

The role of the employer default allocation in defined-contribution retirement plan design

Chester S. Spatt,
Carnegie Mellon University
and TIAA Institute Fellow

Abstract

This paper examines the default asset allocation chosen by employers for those employees who do not select their own customized asset allocation in 401(k) and 403(b) plans. We show that the optimal default allocation entails a risky asset allocation rather than allocation of the funds to cash or a money market fund. We highlight the selection of the default allocation versus a customized allocation by the employee, examining the dependence of the employee's use of the default allocation versus a customized allocation as a function of his risk aversion, funding in the employer account and his cost of implementing the customized allocation (sophistication). Much of the use of the default asset allocation is by relatively less experienced employees with modest funding in the employer's account.

Tepper School of Business, Carnegie Mellon University and the National Bureau of Economic Research. The author gratefully thanks Jim Poterba for helpful conversations as well as research support from the Teachers Insurance Annuity Association of America (TIAA) Institute.

Any opinions expressed herein are those of the author, and do not necessarily represent the views of TIAA, the TIAA Institute or any other organization with which the author is affiliated.

1. Introduction

An important feature of most current employer 401(k) and 403(b) retirement plans is the presence of a “default” provision, which specifies the asset allocation for employees who do not select an explicit allocation and ensures the employee’s participation in the retirement plan.¹ The default provision assigns the same default asset allocation to all employees in a relevant class, such as the same age, who do not directly designate their own asset allocation.² Many individuals do not understand the nature of risk and risk-bearing and, consequently, find it challenging to determine or to implement their own asset allocation. The presence of a “default” option finesses aspects of this challenge by assigning and implementing an asset allocation for the employee’s retirement plan assets without requiring explicit decisions by the employee. Instead, an employee can utilize the default portfolio chosen on his behalf by the employer. This would be very useful in those instances when the employee recognizes the limitations of his own skill relative to that of his employer. In some situations, the employee (sometimes mistakenly) assumes that he knows less about his desired exposure than does his employer. While paternalism by the employer can be useful, it also can be problematic. It is not an unreasonable presumption that an employer would possess greater expertise than many of the employees interested in a “default portfolio,” even if it does not possess greater expertise than its more sophisticated employees or even its average employee. Those employees who possess limited expertise are likely to presume greater sophistication by their employers and are most likely to rely upon their employers.

For some employees, the costs of choosing or implementing the individual’s portfolio looms very large,

and so the use of a “default” portfolio (without any or little personal costs) can be optimal. But at the same time, there are a number of disadvantages of using a “default” portfolio as the portfolio does not reflect the individual’s preferences, such as risk aversion and intertemporal preferences, and characteristics, such as his sophistication, financial wealth, human capital, past service, the mix between taxable and tax-deferred funds, and perhaps in some instances even the individual’s age (even when the default portfolio reflects the employee’s age, it may do so in a different manner than would the employee, such as when the employee desires to retire earlier (or later) than assumed by the designer of the default portfolio, and would therefore have a different optimal target-fund than what would be reflected in the default option, for example).

Furthermore, the presence of a default portfolio would encourage use of it (compared to individual allocation decisions) and inhibit the extent to which the employee improves his investment decision-making expertise, rather than learning to sort through the relevant risk-sharing issues. The official sanctioning by the employer of the default portfolio and the manner in which it substitutes for (or pre-empts) the individual’s choice undercuts the incentive for the employee to develop expertise about lifelong financial security. In this sense, the presence of a default portfolio (and especially a highly suitable default portfolio) is a barrier to the individual customizing the portfolio and a barrier to learning by the investor. While for some employees there is a substantial direct benefit to the use of a “default” portfolio, for others the default portfolio can have adverse indirect effects. The nature of widely used default portfolios in retirement plans has shifted from cash or money market funds to explicit risky asset allocations, such as a target-date fund (mix of risky and riskless assets) designed for the investor’s age by an

¹ It is interesting that the use of a default portfolio arises in many countries and is not confined to defined contribution plans in the United States. For example, discussion of default options is provided by a) Inkmann and Shi (2016) for Australia, b) Dahlquist, Setty and Vestman (2018) for Sweden, and c) Byrne, Blake, Cairns and Dowd (2007) for the United Kingdom.

