amenity characteristics such as school performance are likely to be determined in part by the characteristics of the households who use the amenity. Thus, if households sort by preferences over geography, and not at all on the basis of public goods, we will find high income households in the more geographically desirable locations, and also likely superior school performance and lower crime.

Recently, boundaries have come to prominence as a way around the identification problems caused by endogenous jurisdiction choice and formation. Black (1999) shows that controlling for all observables, virtually adjacent houses on opposite sides of

school attendance lines within the same jurisdiction reflect quality differences in their associated schools in different prices. Relatedly, Hoxby (1994) uses the number of rivers in MSAs as a source of exogenous variation in the number of jurisdiction to estimate the effect of school choice on school quality. Both methodologies rely on the lack of amenity effects associated with being on one side or the other of these boundaries, independent of the associated difference in public goods.

I use a similar methodology to estimate the extent to which income sorting by jurisdiction is driven by differences in government, rather than locational characteristics. To do this, I compare the extent of income sorting in two types of adjacent zip (postal) code pairs. The first type of adjacent zip code pairs is the set of zip codes which are next to each other, and are in the same jurisdiction (e.g. Cambridge, MA 02138 and Cambridge, MA 02139). The second type of adjacent zip code pairs is the set which are physically next to each other, but are separated by both a postal boundary and a jurisdiction boundary (e.g. Cambridge, MA 02138 and Somerville, MA 02143). Zip code boundaries are drawn by the US postal service to rationalize delivery routes logistically. It is natural to think that these boundaries carry no more information concerning locational amenity than jurisdictional boundaries. If this condition is met, then the difference between the average extent of sorting between across-jurisdiction pairs and the average extent of sorting between within-jurisdiction pairs, corrected for sampling properties should be no smaller than the extent of sorting generated by a combination of geography and government (the average extent of sorting across neighboring jurisdiction pairs) minus the extent of sorting generated by purely geographic differences (the average extent of sorting among within jurisdiction pairs).

In 1990 Census data, I find that, on average, location on one or another side of a zip code boundary explains approximately 2.2 percent of the variance of household income. Controlling for population characteristics of the zip code pair, the addition of a jurisdiction boundary increases the R^2 by an average of 0.4 percent. This leaves the large majority of locational income sorting unexplained by governmental differences.

The second section of this paper discusses methodological issues in the measure-

ment of income sorting. The third section discusses the data I use to estimate income sorting at the jurisdiction and zip code level within MSAs, and the fourth section summarizes the extent of sorting I find. The fifth section presents the decomposition analysis, and the sixth section concludes.

3.2 Measuring Income Sorting

A natural way to measure sorting by any characteristic within subregions of a larger region (here, jurisdictions or zip codes within MSAs, “jurisdictions” hereafter when either can be meant), is to compare the average variance of the characteristic within jurisdictions to the variance at the regional (MSA) level.

Kremer and Maskin (1996) note that such a variance decomposition has a neat interpretation as the R^2 in a regression of the characteristic on a full set of dummy variables indicating individual residence in each of the jurisdictions. Indexing households by h and jurisdictions by j , and labeling income y , we have:

$$R^2 \equiv 1 - \frac{\sum_{j=1}^J \sum_{h=1}^{H_j} \frac{(y_h - \bar{y}_j)^2}{H_j}}{\frac{\sum_{h=1}^H (y_h - \bar{Y})^2}{H}}, \quad (3.1)$$

where \bar{y}_j is mean income in jurisdiction j , and \bar{Y} is mean income in the MSA. The numerator of the second term on the right hand side is the population weighted average of within jurisdiction variance. The denominator is the variance at the MSA level. If jurisdictions are close to homogenous, the fraction is small, and R^2 is large (with a maximum of one). If the expected squared difference between households’ income within jurisdiction is equal to the squared difference between households at the regional level, then there is no sorting, and we have and R^2 of zero. Decomposing total variance, we can also interpret the R^2 measure as the ratio of the population weighted average squared deviations of jurisdiction mean incomes from the population mean divided by total variance.

Assuming that variance is a meaningful measure of heterogeneity, the R^2 statistic can be applied to sample data, but only after two defects are addressed.

