

RETIREMENT SAVINGS—GAMES THAT ASSET MANAGERS, DISTRIBUTORS, AND INVESTORS PLAY

Joseph E. Stiglitz

The decisions individuals make concerning retirement savings—both how much to save and how to invest their savings—are among the most important they face in their life time. With increasing longevity, individuals are spending an increasing fraction of their lifetime in retirement.¹ Just seventy five years ago, few people lived much beyond retirement. Today, it is not uncommon for individuals to be retired for thirty years or more—an amount equal to half or more of their working lives.

Unlike other decisions individuals make, they have little foundation on which to base their decisions. When they choose among beef, chicken, and pork, or between one kind of lettuce and another, they can try one or the other, find out what they like and do not like, and make future decisions using the information acquired.

But in their retirement decisions, individuals do not have a second chance. If they save too little, when they are 90 years old they cannot do it over again. They have to live with the consequences. Making matters even worse is the fact that they cannot learn from their elders. With the marked changes in wages, securities markets, and public social security programs the choices facing individuals today and the consequences are markedly different from those of previous generations.

Making matters worse, such decisions are among the most complicated facing an individual. At a technical level, they involve intertemporal optimization with risk. While economists have begun to formulate models describing how such decisions should be made, the models are complicated and clearly beyond the scope of ordinary individuals to solve. Individuals in making their decisions will thus have to rely on simple heuristics, hunches, hearsay, advice, and judgment. They may look to peers for ‘norms’. If they see everyone around them saving 10%, they will be inclined to save 10%. A few people who see themselves as particularly healthy and lived, or particularly worried about old age poverty, might save a little more than 10%. But 10% becomes the norm.

How norms get established, or changed, is a matter that would take us beyond the scope of this paper. But analysis may play an important role. If individuals see that, at current interest rates, increases in wage rates, etc. they are likely to be left with what they view as insufficient funds in their retirement at a 10% savings rate will be induced to save more.

This provides those in the investment community trying to provide products and advice to retirees both an opportunity and a responsibility. Individuals are more vulnerable,

¹ See, for example, the United Nations report "World Population Ageing, 1950-2050," available at: <http://www.un.org/esa/population/publications/worldageing19502050/index.htm>

perhaps more easily misled—because it will be years before mistakes become evident. There are many who would willingly take advantage of this limited information.

For those who take their responsibilities seriously, the task of providing guidelines for retirement savings is not easy. They want to guide individuals, but they recognize that they are not fully informed about the preferences or circumstances of their customers; and that their customers may not fully understand their circumstances and preferences. So while they may wish to provide guidance, they must also provide choices. The guidance they provide must be based on what we know of ‘average’ individuals, or ‘average’ individuals in different circumstances, how their incomes and assets and needs evolve over time.

For instance, some individual’s career paths and wage profiles may be highly predictable; others may not. Some individuals may own a home, an important asset; others may not.

This paper looks at the overall problem of providing retirement security from the perspective of the three major ‘actors’ in the market—the individual investor/retiree, the distributor, and the asset manager.

The Investor

At the center of any discussion of retirement is the worker, who will eventually retire. Today, public pensions provide only a part of his or her income. For simplicity, we will assume that that part is known.² We focus on how the individual will supplement the public program, both through savings through the firm and on his own account.

While there has been an extensive literature on individual savings and portfolio decisions, that literature has faced three important limitations. First, it ignores human capital, the individual’s most important asset. This is an asset whose value changes over time; and while there is considerable uncertainty, the degree of uncertainty too changes over time. For instance, by the time the individual reaches retirement age, there is no human capital left (by definition), and no uncertainty about its value; as the individual approaches retirement, the value of this asset and typically, uncertainty about that value, diminishes. Early in an individual’s career, for instance, individuals may have limited information about their life prospects; subsequent promotions—or being passed over for promotion—resolves the uncertainties.