² One could imagine an alternative type of default portfolio (e.g., for retirement investing) that would be dependent upon the split of the individual’s taxable and tax-deferred wealth, provided that the employer had access to this information and could implement that. In contrast, the investor’s age is directly in the employer’s information set, so that facilitates conditioning the default allocation upon the investor’s age. I highlight portfolios dependent upon the split of wealth between taxable and tax deferred because of the central role that this split of wealth plays in asset location and overall asset allocation (see the analysis in Dammon, Spatt and Zhang (2004)).

asset manager or the employer given the retirement age and investor's risk/asset allocation preferences.³ The risk allocation in such target-date funds, as well as other default allocations chosen by employers, tends to decline with the investor's age. The change in the underlying default allocation from a riskless investment (such as money market funds or cash) to a risky investment (such as a target-date fund) reflects a desire by the employer to reduce the costs to employee investors of not bearing any exposure to risk. Of course, the extent to which the employer possesses relevant expertise for setting this is ambiguous.

Indeed, an important challenge confronting employees is to build their expertise in asset allocation and managing their funds. The presence of a default allocation, especially one that seems credible, can discourage investors from developing this crucial expertise as these appear to be substitutes. Yet the development of such expertise is essential for many participants given the importance of the funds to most plan participants and the heterogeneity in views about asset allocation among these participants (so one cannot rely upon the default portfolio). Of course, improvements in the default allocation (increasing its desirability for many participants) will reduce the likelihood that employees enhance their expertise and decision making about asset allocation within the tax-deferred account.⁴ In effect, a more desirable default allocation serves as a substitute for an increase in employee efforts and sophistication. This is an important consideration that has received insufficient attention in discussion about the use of a default allocation and setting the actual default allocation. This suggests a sense in which there can be important "unintended consequences" associated with the use of a default portfolio or improvements in the attractiveness of the default for most investors.

Along related lines, it is worth noting that the presence of a default portfolio in the tax-deferred account can decrease, as well as increase, the incentive to contribute to the tax-deferred program. The plan design influences participant decisions and reflects how these plans have evolved over time. An interesting illustrative example of the former in a different setting is provided by Choi, Laibson, Madrian and Metrick (2004), which documents that automatic enrollment at a base contribution level actually *reduces* the contributions of many participants.⁵ For example, some participants respond to the positive base (default) contribution as suggesting that level provides an adequate or almost adequate level of retirement plan funding and so contribute that amount (or a modestly higher level) rather than a substantially higher one that they would have otherwise undertaken. This illustrates the motivation for our focus upon the role of the employer's default baseline in the plan and why it can lead to distortions—and in this instance, even a decline in employee savings. The example highlights that paternalism can be problematic and, indeed, even discourage the very behavior (such as savings) that it sought nominally to encourage. Of course, automatic enrollment at a base contribution level also can increase the incentive to contribute to the program and increase participation by individuals who would not otherwise have participated. A special case that illustrates such incentives is the case of a matching contribution in which at least some of the employee's contribution is matched subject to a cap, which—due to the additional reward—can heighten the incentive of employees to contribute to the plan.

We offer a basic perspective about the use of a generic default asset allocation chosen by the employer versus a customized asset allocation chosen by the employee and show that the optimal default allocation is a risky

³ Spatt (2017) interprets target-date funds as providing a basis that spans potential risk allocations. Furthermore, Spatt (2017) also shows that the Capital Asset Pricing Model (CAPM) is equivalent to target-date funds being on the mean—standard deviation frontier. This provides an underlying foundation for the use of target-date funds that does not require that the investor optimally purchases the target-date fund designed for the investor's specific age. This strengthens the foundation for target-date funds as it does not require that the designer chooses the target date-fund optimally.

⁴ The optimal investment of funds in tax-deferred accounts in the presence of taxable investing is explored in the context of asset location (what to invest in taxable versus tax-deferred accounts). The foundation of optimal asset location is developed in Dammon, Spatt and Zhang (2004), and the implications for asset location are discussed further in Dammon, Poterba, Spatt and Zhang (2005).

⁵ See Choi, Laibson, Madrian and Metrick (2004), p. 95.

portfolio in Section 2. Section 3 introduces a formal framework to further examine the selection of the generic default versus customized asset allocation and the selection's dependence upon the risk aversion, funding and sophistication of the employee. The dependence of the choice between the generic default and customized asset allocation is further explored in Section 4. We conclude in Section 5.