3.2.1 Adjustment for Sampling Without Replacement

A well known problem associated with the R^2 measure is that increasing the number of regressors increases the expectation of R^2 in finite samples, even if the added regressors are orthogonal to the dependent variable (here, income). That is, in a world with no behavioral income sorting, MSAs with more jurisdictions would have greater R^2 values mechanically.

To make the R^2 measure an estimate of behavioral sorting, we can observe that the expectation of variance within a jurisdiction, when households are randomly taken from a sample of the MSA without replacement, is given by:

$$E \frac{1}{H_j - 1} \sum_{h=1}^{H_j} (y_h - \bar{y}_j)^2 = \frac{1}{H - 1} \sum_{h=1}^{H_j} (y_h - \bar{Y})^2.$$

Thus, if we replace H_j with $H_j - 1$ in the numerator and H with $H - 1$ in the denominator of equation (3.1), with random assignment of households to jurisdictions, we would obtain an expected R^2 of zero. With behavioral sorting, the expectation will be greater than zero.²

In the data I consider, populations are too large for this adjustment to make a significant difference. It should be observed, however, that the correction comes from the property that variance of sample means decrease in sample size. Other widely used measures such as the Index of Dissimilarity and Thiel's index suffer from the same bias towards observed sorting when jurisdiction sizes are small. Any finding of increasing segregation with increasing fragmentation should be presented and interpreted with care.

3.2.2 Measurement Error

It is well known that mean zero measurement error in the dependent variable, uncorrelated with either the right or left hand side variables in a regression will not bias

²In general, adding more jurisdictions, or equalizing the population share of jurisdictions allows for a smaller value of R^2 : this only affects the expectation if there is behavioral sorting. A finding that adding jurisdictions yields larger estimated adjusted R^2 's means only that there is sorting, not necessarily that sorting behavior is more pervasive in more fragmented regions.

estimated coefficients. However, such error will bias down estimated R^2 . Putting aside the small denominator adjustment, suppose that reported income is $y_h + v$, where v is mean zero and i.i.d. across households with variance $\sigma_v^2 = Ev^2$. Our estimate of R^2 becomes

$$1 - \frac{\sum_{j=1}^J \sum_{h=1}^{H_j} \frac{(y_h - \bar{y}_j + v_h)^2}{H_j}}{\frac{\sum_{h=1}^H (y_h - \bar{Y} + v)^2}{H}} \quad (3.2)$$

As the signal to noise ratio $\frac{\sigma_y^2}{\sigma_v^2}$ approaches zero, (3.2) approaches zero, even if there is perfect sorting by true income. If the signal to noise ratio were known, then an unbiased estimate of R^2 could be obtained. Our expectation of mean squared error within jurisdictions is equal to

$$\sigma_{yj}^2 + \sigma_v^2 + 2\sigma_{y_j v},$$

and our expectation of mean squared error at the MSA level is

$$\sigma_y^2 + \sigma_v^2 + 2\sigma_{yv}.$$

Under the assumption that the covariance terms $\sigma_{y_j v}$ and σ_{yv} are zero at both the jurisdiction and MSA level (so that measurement error is no more severe in any jurisdictions than in others), we can subtract an estimate of σ_v^2 from both expressions to obtain unbiased estimates of R^2 . That is, a measurement error adjusted R^2 is equal to

$$R_{me}^2 = 1 - \sum_j \frac{H_j}{H} \frac{\hat{\sigma}_{yj}^2 - \hat{\sigma}_v^2}{\hat{\sigma}_y^2 - \hat{\sigma}_v^2}. \quad (3.3)$$

Measurement error v can come from several sources in cross sectional survey data. First, we are typically interested in a measure of sorting by wealth rather than income, but annual, rather than lifetime, income is reported in most survey data. This would not be a problem if annual income were simply equal to a constant fraction of lifetime income. However, this relationship is violated both by year-specific shocks

to income and by a generally upward trending age-earnings profile. A young graduate student may exert at least as great a positive externality on neighbors as an old tenured professor, but will show up in cross sectional data as low income. Second, households may misreport their earned income in the survey year. Third, and particularly important in the census data, income is reported in bins, so that we must guess an income for each household with income reported in that range.

