The Crowd-out Effects and Welfare Implications of a Retirement Savings Nudge*

Derek Messacar†

November, 2015

Abstract

This paper investigates whether a savings 'nudge' is a viable mechanism to help people overcome knowledge barriers associated with saving for retirement. Using administrative tax records from Canada, I begin by estimating whether a nudge in workplace pensions increases net savings or induces workers to reduce contributions in other accounts. The analysis controls for correlations in savings across accounts due to unobserved tastes for saving by exploiting the fact that employer contribution rates increase discontinuously on earnings above the average industrial wage, a unique feature of occupational pensions in Canada, the effect being estimated in a regression kink design. I find that a $1 increase in employer contributions crowds out other savings by approximately $0.50. Hence, a nudge raises net savings on average but many workers actively respond to it, which leads me to consider underlying heterogeneity in more detail. Specifically, the effects of education on savings adjustments to a nudge and savings behavior are examined, using compulsory schooling reforms in the identification. I find that workers with low education save inadequately for retirement on their own but benefit from a nudge by remaining passive, whereas those with high education are forward-looking and also respond properly to the distortion by re-optimizing savings across accounts. These results indicate that a nudge likely increases private welfare in practice. However, education reform is also a viable policy intervention for raising savings, a finding that should be considered when deciding on socially optimal policies designed to help people save.

JEL Codes: E21, H30, J32, I21
Keywords: Pensions, Nudge, Crowd-out, Education, Welfare, Regression Kink Design, Instrumental Variables

*I am grateful to Robert McMillan, Kory Kroft, and David Seim for their support on this project. I also thank Philip Oreopoulos, Serdar Ozkan, Laura Turner, Mike Veall, participants of the University of Toronto CEPA Seminar and Public Economics Workshop, and members of the Social Analysis and Modelling Division at Statistics Canada for the discussions, comments and advice on earlier versions of this work. I am indebted to Winnie Chan, Marc Frenette, René Morissette, Brian Murphy, Yuri Ostrovsky, Hung Pham, Paul Roberts, and Grant Schellenberg at Statistics Canada for their assistance. Any errors are my own.

†Ph.D. Candidate. Department of Economics, University of Toronto, 150 St. George Street, Toronto, Ontario, Canada, M5S 3G7. E-mail: derek.messacar@utoronto.ca.

1 Introduction

Private savings rates in some countries have declined sharply in recent years, raising concerns about the future retirement prospects of today's workers (de Serres and Pelgrin, 2003; Harvey, 2004). The traditional approach of offering tax incentives on assets held in designated accounts is often regarded as an ineffective way to boost wealth (Attanasio et al., 2004). Instead, many economists now advocate greater use of programs that nudge people in the direction of saving more. A growing literature in behavioral economics finds that workplace pension plan features such as default options or savings rate escalators raise contributions within workplace plans (Madrian and Shea, 2001; Thaler and Benartzi, 2004). Nudges are increasingly viewed as effective ways to influence the savings of individuals who prepare inadequately for retirement, while permitting those who wish to opt out of such arrangements the chance to do so (Thaler and Sunstein, 2008; Iwry and John, 2009). Recent evidence also indicates that automatic contributions increase net savings, even controlling for crowd-out responses across savings accounts (Chetty et al., 2014).

Although a nudge effectively shifts consumption from working years to retirement, the life-cycle consumer welfare effects of such a distortion remain unclear. For example, the costs of inducing some people to over-save can be as great as the consumption losses to others in retirement from myopia (Whitehouse, 2013). Whereas rational agency predicts that agents would respond to a nudge by actively adjusting contributions within or across savings plans instead of remaining passive to the distortion, passive behavior may reflect a wider problem that individuals do not understand how to save adequately on their own. In this case, a nudge may be a viable policy instrument for correcting this behavioral problem and raising welfare. A better understanding of the relationship between how individuals respond to a nudge and how they save for retirement has implications for the optimal design of retirement income systems.

In this paper, I investigate the viability of a nudge as a mechanism to help people overcome knowledge barriers associated with saving adequately, from a life-cycle welfare perspective. In a stylized consumer model, I begin by showing that the following conditions are necessary for a nudge to increase consumer welfare: (1) passive agents must be the most likely to save inadequately, and therefore to benefit from this intervention; and (2) agents who already save adequately must be the most likely to respond actively to the distortion by re-optimizing savings across accounts *at low cost*. Most importantly, under common assumptions specified in the model these conditions can be assessed empirically, which I subsequently carry out using administrative tax records and Census data from Canada spanning 1991-2010. Whereas previous research has shown that some individuals respond less to changes in pension wealth

than others (Chetty et al., 2014), this study is the first to directly assess the interconnection of active versus passive choice and savings adequacy.