2. Default versus individual asset allocation: A basic perspective

While investors are heterogeneous in their risk preferences and desired portfolio allocation, there is a broad recognition that only owning riskless assets is not optimal for many (or perhaps any) investors. Some of this view reflects the substantial historical realized returns on equity, including the presence of an “equity premium puzzle,” suggesting that realized equity returns have exceeded the benchmark returns implied by relatively simple frictionless models. The optimality of positive equity holding by (all) investors can be rationalized by a number of model frameworks. First, consider a risk-averse investor solving a portfolio problem with a riskless asset and a single risky asset (portfolio). Then, as long as the expected return on the risky asset exceeds the risk-free rate, the optimal holding of the risky asset is positive as the risk-averse investor is locally risk-neutral when he holds a zero amount of the risky asset (and so would hold optimally at least some risk, since the risky asset offers a higher expected return). An alternative (equilibrium) perspective that points to the optimality of positive holdings of the risky asset is that as long as the aggregate supply of the risky asset is positive, then optimal risk sharing suggests that in equilibrium *all* investors should hold positive amounts of it.⁶ Given the conclusion that the optimal allocation of risky assets is positive for all investors, the default investment portfolio should involve optimally holding a positive amount of the risky asset rather than owning only the risk-free asset. In effect, this provides a theoretical foundation for the

default portfolio not being invested exclusively in the risk-free asset in many employer plans.⁷

Proposition 1

If the expected return on the risky asset exceeds the risk-free rate, then the optimal holding of the risky asset is positive for all investors. Hence, the “optimal default” investment portfolio is risky.

This is not to suggest that every risky portfolio (or even every efficient risky portfolio, such as a combination of the market portfolio and a riskless asset) would be beneficial for the default portfolio relative to the riskless asset, but rather that at least modest inclusion of risky assets would be beneficial for all participants (and potentially, more substantial inclusion of risky assets would be beneficial for many).

Proposition 1 does highlight the improvement that emerges when retirement plans switch their default allocation from a money market fund or cash to a risky portfolio. The greater attractiveness of the default portfolio would imply greater use of the default option after the movement of the default choice from the riskless portfolio and then less focus upon selecting a customized portfolio. An interesting prediction to explore is whether and when this arises in practice (and even how it depends upon the specific default portfolio). Because retirement plans have changed their default portfolio at a variety of dates, it seems plausible that one could use a difference-in-difference approach to tease out the effect of the change from the default portfolio being riskless to it being somewhat risky upon the use of a customized portfolio and the choice of it. As the change in default portfolio occurred at different times, it would appear straightforward to control time trends in the use of the default portfolio.

In assessing the potential benefits and consequences of a default portfolio selected by the employer, it is important to understand how the employer would

⁶ For example, under the CAPM, individual investors would not sell short the risky market basket in light of the risk premium for the market portfolio and market clearing.

⁷ The analysis does not provide a direct explanation for the change in the default portfolio in many employer plans (going from riskless investing to a risky allocation) because it does not account for the prior use of the risk-free default allocation.

determine the default portfolio and which employees would be most likely to select it. At a minimum, in the presence of considerable heterogeneity, we would not expect all individuals to select the default portfolio (unless investors were identical). Indeed, if all individuals selected the default portfolio, that could not reflect the full diversity in employee investor circumstances if there were considerable heterogeneity.

Individual employees differ in many ways that would be relevant to their investment decisions in the tax-deferred account, including their age,⁸ sophistication and costs of decision making, risk aversion and wealth (including the split between tax-deferred and taxable wealth). This raises an interesting question: Which investors would be most interested in departing from the default portfolio selected by the employer? The employer's choice of a default portfolio should not reflect the full distribution of employee investor types, but rather those who would select it and avoid the costs of implementing a customized choice. This highlights the importance of "selection" as to the employees who choose the default portfolio versus making a customized choice.

Certainly, sophisticated investors would feel that they could make a more appropriate overall asset allocation selection/choice. Investors with relatively more wealth in their employer tax-deferred account would be less likely to delegate the decision to the employer (at least for a given level of outside wealth, a given age and extent of past service). That's because the decision would be of relatively more consequence to them. Analogously, we would expect that relatively higher income individuals would tend to be more proactive in their allocation choice (due to larger absolute amounts in the employee's account and due to greater sophistication and expertise on average). While higher income individuals have greater value to their time and to their human capital,

nevertheless they should be more likely to be proactive and less likely to use the default allocation (at least absent wealth outside the employer plan) as the time required for choosing would seem relatively modest for higher income individuals (at least to make a basic portfolio decision). A mitigating factor for higher income individuals is the extent to which they possess other resources that would not be affected by the choice of portfolio in the employer plan and which would reduce the relative importance of the plan assets.