To estimate σ_y^2 and σ_v^2 separately, we recall the formula for attenuation bias in a regression where a single right hand side variable is measured with error. If we regress some variable Z on reported income y , which is a noisy measure of true income y^* :

$$Z_h = a + b(y^*_h + v) + \epsilon, \quad (3.4)$$

then we have

$$plim(b_{OLS}) = \frac{Cov(y, Z)}{\sigma_y^2 + \sigma_v^2}.$$

By contrast, if we find an instrument that is correlated with y^* , but not with v or ϵ , then the two stage least squares estimator b_{IV} has the true coefficient on income as a probability limit:

$$plim(b_{IV}) = b = \frac{Cov(y, Z)}{\sigma_y^2}.$$

Comparing the OLS and IV estimators yields the relationship

$$plim\left(\frac{b_{IV}}{b_{OLS}}\right) = \frac{\sigma_y^2 + \sigma_v^2}{\sigma_y^2}. \quad (3.5)$$

Since we observe the combined signal and noise variance of observed income $\sigma_y^2 + \sigma_v^2$, if we can find an auxiliary regression of the form described above, then we can estimate measurement error and the unbiased sorting estimator R_{me}^2 .

3.3 Data

I estimate the extent of income sorting using 1990 US census (STF3A) data on the distribution of household incomes at the MSA, jurisdiction and zip code levels within 207 MSAs. For each of these geographic entities, I observe the number of sampled households (approximately 15 percent on average) reporting 1989 income in each of 25 income categories. I assume that all households reporting income in any income bin reported the midpoint income of the bin. For example, I deem a household in the income category of \$10,000 to \$15,000 to have an income of \$12,500.

The Census Bureau aggregates data at the level of metropolitan area, which are physically continuous areas that can plausibly be considered a single market. Incorporated jurisdictions typically compose only a portion of their encompassing MSA, because some areas are not incorporated into political units below the county level. Depending on the state, following Census descriptions of the primary unit of local government, I define jurisdictions as either “incorporated places” or “minor civil divisions.” Such entities undertake a significant fraction of all state and local spending.

To estimate measurement error in income, I use 1990 Census microdata (not available with jurisdictional or zip code detail) on household incomes, education and house price. The microdata reports a bounded, continuous (integer) value for income, which I transform into the midpoint of the corresponding bin that would be reported in the geography-specific STF3A data (so that a household reporting income of \$12,300 is assigned the \$12,500 midpoint of the \$10,000 to \$15,000 bin). I regress the reported value of housing (among homeowners) on the transformed income variable to obtain the OLS estimate. The IV estimate is obtained by instrumenting for the transformed income variable with the mean income for the industry and MSA cell in which the household head works (so that all households in Pittsburgh headed by a real estate broker receive the same aggregated instrument). I perform these regressions in both logs and levels, at the MSA level. The assumption is that the industry mean affects housing purchases only through household income, and not through any unobserved variables included in the error term ϵ in equation (3.4). Use of education as an

alternative instrument can corroborate this assumption.

3.4 Results

With no correction for measurement error, across 207 US MSAs, at the jurisdiction level I find a mean R^2 for the level of income of 0.04. At the zip code level, I find a significantly larger value of 0.07. As reported in Table 3.1, the results are approximately identical for sorting by the natural log of income.³ We can infer that zip codes are considerably more homogenous than are jurisdictions. Already we see that there is sorting that is not driven entirely by shared government. Neighborhoods (zip codes) within jurisdictions may have different service levels, but these do not arise from the simple Tieboutian voting with feet mechanism.⁴

However, I find considerable measurement error in income, which is worsened by assuming bin midpoint values for households' incomes. Running OLS and IV equations of the reported value of homeowners' homes on transformed income for each MSA, I find a mean OLS coefficient of 0.75 in levels and 0.37 in natural logs. By contrast, instrumenting for income with industry mean income, I find a mean coefficient of 1.45 in levels and 0.70 in logs. Taking means over MSAs, I estimate noise to signal ratios ($\frac{\sigma_v^2}{\sigma_y^2}$) of 0.95 in levels and 0.91 in logs. Encouragingly, I obtain almost identical results using education as an instrument for income - the similarity of IV results suggests that the observed positive effect of education on housing value observed in demand regressions may be purely attributable to unobservable components of lifetime income. Both education and industry mean wages should overcome the problem associated with unobservable age earnings profiles.