The existing literature looks carefully at the correlations among financial assets. Much of the conventional wisdom argues that as individuals approach retirement, they have fewer periods over which to average out outcomes, so that they should act in a more risk averse manner, investing more money into bonds. Much of the formal literature³ questions that

² Of course, with pension rules changing so much and so often, many individuals do not feel confident about what they will receive.

³ See, for example, John Y. Campbell, Joao F. Cocco, Francisco J. Gomes, and Pascal J. Meanhout, “Investing Retirement Wealth: A Life-Cycle Model,” NBER Working Paper No. 7029, March 1999.

presumption, suggesting that in the benchmark case of constant relative risk aversion, individuals should keep the same fraction in equity as they get older.

Bringing human capital puts another twist into this advice. It says that if human capital risk is diminishing over time, then perhaps individuals should be willing to undertake more financial capital risk, so that the fraction of wealth invested in equities should *increase* as individuals get older. (There are, in addition, industry or firm specific risks that individuals should take into account. An individual working in the automobile industry might want to short auto stocks, to provide some insurance against the variability in his human capital—recognizing that there are typically large losses in wages when individuals have to move from one job to another.)

Secondly, it ignores other assets, like real estate, which for most families, is the most important asset other than human capital. Again, much of the formal literature ignores housing, even though it has distinctive risk characteristics. In particular, the price of housing is likely correlated with what individuals may have to spend in the future for housing—one of the most important categories of expenditure. Thus, housing provides good insurance against price variations; it is inflation insurance for the most important single component of household expenditures. Thus, in assessing the role of housing in a diversified portfolio, one just treats it like another asset, looking at the correlation of its returns with that of other assets. Rather, one has to take explicitly into account its role in insuring against risks of changes in the relative price of housing.⁴

Such an analysis would suggest that individuals should hold more housing than a simple portfolio model would have predicted—perhaps partially accounting for the large role that it does play. There are several other important factors that enter into the decision concerning housing—which then have important implications for other aspects of an individual’s investment strategy. Because of problems with the rental market—partially related to issues of moral hazard, partially related, in some countries, to legal protections for renters—even with tax preferences, there can be distinct advantages to home ownership. Moreover, there is almost no equity market in housing; housing finance almost always comes in the form of debt. The result is that individual’s decisions concerning how much to invest in housing may be markedly different from what they would be in a standard portfolio model, even if there were no risks associated with relative price changes. At critical points in their lives, individuals may own more housing than an optimal portfolio allocation might suggest is appropriate, were these factors not present.⁵ The fact that they own more housing can have important implications for other aspects of portfolio allocation. Because housing demand is likely affected by factors like family size, that are of little relevance to other aspects of portfolio allocation, inevitably housing will need to be treated quite differently.

⁴ There is fact remarkably little literature exploring how this should best be done. See, for example, Joseph E. Stiglitz, “Behavior Toward Risk With Many Commodities,” *Econometrica*, 37(4), October 1969, pp. 660-667.

⁵ Two other financial factors affect investment decisions in housing. Housing is often tax preferred, leading individuals to hold more housing than they might otherwise hold. Moreover, individuals who wish to be more leveraged can borrow more easily and at lower interest rates using housing as collateral.

A third factor that has often been left out of the analysis is family structure. Family structure affects both risk and ‘insurance.’ For instance, families in which there are two wage earners have a more diversified, less risky ‘human capital’ than those with a single earner. This should affect the risks which they are willing to take. At the same time, families face more ‘risks’—more risks of an accident, of an illness, of the kinds of needs that draw upon savings prior to retirement. Still, so long as risks are not perfectly correlated, pooling together risks and risk absorption reduces the overall risk. Two earner households are better able to absorb risks. On the other hand, single earner households with two members may be less able to absorb risks.

A fourth factor which seems to be important in the determination of savings behavior is whether the individual is an entrepreneur. Such individuals often are credit constrained, and believe that investments in their small businesses yield high returns, sufficiently high that they offset the benefits of diversification. Such individuals are reluctant to put aside money in a separate savings account, because of the high opportunity costs that they see to such funds. But to the extent that they decide to put aside some money, one has to be aware of the undiversified nature of their (financial and human) capital.