Of course, risk aversion has a major impact on asset allocation. Conventional theory teaches that the less risk averse the investor, the greater the holdings of the risky asset relative to the riskless asset. For example, under constant relative risk aversion (either power utility or log utility), the more risk averse the investor (i.e., the greater the coefficient of relative risk aversion), the smaller the fraction of his portfolio that he allocates to the risky asset. Under what circumstances would this investor then be willing to rely upon the default portfolio? When the portfolio desired by the individual participant is close to the default portfolio (so that it is not worth the cost of customizing the portfolio), then the participant would rely upon the default portfolio. In effect, if the coefficient of relative risk aversion is relatively low or relatively high, the individual will implement his desired (customized) allocation within the employer plan. The resulting coefficient of relative risk aversion cutoffs depend upon the other parameters, such as the amount that the individual is investing through the employer plan.

There also is an interesting dynamic to choosing an actual individual allocation rather than relying upon the default allocation. The relevant decisions are to a degree long-run decisions (even though easily changed) rather than just one-time decisions; hence, a decision may be very significant for someone with a small current balance (who recently started employment, for example) due to

⁸ The standard target-date fund formulation for the default portfolio would take into account the investor's age (though not necessarily in the manner desired by the investor). The importance of age for the default portfolio is highlighted by the analysis of Inkmann and Shi (2016). As the investor ages, the proportion of risky assets in the default fund should decline. They also show that given the sensitivity of the risk allocation to age, plans with considerable age dispersion but an inability to condition the default allocation upon age should find that relatively more of the investors would choose customized portfolios. The authors find that the behavior of a panel dataset of Australian defined contribution plans is consistent with both of these hypotheses.

the future cumulative effects. Still, we would expect that younger individuals (who would have smaller balances and less experience) would be more likely to rely on default allocations. Furthermore, the cumulative aspect of these decisions suggest that once individuals make an active asset allocation, they are more likely either to continue those decisions or make new decisions after changing employers if the incremental cost associated with a proactive allocation decision is diminished after the individual has already made such decisions.⁹ These hypotheses reflect a variety of implicit costs to decision making. Collectively, these highlight that the use of the default portfolio should be much greater for young workers and workers early in their professional career. This reflects limited wealth in the employer plan, limited sophistication by these individuals and that much of the costs arise from the initial allocation decision rather than a sequence of asset allocation decisions over time.

3. Formal framework

For simplicity, we will assume initially that all of the employee's wealth is invested through his retirement plan and that the investment decision covers a static one-period problem in which the investor's funds arise in a single account. The employee investor has wealth W in this retirement plan; the asset allocation in the retirement plan is set by the employee investor at a cost c —unless the investor chooses to adopt the default allocation as structured by the employer. The investor is assumed to have a constant coefficient of relative risk aversion equal to R . We let $\alpha(R)$ denote the fraction of wealth that the employee with risk aversion R would invest in the risky asset if he incurs cost c , and α_{employer} is the fraction of wealth in the risky asset in the default portfolio selected by the employer, which is known by the employee.

If the investor incurs the cost c , then his optimal risky portfolio fraction, $\alpha(R)$, decreases with his risk aversion, R (this is a standard feature of the portfolio problem with a riskless asset and single risky asset under the assumption of constant relative risk aversion). When would the investor choose to incur the cost, c , rather than employ the default portfolio? He would do so when his risk aversion is sufficiently high or sufficiently low—i.e., when his optimal portfolio mix is either far above or far below α_{employer} , which itself depends upon the distribution of preferences as perceived by the employer.

Proposition 2

The individual employee investor chooses his optimal portfolio when his coefficient of relative risk aversion is sufficiently high or low and relies upon the default portfolio chosen by the employer for intermediate levels of risk aversion.