Correcting the sorting measures accordingly, I find a mean R^2 of 0.09 in levels and 0.11 in log income at the jurisdiction level, and 0.14 in levels and 0.17 in logs at the zip code level.

³We might think that ratios of income are more relevant to choice than levels

⁴One could imagine more involved political models which result in heterogeneous services within jurisdictions such that households are satisfied, but this is outside of models such as Epple and Sieg (1999)

Table 3.1: R^2 Estimates of income sorting

	(1)	(2)	(3)	(4)
Zip code	0.07	0.14	0.08	0.17
Jurisdiction	0.04	0.09	0.05	0.11
Income Specification	Level	Level	Log	Log
Measurement Error Corrected?	No	Yes	No	Yes

Note: Values reported are population weighted average ratios of within jurisdiction or zip code variance to variance of MSA income, averaged across 207 US MSAs.

The corrected R^2 estimates provide unbiased estimates of the extent of sorting by income into zip codes and jurisdictions. Two notable facts are, first, that while the extent of sorting is statistically significant in almost every MSA, the large majority of income differences across households survive to the level of jurisdiction and the finer level of zip code. Second, a significant fraction of the variation in income not explained by differences across jurisdictions is explained by differences across zip codes within jurisdictions. This implies that there are important differences in conditions within jurisdictions, so that government is imperfectly correlated with other local amenities.

Notably, matching the result of Epple and Sieg (1999), I obtain an R^2 measure of 0.11 in the level of income for the Boston MSA with no correction for measurement error. This estimate doubles to 0.22 with the correction.

The largest value of sorting by level of income is 0.65 in Savannah, Georgia, and 0.80 in levels in Decatur, Illinois. This matches a general fact that there is substantial sorting in Southern and Midwestern MSAs. These results are consistent with racial preferences driving income sorting (explaining the South), as in Alessina, Baqir and Hoxby (1999) or with proliferation of local governments (as in the Midwest) indicating that where governance is relatively important to households, income sorting is stronger. Beyond such impressionistic evaluations, identifying the underlying causes of sorting rigorously would be quite difficult.

In addition to measurement error, “perfect” income sorting is hindered by the fact that the large share of population located in some jurisdictions implies that there must be some mixing of income. For example, a metropolitan area with just two jurisdictions cannot feature an R^2 value of one as long as there are more than two