We can simplify the analysis by thinking of a three period model—early work, later work, and retirement. For simplicity, we assume the individual begins with no assets; he knows his wage this period, but not next period. After his first period, uncertainty about second period wages is resolved.⁶ The individual buys a house. Unfortunately, because of the nature of the housing market, for most individuals, the consumption of housing services is directly related to the size of the investment in housing. If they are worried about being ‘overexposed’ to housing, they reduce their consumption of housing services; this is one of the ‘costs’ to getting a better portfolio diversification. But even after cutting back on housing services, individuals may end up with more housing assets than they would otherwise like to have had.

Parameterizations of this model can be used to provide quantitative analyses of both savings and portfolio allocations. But even short of such a quantitative analysis, qualitatively, there are two questions of interest: under what conditions should individuals save more later in life than earlier in life; and under what conditions should individuals allocate a larger fraction of their financial portfolio to equities. The answer to the first question will depend heavily on wage and ‘needs’ profiles. If wages are much higher in the second period, for instance, then over the life cycle, it makes little sense to constrain consumption early on even more than it is as a result of low wages. If, as in most countries, individuals have to put down substantial amounts in housing, and if there are decided tax advantages to home ownership, then more of the individual’s total wealth will go into housing than otherwise would be the case, and savings first period may be higher than it otherwise would be, to enable individuals to take advantage of these tax advantages (as well as the other risk advantages of home ownership.)

⁶ Other uncertainties, affecting both needs and assets, get resolved as one goes through life as well, such as divorce, children, etc.

Distributors and Asset Managers

Those providing products to retirees see their main objective of maximizing profits, which means providing products which will be bought. Investors want high returns, low risk, high liquidity (the ability to sell assets easily). They want to be sure that they are not being cheated, that the money that they have been promised will be there when it comes time for retirement. Unfortunately, individuals may have limited basis on which to make such judgments. That provides an opportunity for opportunism, for those who wish to take advantage of uninformed investors to do so. Historically, there have been many instances of such abuses; fortunately, some of the worst abuses have been curtailed by regulation and, in many countries, government provided insurance.

Many firms pride themselves on their reputations, and believe that their reputation—including their track record—is what sets them apart from their competitors. Their track record includes both the advice they give and the returns on the investment.

But for reputation mechanisms to really work, there has to be a *complete* track record, and as we noted earlier, individuals will not really know whether the advice that has been proffered is good or not until it is too late. Given the huge uncertainties, it may not even then be possible to judge whether the advice was misguided or the investments poor. It is not a mistake to sell someone earthquake insurance simply because the earthquake did not occur.

Unfortunately, both asset managers and distributors (and those that work for them) often seem to have incentive to increase their profits by exploiting investor ignorance and living (mildly) off reputation. The fact that individuals will find it difficult, if not impossible, to tell how good the advice of the firm is or how well managed the assets are until it is too late—if then—provides even greater scope for abuse.

For instance, an asset manager might create a large number of funds, each based on random picks. (He might tell some stories about the picks, trying to convince the clients that they were not random. He might even believe some of the stories he tells.) Randomly, some of the funds will do better than average for one year; some might even do better than average over a two year period. Some might do better than average in each of the first two years. By selectively winnowing out the losers, the asset manager can end up with a portfolio of funds with an impressive track record. But given the randomness of the selection (and evidence of absence of serial correlation in markets) there is no reason that a good track record has any relevance for future performance. But the asset manager is unlikely to reveal how he has gamed the market. And the unwary investor is likely to fall for the ‘track record.’

The recent scandals in the United States, and elsewhere, involving almost every major investment bank and several major mutual funds shows that either they believed that they could exploit uninformed investors indefinitely, or that the short run gains that they obtained while they were exploiting them, was worth the costs, should they be found out, including the fines and the loss of reputation. The explicit cynicism on the part of

investment advisers about the ease with which they could deceive investors was obviously disturbing; but so too was the perspective of those in the industry that they were doing nothing wrong, not only in providing the deceptive information, but in the allocation of IPO's.