The decision of the employee investor as to whether to incur costs rather than using the default portfolio depends upon the ratio of W/c ; if there is sufficient wealth to be invested per unit of cost, then the employee investor will incur the cost and make his own asset allocation choice (reflecting his own risks), while if the wealth to be invested is modest per unit of cost, then the employee investor will rely upon the default portfolio. Fixing c , employee investors with sufficient wealth select their own portfolio mix, while investors with more modest wealth rely upon the default portfolio. An interpretation of the parameter c is that higher values of c reflect the investor being less sophisticated (so more costly for the individual to select his portfolio).

⁹ An important regularity in asset allocation data is that individuals rarely switch their active allocation of new funds or rebalance existing retirement plan investments (Ameriks and Zeldes (2004)). This would be consistent with only modest incremental benefits arising from new allocation decisions so that it would not be worthwhile to incur the costs of implementing proactive changes. Indeed, financial theory highlights that in equilibrium, the benefits of active portfolio rebalancing are limited when price and valuation changes do the adjusting. For example, in a situation in which all of the investment is in the market portfolio, there is no reason to adjust the asset allocation because the individual would continue to hold the market portfolio.

Proposition 3

The individual employee investor relies upon the default portfolio chosen by the employer if the individual's wealth in the retirement plan or sophistication is sufficiently low ($W/c < k^*$) and otherwise chooses a customized portfolio. The individual selects the default portfolio when W/c is below the critical value, k^* .

Of course, the composition of employee investors who the employer perceives should select the default portfolio influences how the employer selects the appropriate allocation for it. The selection of the default portfolio by the employer should reflect only those employees who will use the default (of course, the composition of the default portfolio may influence those on the margin of selecting the default portfolio versus a customized portfolio). This highlights that the employer should be especially focused on setting the default for those with relatively modest funds and those who are relatively less sophisticated (high cost, c) as these will be the employees who utilize the default portfolio.¹⁰ In effect, the "selection effect" suggests a foundation for a paternalistic focus by institutions on those with modest funds for designing the default portfolio.

An additional point to highlight is that the use of the default allocation in the employer account would decline over time, assuming that the cost structure of choosing and implementing an active portfolio would decline over time and the wealth being managed through retirement contributions grows over time, so that use of the active choice by the employee would increase over time.

Proposition 4

The individual employee investor is more likely to rely upon the default portfolio chosen by the firm during his early years with a firm.

Much of the use of the default asset allocation is by relatively less experienced employees with modest funding in the employer's account. As the employee amasses assets in the default allocation and potentially

increases his sophistication, we would expect that the employee would substitute to a customized portfolio. The incentive to incur the costs to reallocate his investment fund increases with the employee's retirement wealth, sophistication, experience and age.

4. Default versus individual asset allocation: Further perspectives

In the formal analysis in the prior section, we did not explicitly condition upon the investor's age. For a variety of reasons, including the extent of future human capital and the remaining horizon over which the individual plans to spend his resources, the individual's optimal portfolio allocation would depend upon his age. On the other hand, the target-date funds approach also can lead the default portfolio to depend upon the investor's age, though in a particular manner that may not line up with the specific preferences of the individual. The dependence of age in the target-date funds approach may not align so closely with how the individual employee investor conditions upon age in light of the individual's specific preferences, which would reflect his anticipated retirement age and the nature of the investment horizon that he anticipates (including the extent to which he is investing indirectly on behalf of his heirs). In this sense, the investor's age would potentially influence whether the individual chooses to make his own customized portfolio determination rather than rely upon the default portfolio.

Another important aspect in practice governing the possible use of the default portfolio is that such a structure would only apply to the employer's 401(k) and 403(b) plans and not to either other tax-deferred retirement plans or the employee's personal taxable funds. The employee investor should optimally first hold equity in his taxable accounts and riskless assets first in his tax-deferred accounts (see Dammon, Spatt and Zhang (2004) and Dammon, Poterba, Spatt and Zhang (2005) for related discussion on asset location in taxable and tax-deferred accounts). The discussion here suggests that similar comparative statics should obtain with respect to the use of the default portfolio in the

¹⁰ However, the employer may weigh relatively more those with relatively larger accounts (due to their larger investments) in situations in which the employer perceives they would actually select the default portfolio.

employer account, taking employer accounts as given. One important caveat to the earlier conclusion that the optimal design of the default portfolio is risky is that it would be optimal to hold only riskless assets (bonds) in individual tax-deferred accounts if the employee has a sufficiently low fraction of his overall wealth in the tax-deferred account.