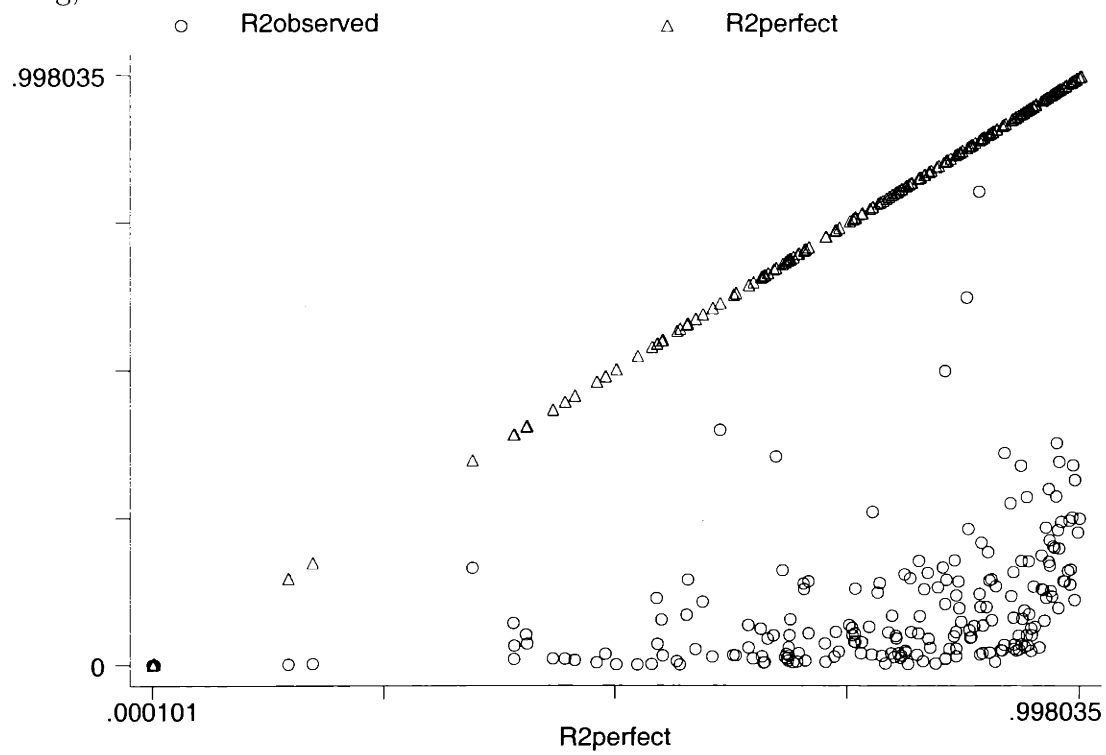
income categories which positive population. Large cities such as New York City are typically larger than the population in any single income category, and hence must feature some income mixing.

Figure 1 plots two values of R^2 for each metropolitan area. The smallest value is the observed extent of sorting by jurisdiction, corrected for measurement error and summarized in the second column of Table 3.1. The second value is income sorting which would occur if “perfect” sorting were accomplished by locating the wealthiest households in the smallest jurisdictions and the poorest households in the largest jurisdictions. If sorting were complete, the observed data (circles) would lie along the 45-degree line with the maximized R^2 values. As noted above, in most MSAs, this observed value is substantially less than one, with a mean value of 0.76. Evidently, observed sorting, with a mean of 0.09 on average falls far short of this value, so the lack of income sorting can be attributed only in small part to feasibility constraints. Further, the correlation between perfect sorting and actual sorting is highly imperfect, with a value of 0.40. Again, it would be interesting in future research to understand the basis for difference across metropolitan areas in the extent of sorting. Part of such an understanding will include distinguishing government-based sorting from sorting based on other determinants of jurisdiction choice.

3.5 Decomposing Income Sorting

As a start at understanding what drives income sorting, it would be interesting to disentangle governance and public goods from other locational attributes as causes of sorting. The fact that zip codes are more homogeneous than jurisdictions implies that there are neighborhood effects. This suggests that some fraction of the observed jurisdictional sorting, too, is driven by extra-governmental characteristics of jurisdictions, since jurisdictions, like zip codes, are locationally homogeneous relative to the larger metropolitan area. In theory, all of the observed sorting into jurisdictions could be driven by locational factors other than governance. However, the results are at least theoretically consistent with a world in which households have lexicographic

Figure 3-1: Income Sorting: Observed R^2 values and maximized values with “perfect sorting,” with and without measurement error



preferences, such that they choose first a governmental package (and hence a jurisdiction), and next choose the neighborhood they like best within the jurisdiction. In this case, all of the sorting at the jurisdiction level would be attributable to differences in governance.

Putting any structure on the basis of jurisdiction choice is a risky enterprise. As Epple and Sieg (1999) illustrate, it is not trivial to prove that households will sort into jurisdictions by income on the basis of governance, even under extremely strong assumptions. I thus do not attempt to recover underlying preference parameters. Rather, I simply ask: how much of the income sorting we observe at the jurisdictional level can be attributed to the fact that there are governmental differences between jurisdictions?

To answer this question, I estimate a sorting measure for “regions” defined as physically adjacent pairs of zip codes. The large majority of zip codes are located almost entirely within a single jurisdiction, and I confine the analysis to such zip codes. This way, the border between two zip codes can either be a mere postal division, or can be both a postal division as well as a jurisdictional division. Under an assumption on the nature of zip code and jurisdictional boundaries, the difference between the measure of sorting observed between across-jurisdiction zip code pairs and the measure of sorting observed between within-jurisdiction pairs, on average, can be interpreted as the sorting “value added” of governmental differences.