To carry out this investigation, the empirical analysis proceeds in three steps. First, I investigate whether a change in employer pension contributions increases net private savings or simply induces workers to adjust contributions in other retirement accounts. This analysis adds to a large literature examining the effects of public and employer-sponsored pensions on savings outcomes (see Bernheim (2002) for a survey). I provide credible new insight into this unresolved issue by estimating the crowd-out effect of random variation in workplace savings that closely resembles a true nudging mechanism.[1] Specifically, the identification exploits the fact that employer contribution rates tend to increase discontinuously on earnings above the average industrial wage, a unique feature of occupational pensions in Canada. The magnitude of this change in employer contributions and its displacement effect on other retirement savings are jointly estimated in a two-stage regression kink design (RKD). The results show that workplace pensions *partially* displace other retirement savings by approximately \$0.50 per \$1, indicating that a nudge raises net savings but that many workers are very responsive to the distortion.

In the second step of the analysis, I provide the first quantitative test of the behavioral mechanisms behind active versus passive choice. Specifically, I investigate the extent to which education affects such behavior, motivated by evidence that schooling correlates with financial literacy and retirement preparedness (Lusardi and Mitchell, 2010; Mullock and Turcotte, 2012; Venti and Wise, 2014). To this end, I obtain schooling measures from unique data sets that link the tax records of nearly 800,000 individuals to their 1991 or 2006 Census of Canada responses. Following the returns-to-education literature, compulsory schooling reforms are used to obtain exogenous variation in education to identify this effect (Acemoglu and Angrist, 2000; Milligan et al., 2004; Oreopoulos, 2006b). The results show that workers with high education respond to a nudge by crowding out other savings whereas those with low schooling are passive. Hence, passive behavior in part results from knowledge-based traits that are amenable to change through education policy. To the extent that education is a factor affecting the human capital costs of active choice as these results seem to indicate,

[1] Previous studies look at the crowd-out effect of automatic contributions arising from changes in: pension eligibility (Poterba et al., 1995; Engen et al., 1996; Venti and Wise, 1996; Engen and Gale, 2000; Gelber, 2011); worker mobility (Chetty et al., 2014); and mandatory public pensions (Euwals, 2000; Arnberg and Barslund, 2013; Chetty et al., 2014). These studies provide important insight into the roles of public and employer-sponsored pensions in the wealth formation process. However, these sources of variation are typically associated with wide-ranging events relative to a nuanced nudging mechanism (e.g. default option) studied in behavioral public finance (Madrian and Shea, 2001; Choi et al., 2005). The primary source of variation that I consider resembles a savings rate 'escalator' that Thaler and Benartzi (2004) study in a controlled experiment as a savings commitment device.

then workers with low adjustment costs are indeed more responsive to a nudge than workers with high adjustment costs. This finding provides new insight into the underlying reasons why nudges work—a topic that is poorly understood (Bernheim et al., 2013).

The third step of the empirical analysis is to consider the welfare implications of nudges. Since under-responsiveness to a nudge is to some extent explained by low levels of schooling, then workers with low education must also save inadequately relative to workers with high education in order for a nudge to raise consumer welfare. To test this requirement, I posit that savings adequacy can be assessed using exogenous variation in education. Controlling for the standard channels through which education indirectly affects savings, providing workers with more schooling should not alter their savings decisions if they were acting optimally in the first place. In contrast, if education were found to directly affect savings behavior, this finding would imply that workers affected by the reforms were helped in overcoming knowledge barriers to saving adequately. The results show that, on average, completing high school raises individuals' annual retirement savings rates by 4 percentage points, an effect that persists over the life-cycle.

Taken together, the findings indicate that a nudge likely increases consumer welfare in practice. Such an intervention will boost the savings of workers with low education who appear to benefit the most from the intervention. In addition, although workers with high education are more forward-looking on average, they tend not to be significantly disadvantaged by a nudge given that they properly re-optimize savings across accounts at low cost. However, these results also indicate that education reform is a viable substitute for a nudge given that both programs raise savings. More schooling reduces the need for, and effectiveness of, a nudge as workers learn to save on their own, a finding that policy-makers should consider when deciding on socially optimal programs intended to help people save.