Recent regulations and laws in the United States have attempted to address some of the most egregious abuses. It is important to reduce the scope for conflicts of interest, to circumscribe certain behaviors (such as market timing), to require disclosures (as it is sometimes put, sunshine is the strongest antiseptic.) The abuses were systemic, and will only be addressed (and even then only partially) by systemic regulations. Such regulations and laws need to be viewed as pro-business, for they will help restore confidence by market participants in the markets.

But while regulations can do something about the worst abuses, they can do little about the more fundamental problems—advice which is inappropriate, especially when certain precepts come to be well accepted. Indeed, firms have an incentive to play into conventional wisdom or common prejudices. If the conventional wisdom is that individuals as they grow old should hold a larger fraction of their portfolio in bonds, then a firm that caters to that belief is likely to do better—at least in the short run—than one that runs counter to that wisdom.

Incentives at the individual and organizational level exacerbate the problem. American students are typically graded on the curve—the top 10% get A's, the bottom 10% D's. So too the market grades on the curve. But this can provide a strong incentive to “follow the herd.” Assume the conventional wisdom is that telecom stocks are going to outperform the market. If one follows the herd, and invests relatively heavily in telecom stocks—trying to “pick winners” among them—then if that conventional wisdom proves wrong—telecom stocks perform badly—one will not be blamed: one was only doing what everyone else was doing, and one's performance will be commensurate with theirs (if one managed to pick good stocks within the telecom sector, one in fact will do better than average). But if one runs against the stream, and it turns out the conventional wisdom is correct, one will be harshly criticized. One's performance will be below par, and the market will take its toll.⁷

The separation between distribution and asset management makes addressing the potential problems more difficult. Asset managers want to produce products that please distributors, i.e., that distributors believe they can sell to their customers. Thus, if most distributors are telling their clients that they should diminish the ratio of equity to debt in their financial portfolio as the individual grows older (say according to a certain formula), then the asset manager needs to provide an ‘automatic aging product’ which shifts the mix over time—even if that portfolio mix is *wrong* for the long term well being of the individual, especially since the individual may never know that it is wrong.

⁷ For a discussion of the consequences of pay based on relative performance, see Barry J. Nalebuff and Joseph E. Stiglitz, “Prices and Incentives: Towards a General Theory of Compensation and Competition,” *Rand Journal of Economics*, Vol. 14, No. 1, Spring 1983, pp. 21-43.

Similarly, a distributor might want to provide better advice to his customers, but unless there are products which ‘match’ his advice, he is unlikely to do so.

This problem is referred to as one of ‘coordination failure.’ It is thus imperative that an asset manager trying to create products and advice that better serves the long run interests of workers do so in close collaboration with distributors.

The open question is: is it good business to do well? Can an asset manager that devises products that better meet the long run needs of the customers find a niche in the market?

I believe that while most individuals may not be able to follow the complicated mathematics of intertemporal utility maximization under uncertainty, they can understand the simple maxims which can be derived from this framework. They can understand what is wrong with some of the maxims that have become standard fare in the field if those errors are well and simply explained.

This is especially so when there are other responsible intermediaries between the worker and the asset manager/distributor—when the firm or the union must decide on a retirement program. (This introduces still a third level in the agency relationships, but I shall ignore the problems with this agency relationship, and assume that the firm or the union acts in the best interests of the worker.)

A Strategy for Responsible Investment and Advice

Given that individuals cannot base their savings/investment/portfolio decisions on experience—or on the experiences of their elders and peers—advice and guidance is particularly important. It is especially important given the complicated calculations and the large amount of information which are, in principle, required to make an optimal allocation. What we have already said provides the core ingredients of a strategy for responsible investment and advice.

First, we need to identify the major ways in which individuals differ from each other that are relevant to the investment/savings decisions, which include (a) home ownership; (b) entrepreneurship; (c) family structure; (d) age; and (e) occupation and other factors related to the ‘riskiness’ of human capital.