The critical assumption is that within jurisdiction zip code boundaries signify the same degree of extra-governmental neighborhood differences, such as differentiated topography or housing characteristics, as jurisdictional boundaries. I consider this to be a weaker condition than two other boundary-related identifying assumptions that have come to prominence in the literature on local public goods. Black (1999) assumes that if differences in prices of houses that are virtually adjacent, but are located across school attendance lines within the same jurisdiction, are orthogonal to observable characteristics and correlated with differences in school quality across the attendance lines, then the differences are caused by the differences in school quality. This implicitly relies on the assumption that school attendance lines convey

zero unobservable information about non-school neighborhood characteristics. This is stronger than my assumption, because I require only that, the magnitude of neighborhood differences across zip code lines are on average equal to the magnitude of neighborhood differences across jurisdiction lines. I do not require an absence of neighborhood differences as Black does.

Hoxby (1994) argues that the number of rivers in a metropolitan area causes jurisdictional fragmentation, but does not cause economic segregation. This assumption allows an interpretation of a relationship between the number of school districts and school quality to be interpreted as causal. If variation in school quantity were generated by a mechanism which generated correlated economic segregation, then we could not reject the alternative interpretation that the differences in average school quality is driven by economic segregation. For such an interpretation to be ruled out, it must be the case that rivers do not generate economic segregation. In turn, this requires that rivers do not mark changes in neighborhood characteristics. Again, I regard this as a stronger assumption than the assumption that zip code boundaries convey approximately the same amount of information as jurisdictional boundaries.

The Census Bureau provides the longitude and latitude of each zip code's centroid. For each zip code in each MSA, this allows me to find the closest different zip code, measured in centroid-to-centroid distance.⁵ For each zip code and its nearest neighbor, I can calculate an R^2 measure as if the two zip codes jointly formed a metropolitan area: the larger the estimated R^2 the greater the difference in mean household incomes across zip codes relative to total household income variance, and hence the greater the extent of sorting. Census data also indicates what fraction of each zip code is located in which jurisdiction. I consider only zip codes for which at least 95 percent of the population is located within a single jurisdiction. Across the 207 MSAs, approximately half of the nearest neighbor pairs are located in the same jurisdiction. Same jurisdiction zip code pairs will tend to be in larger cities, on average, since small jurisdictions tend to have just one zip code.

With measurement error corrections, I find a mean R^2 of 0.0217 in the level of

⁵This relationship need not be reflexive in general

Table 3.2: Regressions of Income R^2 for Neighboring Zip Codes on a Variable Indicating Different Jurisdictions

	(1)	(2)
Different Jurisdiction	0.0037** (0.0014)	0.0078** (0.0030)
R^2 for income in	level	log
Mean R^2	0.0217	0.0455
Observations	4,685	4,692

Notes: Dependent variable is the 1 minus the population-weighted average ratio of within zip code income variance to total variance in neighboring zip code pairs. Both regressions include MSA fixed effects as well as polynomials and interactions of mean income level, total population of the combined zip codes, and population share of the larger zip code. ** Significant at 5 %.

income, and 0.0455 in log income. Without covariates, I find that the mean R^2 values are indistinguishable whether the zip code boundary is a within jurisdiction boundary, or also a jurisdictional boundary. Adding demographic covariates, however, I find a small positive effect on R^2 of jurisdictional difference. Table 3.2 shows that “adding” a jurisdictional boundary increases the extent of sorting by approximately 0.004 in levels and approximately 0.008 in logs. Evidently, different governments explain only a small fraction of locational sorting on income.

3.6 Conclusions

Jurisdictions are segregated by income relative to metropolitan areas. Correcting for measurement error in income approximately doubles estimates of the extent of sorting, but sorting remains far from complete. Observed sorting at the jurisdiction level may be generated by differences in tax and spending policies, or by differences in extra-governmental amenity, or, most likely, by a combination of the two. The evidence presented here suggests that extra-governmental amenity plays a very large role in the sorting process, and that these amenities vary not only across jurisdictions, but also within.