This study contributes to three related literatures. The first addresses the longstanding issue of whether workplace pensions raise or redistribute total savings. Whereas some studies find that these plans creates new savings (Poterba et al., 1994; Venti and Wise, 1996; Gelber, 2011; Chetty et al., 2014), others find that they largely crowd out contributions in other accounts (Engen and Gale, 1994; Gale, 1998; Veall, 2001; Benjamin, 2003). The lack of consensus on this issue may arise due to limitations in the availability of reliable data or suitable research designs (Hubbard and Skinner, 1996; Bernheim, 2002).[2] Second, this study

[2]There are three identification problems that beset this literature. First, measurement error in reported pension wealth from survey data may cause researchers to overstate the extent to which workplace plans generate new savings (Engelhardt and Kumar, 2011). Second, unobserved heterogeneity in individuals' tastes for saving would bias ordinary least squares estimates of crowd-out upward, since some individuals contribute more to all types of retirement savings plans. Third, workers may sort into firms based on unobserved preferences for pension coverage, which would introduce bias in comparisons of savings outcomes between eligible and ineligible members (Ippolito, 1997). Firms may also offer pensions based on the demands of their

advances current research on the viability of a nudge as a policy instrument to raise savings (Madrian and Shea, 2001; Choi et al., 2004; Thaler and Benartzi, 2004; Carroll et al., 2009; Chetty et al., 2014). The finding that education affects active versus passive choice is also consistent with existing evidence on how bounded rationality and inertia affect savings (Sethi-Iyengar et al., 2004; Choi et al., 2011; Beshears et al., 2013). Third, this study adds to a large literature on the returns to schooling. Previous research finds that less-educated individuals have lower earnings (Angrist and Krueger, 1991), are more likely to be incarcerated (Lochner and Moretti, 2004), and are less healthy or happy (Oreopoulos, 2007) than highly-educated individuals. More schooling is also associated with better financial literacy which, in turn, is thought to increase savings (Bernheim and Garrett, 2009; Bayer et al., 2009; Lusardi and Mitchell, 2014). Although the relationship between education and retirement preparedness has been documented (Venti and Wise, 2014), no prior study has evaluated whether education affects both savings adjustments to a nudge and life-cycle savings behavior.

The paper proceeds as follows: the next section presents a stylized conceptual framework to provide a lens through which to interpret the empirical results; Section 3 reviews Canada's retirement income system, with a discussion of key components used in the empirical analysis; Section 4 describes the administrative tax data and sample selections used; Sections 5, 6, and 7 show the results on crowd-out, heterogeneity by education, and returns to education; and Section 8 concludes.

2 Conceptual Framework

This section explores the consumer welfare implications of a savings nudge. I extend a framework for active versus passive choice by Chetty et al. (2014) and Bernheim et al. (2015) to endogenize agents' costs of active choice. The model posits that how agents respond to a nudge and how they save are co-determined by their human capital investments, delivering a role for a savings nudge *only* under specific conditions.

2.1 Model Setup

Agents live for two periods, making decisions sequentially about human capital investments and then life-cycle consumption. For simplicity, I describe the model in reverse-chronological order. In the second period, an agent inelastically supplies one unit of labor in exchange for compensation comprising labor earnings E and a pension contribution P. With this income,

workers. By exploiting rich tax data and unique institutional details of the Canadian retirement income system, I am able to estimate the crowd-out effect of workplace pensions on other retirement savings in a way that resolves these identification issues.

the agent uses investment instruments to purchase a consumption portfolio $\{C_1, C_2\}$ by investing into taxable and tax-deductible savings plans, S and R respectively.[3] The investment instruments map into the consumption bundle according to $c_1 : \{E, S, R, P, r_0\} \to C_1$ and $c_2 : \{E, S, R, P, r_0\} \to C_2$, given by:

$$c_1(E, S, R, P, r_0) = E - (1 + \phi)S - R - P \tag{1}$$

$$c_2(E, S, R, P, r_0) = r_0(S + R + P) \tag{2}$$

Equations (1) and (2) indicate that C_1 equals income minus savings, and C_2 equals the wealth accumulated from the first period. Note, $\phi \in (0, 1)$ ensures R and P yield a higher rate of return than S. I impose that agents require a minimum level of taxable savings, $\bar{S} \in [0, E)$, for liquidity given that R strictly dominates S in this two-period setting.[4]