On the basis of this information, we form a set of key “prototypes.” A simple set of questions should allow the identification of which prototype best describes any given individual. The number of prototypes should be chosen to be the smallest number that still captures a large fraction of the variability across individuals.

Within each prototype there may still be large differences, e.g. in their willingness to take risks. Individuals may have perceptions about whether they are riskier than average for someone of their prototype; but these perceptions may or may not be accurate. One can attempt to verify these perceptions by looking for behavioral indicators of willingness to take risks, e.g. purchases of insurance, willingness to gamble.

Much of the “testing” of the *reasonableness* of the prototypes can be done by simulation. For example, one could take representative utility functions (e.g. logarithmic utility functions) which reflect commonly observed degrees of risk aversion, and see how (given observed risk characteristics of human, financial, and housing capital) they lead to different patterns of savings and portfolio allocation for different prototypes over time. Different prototypes will presumably have different lifetime investment strategies.

Furthermore, one can contrast these ‘normative’ models against observed patterns: do those who have invested more in housing or those with riskier human capital systematically invest more or less of their financial assets in equity? If they do not, it does not necessarily mean that the ‘model’ is inappropriate. It may only indicate that there is an even greater need for advice. But while we should not expect all individuals to respond in the way that the normative model predicts, we should hope that at least on average there is some tendency to respond in the way predicted. One might see different behaviors in different countries, or even within the same country in different regions—differences that cannot be accounted for by differences in economic and demographic conditions.

A critical question arises—what is the best way to provide the advice, the guidance? Here it is natural to take advantage of another systematic aspect of observed behavior—a characteristic, too, which can be abused. Even when given a choice, if one choice is singled out as a default, individuals tend to pick that choice. Thus, for instance, if individuals are given three saving rates, they will tend to pick the one that is chosen as the default.⁸ Confronting randomly chosen individuals with a low default saving rate will result *on average* in lower savings, even if the choice set is exactly the same. This is consistent with the observation made earlier that individuals often base their behavior on what they view as *norms*.

But what are appropriate norms for one group may not be for others. We have identified earlier several of the critical factors that are likely to be important in determining appropriate savings rates and portfolio allocations. Empirical evidence on savings and portfolio behavior of different groups would verify that these are indeed important determinants of behavior.

For a variety of reasons, the number of subclasses into which individuals should or can be divided will inevitably be limited. The objective information will not fully identify all of the relevant characteristics. Both because of information asymmetries—the investors should have more information available than the distributor/asset manager, at least about key attributes—and because individuals value the *right* to choose, individuals should be given choices; but their choices should be oriented around norms/defaults that provide benchmark choices for average individuals with the given characteristics.

⁸ See, for example, James J. Choi, David Laibson, and Brigitte C. Madrian, "Passive Decisions and Potent Defaults", in *Analyses in the Economics of Aging*, David Wise, ed., Chicago: University of Chicago Press, 2005, pp. 59-78.

In providing a set of choices around this benchmark, there are at least two approaches. One identifies say three variants—how an individual of the given prototype following the conventional wisdom would behave; how an individual who is somewhat more risk averse would behave, and how an individual who is somewhat less risk averse would behave.

The second approach would be to provide “dials”, e.g. fraction of the portfolio invested in equities. The default would be set as the ‘optimal’ number for the prototype, but individuals would be allowed to dial up or down. Once the system is going, information could be provided about the distribution of choices, e.g. 50% choose the benchmark, 30% chose a higher fraction, 20% a lower fraction, 10% choose a number below xx, and so forth.

In this approach, one might identify not only the benchmark (the percentage allocation to equity based on reasonable assumptions concerning risk aversion and risk), but also the mean and modal choices, and the choices that might correspond to the conventional wisdom, prevailing views for instance about the fraction of portfolio that should be allocated to equity as individuals get older. It will be interesting to see what kinds of information and events will lead to gradual changes in choices and norms.