The conclusion that neighborhood effects are an important source of sorting within jurisdictions is supported by the fact that zip codes are considerably more homoge-

neous than jurisdictions. Further support comes from the fact that neighboring zip codes are only slightly more sorted by income if they are in different jurisdictions than if they are in the same jurisdiction. The existence of significant income differences across zip codes suggests that there are likely to be differences associated with other boundaries, such as rivers and school attendance lines which have previously been assumed to be innocuous with respect to neighborhood effects. The evidence suggests further that there are likely to be interactions between neighborhood characteristics and government. Different zip codes in the same jurisdiction likely enjoy different level of public goods quality (by virtue of different spending on different schools, or different access to public facilities. The fact that different governments explain only a small fraction of income sorting implies only that government is not the sole factor driving sorting. It does not imply that quality of access to public goods is relatively unimportant.

Bibliography

- Alessina, Alberto, Reza Baqir, and Caroline Hoxby (1999) 'Political jurisdictions in heterogeneous communities'
- Benabou, Roland (1993) 'Workings of a city: Location, education, and production.' *Quarterly Journal of Economics* 108(3), 619–652
- Berkovec, James, and Don Fullerton (1992) 'A general equilibrium model of housing, taxes and portfolio choice.' *Journal of Political Economy* 100(2), 390–4429
- Bernheim, B. Douglas (1991) 'How strong are bequest motives? evidence based on estimates of the demand for life insurance and annuities.' *Journal of Political Economy* 99(5), 899–927
- Black, Sandra (1999) 'Do better schools matter? parental valuation of elementary education.' *Quarterly Journal of Economics* pp. 577–599
- Bound, John, and Alan B. Krueger (1991) 'The extent of measurement error in longitudinal earnings data: Do two wrongs make a right?' *Journal of Labor Economics*
- Brown, Jeffrey R. (2001a) 'Private pensions, mortality risk, and the decision to annuitize.' *Journal of Public Economics*
- Brown, Jeffrey R. (2001b) 'Redistribution and insurance: Mandatory annuitization with mortality heterogeneity.' Technical Report, Boston College Center for Retirement Research. Working Paper 2001-02
- Campbell, John, and Joao Cocco (2001) 'Household risk management and mortgage choice'

- Cocco, Joao (2000) 'Hedging house price risk with incomplete markets'
- Davis, Steven J., and Paul Willen (2000) 'Occupation-level income shocks and asset returns: Their covariance and implications for portfolio choice.' NBER Working Paper W79025
- Diamond, Peter, and Daniel McFadden (1974) 'Some uses of the expenditure function in public economics.' *Journal of Public Economics* 3, 3–21
- Epple, Dennis, and Holger Sieg (1999) 'Estimating equilibrium models of local jurisdictions.' *Journal of Political Economy* 107(4), 645–681
- Flavin, Marjorie, and Takashi Yamashita (2001) 'Owner-occupied housing and the composition of the household portfolio over the life-cycle.' *American Economic Review*
- Glaeser, Edward, and David Cutler (1997) 'Are ghettos good or bad?' *Quarterly Journal of Economics* 112, 827–872
- Griliches, Zvi, and Jerry Hausman (1986) 'Errors in variables in panel data.' *Journal of Econometrics* 31(1), 93–118
- Heaton, John, and Deborah Lucas (2000) 'Portfolio choice and asset prices: The importance of entrepreneurial risk.' *Journal of Finance* 55(3), 1163–1198
- Hoxby, Caroline Minter (1994) 'Does competition among public schools benefit students and taxpayers?' Technical Report, NBER, December. W4979
- Hurd, Michael (1989) 'Mortality risk and bequests'
- Ioannides, Yannis M., and Stuart S. Rosenthal (1994) 'Estimating the consumption and investment demands for housing and their effect on housing tenure status.' *The Review of Economics and Statistics* 76(1), 127–141
- Jousten, Alain (2001) 'Life-cycle modelling of bequests and its implications for annuity valuation.' *Journal of Public Economics* 79, 149–177