Indeed, in this spirit, Dahlquist, Setty and Vestman (2018) document considerable empirical heterogeneity among Swedish investors, suggesting “that it may be beneficial to carefully design the default fund to suit each investor’s specific situation rather than imposing one allocation on all.” They argue that the asset allocation in the default portfolio should condition on more than just the investors’ age. This is consistent with the analysis in Dammon, Spatt and Zhang (2004), in which the split of wealth between the tax-deferred and taxable accounts plays a central role in asset location and allocation. The evidence in Dahlquist, Setty and Vestman (2018) suggests that passive investors tend to be less educated, have lower wealth and labor income, and are less sophisticated on an overall basis.

5. Concluding comments

The employer’s default portfolio allocation influences which employees choose to bear the costs associated with determining a more customized asset allocation in his retirement plan. Our analysis offers several important insights, including explaining why the optimal default allocation is not a riskless allocation; why the optimal default allocation should not reflect the full joint distribution of employee characteristics but instead those who are anticipated to select the default portfolio, such as employees who are relatively new and have modest plan accumulations; the nature of systematic differences over which employees will choose a customized allocation and which employees rely upon the default allocation; and why improvement in the default allocation can damage the individual’s ability to manage his retirement funds over time.

References

- Ameriks, J. and S. Zeldes, 2004, "How Do Household Portfolio Shares Vary with Age?", unpublished manuscript, Vanguard Group and Columbia University.
- Byrne, A., D. Blake, A. Cairns and K. Dowd, 2007, "Default Funds in UK Defined Contribution Pension Plans," *Financial Analysts Journal*.
- Choi, J., D. Laibson, B. Madrian and A. Metrick, 2004, "For Better or for Worse: Default Effects and 401(k) Savings Behavior," in *Perspectives on the Economics of Aging*, David Wise (ed.), University of Chicago Press.
- Dahlquist, M., O. Setty and R. Vestman, 2018, "On the Asset Allocation of a Default Pension Fund," *Journal of Finance*, forthcoming.
- Dammon, R., J. Poterba, C. Spatt and H. Zhang, 2005, "Maximizing Long-Term Wealth Accumulation: It's Not Just about "What" Investments to Make, but also "Where" to Make Them," TIAA Institute *Research Dialogue* (September) 1-12.
- Dammon, R., C. Spatt and H. Zhang, 2004, "Optimal Asset Location and Allocation with Taxable and Tax-Deferred Investing," *Journal of Finance* 59, 999-1037.
- Inkmann, J. and Z. Shi, 2016, "Life-Cycle Patterns in the Design and Adoption of Default Funds in DC Pension Plans," *Journal of Pension Economics and Finance* 15, 429-454.
- Spatt, C., 2017, "Target-Date Funds, Annuitization and Retirement Investing," TIAA Institute *Research Dialogue* No. 134, May.

About the author

Chester S. Spatt is the Pamela R. and Kenneth B. Dunn Professor of Finance at the Tepper School of Business at Carnegie Mellon University, where he has taught since 1979. He served as chief economist of the U.S. Securities and Exchange Commission and director of its Office of Economic Analysis from July 2004 through July 2007.

Spatt is a well-known scholar studying financial economics with broad interests in financial markets. He has analyzed extensively market structure and trading, financial regulation, pricing and valuation, and the impact of information in the marketplace. He has been a leading expert on the design of security markets in various settings, mortgage valuation, and taxation and investment strategy. His co-authored 2004 paper in the *Journal of Finance* on asset location won TIAA's Paul A. Samuelson Award for Outstanding Scholarly Writing on Lifelong Financial Security.

Spatt has served as executive editor and one of the founding editors of the *Review of Financial Studies*, president and a member of the Founding Committee of the Society for Financial Studies, and president of the Western Finance Association. He is currently an associate editor of several finance journals and a member of the Systemic Risk Council, a research associate of the National Bureau of Economic Research and a senior economic adviser to Kalorama Partners. He was one of the initial members of the Federal Reserve's Model Validation Council and also previously served as a member of the SEC's Equity Market Structure Advisory Committee, the Advisory Committee of the Office of Financial Research and the Shadow Financial Regulatory Committee. He also is a member of the Financial Economists Roundtable and a Fellow of the TIAA Institute. He earned his Ph.D. in economics from the University of Pennsylvania and his undergraduate degree from Princeton University.