A similar kind of analysis applies one stage further ‘upstream.’ Asset managers can develop a set of products, building blocks, which can be combined in a wide variety of ways to provide tailored products for each individual or prototype. Again, one of the tasks is to look for the minimum number of such products which effectively ‘spans’ the range of desired combinations, identified not just by, say, means and variances, but by correlations, e.g. with the housing market, wage levels, etc. In the traditional mean variance model, ‘dialing up’ simply means choosing the fraction of financial wealth invested in the ‘market’ portfolio of equities, and the asset manager simply provides two funds, the fund with the minimum variance, and the fund representing the ‘market.’ Once the additional factors that we have defined are taken into account, a larger number of ‘funds’ is required.⁹ Simulation exercises can be conducted to identify how close to the ‘efficient’ frontier one can go, given a range of prototypes, with say 5 or 6 well chosen funds.

There is another problem that asset managers need to address. We have repeatedly noted that the optimal mix of financial assets will change over an individual’s life, as a result of aging, changed family needs, changes in the riskiness of human capital, changes in wealth, changes in housing ownership, etc. At each stage of an individual’s life, we can think of him as solving a complicated dynamic programming problem, taking into account the probability of a transition from one state (say his current prototype), into another. It would seem both convenient and efficient (economizing on transactions costs) if the ‘product’ provided to the individual corresponded to the natural (average) evolution of someone of his prototype and age. Thus, an individual of prototype A of age 36 would have a different proportion of wealth in equities than an individual of prototype A of age

⁹ See, for example, David Cass and Joseph E. Stiglitz, "Risk Aversion and Wealth Effects on Portfolios with Many Assets," *Review of Economic Studies*, Blackwell Publishing, vol. 39(3), 1972, pp. 331-54.

37; but the individual would not have to tell the fund to reallocate his portfolio. It should happen automatically. The investment vehicles need to have sufficient flexibility that at any point, the individual may ask to have the proportion in equities dialed up or down; and it should be easy for an individual to switch 'plans'; as he sells his house, he may change from being prototype A age 36 to prototype B age 36, and this will require a portfolio reallocation. Again, for such changes to be effected easily and with little cost, each portfolio should consist of relatively few elements which can quickly be rebalanced as life circumstances change.

Further Elaborations

We have focused extensively on identifying key characteristics of individuals which do, or should, affect their portfolio decision. There is another in which individuals may differ from each other: their perceptions concerning the future. There is considerably uncertainty. Historical data are of only limited relevance for predicting the future. Some individuals may be more optimistic, others more pessimistic. Some may believe that the equity premium is likely to diminish over time, others that it will remain the same, others may not understand what is meant by the equity premium. Some see the passing of the baby boomer generation as having a major effect on housing prices, others are less certain. A further elaboration of the "model" would help individuals think through the consequences of various future scenarios of the economy.

Concluding Remarks

Few decisions are of more moment for the long run well being of individuals than how they plan for the retirement—the amount of money they save for their retirement and how they manage that money. Putting too little aside, or investing it in risky ventures which do not turn out well will force major cutbacks in standards of living during retirement; in some cases, individuals may even be forced to turn to the state or their family for maintaining a barely acceptable standard of living. The dangers that arise from excessive prudence are not so dire; but they mean that the individual and his may have sacrificed unnecessarily during productive working years.

Normal economic forces that would ensure that firms provide products that maximize the well being of individuals may not work well. There is a risk that firms may prey on individuals misperceptions, including on systematic biases in their judgments.

This paper is predicated on the belief that socially responsible financial firms will, nonetheless, not only do well, but in the long run prosper. Such firms will help individuals think through more clearly the uncertainties they face and the consequences of their decisions, and design products that enable them to adapt to their changing circumstances in ways that reduce transactions costs. This paper has provided an ambitious work program: it has outlined some of the key choices that need to be faced in designing such products and some of the key information that has to be obtained in providing appropriate advance. It has also outlined a research program that will help financial firms and advisors wishing to design retirement products and programs that will

be more responsive to the needs and preferences of individuals. It is a program which will not be easy, but the potential benefits are truly enormous.