- Kennickel, Arthur B., Martha Starr-McCluer, and Brian J. Surette (2000) 'Recent changes in u.s. family finances: Results from the 1998 survey of consumer finances.' *Federal Reserve Bulletin* pp. 1–29
- Kotlikoff, Laurence J., and Avia Spivak (1981) 'The family as an incomplete annuities market.' *Journal of Political Economy* 89(21), 372–391
- Kremer, Michael (1997) 'How much does sorting increase inequality?' *The Quarterly Journal of Economics* pp. 115–139
- Kremer, Michael, and Eric Maskin (1996) 'Wage inequality and segregation by skill'
- Laibson, David, Andrea Repetto, and Jeremy Tobacman (2000) 'A debt puzzle.'
NBER Working Paper
- Mas-Collel, Andreu, Michael Whinston, and Jerry Green (1995) *Microeconomic Theory* (New York: Oxford University Press)
- Milevsky, Moshe A. (2001) 'The real option to delay annuitization: Its not now or never.' Working paper presented at 2001 Meetings of the American Economic Association
- Mitchell, Olivia, Jeffrey R. Brown, James M. Poterba, and Mark Warshawsky (1999) 'New evidence of the money's worth of individual annuities.' *The American Economic Review* 89(5), 1299–1318
- Newey, Whitney K., and Kenneth D. West (1987) 'A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix (in notes and comments).' *Econometrica* 55(3), 703–708
- Ortalo-Magne, Francois, and Sven Rady (1998) 'Housing market fluctuations in a life-cycle economy with credit constraints.' Research Paper No. 1501
- Palmon, Oded, and Avia Spivak (2001) 'Information and the market for annuities.'
Mimeo

- Poterba, James M. (1997) 'The history of annuities in the united states.' Technical Report, NBER, April. W6001
- Rothschild, Michael, and Joseph Stiglitz (1970) 'Increasing risk i: A definition.' *Journal of Economic Theory* 2, 225–243
- Shiller, Robert (1993) *Macro Markets* (Oxford: Clarendon Press)
- Sinai, Todd, and Nicholas Souleles (2001) 'Owner occupied housing as insurance against rent risk.' Manuscript, Wharton
- Tiebout, Charles M. (1956) 'A pure theory of local expenditures.' *Journal of Political Economy* 64(4), 442–444
- Venti, Steven, and David Wise (2000) 'Aging and housing equity.' NBER Working Paper 7882
- Viceira, Luis (2001) 'Optimal portfolio choice for long-horizon investors with non-tradable labor income.' *The Journal of Finance* LVI(2), 433–470
- Vissing-Jorgenson, Annette (2000) 'Towards an explanation of household portfolio heterogeneity: Nonfinancial income and participation cost structure.' Working Paper, University of Chicago Department of Economics
- Wald, A. (1940) 'The fitting of straight lines if both variables are subject to error.' *Annals of Mathematical Statistics* 11(3), 284–300
- Walliser, Jan (2001) 'Regulation of withdrawals in individual account systems.' In *New Ideas About Old Age Security: Toward Sustainable Pension Systems in the 21st Century*, ed. Robert Holzmann and Joseph Stiglitz (The World Bank) pp. 367–390
- Warshawsky, Mark J., Brenda Spillman, and Chris Murtaugh (Forthcoming) 'Integrating the life annuity and long-term care insurance: Theory, evidence, practice and policy.' In *Financial Innovations in Retirement Income*, ed. Olivia Mitchell, Zvi Bodie, Brett Hammond, and Steve Zeldes

Wilson, William Julius (1987) *The Truly Disadvantaged: The Inner City, the Underclass and Public Policy* (The University of Chicago Press)

Yaari, Menahem E. (1965) 'Uncertain lifetime, life insurance, and the theory of the consumer.' *Review of Economic Studies* 32, 137

Yao, Rui, and Harold H. Zhang (2001) 'Optimal consumption and portfolio choice with risky housing and stochastic labor income.' University of North Carolina, Chapel Hill