I assume that $r_0 = 1$ but that this is only known to a social planner. Agents have beliefs about the wealth accumulation process, r_0, given by the random variable r, which has the density function $p_r(r)$ such that $\mathbb{E}_r[r] = \int p_r(r)dr = 1$, so beliefs are on average correct, and $r > 0$. Since r_0 is unknown, this setup is a type of 'complexly framed' problem in which financial incompetence arises because agents do not fully understand how savings map into consumption (Ambuehl et al., 2014). Agents may invest into costly human capital $h \in \{0, \bar{h}\}$ in order to learn about r_0, which delivers a signal $s = r + \epsilon$, where $\mathbb{E}_\epsilon[\epsilon] = 0$.

The second-period consumption problem for a representative agent with human capital $h \in \{0, \bar{h}\}$ and signal $\tilde{s}$ is:

$$\max_{\{C_1, C_2\}} \mathbb{E}_{r|\tilde{s}} \left[U\left(C_1, C_2\right) | \tilde{s} \right] \tag{3}$$

subject to:

$$C_t = c_t(E, S, R, P, r) \text{ for each } t = \{1, 2\} \tag{4}$$

$$(S - \bar{S}), R \geq 0 \tag{5}$$

$$R + P \leq L \tag{6}$$

where $\tilde{s} \in \{\emptyset, s\}$ such that $\mathbb{E}_{r|\tilde{s}}[\cdot|\tilde{s} = \emptyset] = \mathbb{E}_r[\cdot]$, and L is a regulated limit for contributions to tax-deductible accounts. The following standard assumptions on agents' utility are imposed: U is twice-continuously differentiable, has positive but diminishing marginal returns to consumption, and $\{C_1, C_2\}$ are complementary.

[3] The consumption problem is collapsed into one decision period on the basis that this can be viewed as a reduced-form version of a dynamic optimization problem, following Bernheim et al. (2015).

[4] A demand for liquidity could also be modeled via a direct utility relation over taxable savings (Chetty et al., 2014), a multi-period model where tax-deductible assets lock in (Gale and Scholz, 1994), or short positions in taxable accounts (Milligan, 2002).

Human capital decisions are made in the first period. The cost of human capital is u, which is exogenously distributed in the population with density $p_u(u)$, and $u > 0$. The return from human capital is the improved knowledge of how investments map into consumption.[5] This assumption rests on recent evidence that educational attainment is correlated with financial literacy and retirement preparedness (Mullock and Turcotte, 2012; Venti and Wise, 2014), and aligns with the standard frameworks of bounded rationality (Simon, 1955) and rational inattention (Sims, 2006).

Agents' decisions are formed based on incomplete information about the savings opportunities available to them, but this uncertainty affects utility only through its effect on consumption. These features justify removing specific choices from the set of choices used to evaluate consumer welfare (Chetty et al., 2009; Bernheim et al., 2015; Ambuehl et al., 2014). Hence, the burden of this uncertainty can be viewed through the choice-theoretic framework for welfare analysis of Bernheim and Rangel (2009).

2.2 Optimal Behavior and Comparative Statics

An agent solves this problem by backward induction. In what follows, I restrict the analysis to 'interior savers' for whom equations (5) and (6) hold as strict inequalities. This consumption problem collapses into an unconstrained optimization problem by recognizing that $S^\star = \bar{S}$, and then substituting the relevant constraints into the objective function.

Let $U(R) = U(c_1(E, \bar{S}, R, P, r), c_2(E, \bar{S}, R, P, r))$ denote an agent's second-period utility from consumption as a function of the remaining choice variable R. Conditional on h and $\tilde{s}$, the optimal choice of retirement savings is:

$$R^\star(h, \tilde{s}) = \arg\max_R \mathbb{E}_{r|\tilde{s}}[U(R)|\tilde{s}] \tag{7}$$

Let $\hat{R}$ be the optimal retirement savings with perfect information about r_0. A measure of characterization failure resulting from the behavioral costs of human capital and investment uncertainty is $|\hat{R} - R^\star(h, \tilde{s})|$.

In the first period, the representative agent chooses whether to acquire human capital by weighing its cost of acquisition against the expected future return. Specifically, the acquisition rule is: $h^\star = \bar{h}$ if $u \leq \bar{u}$ and $h^\star = 0$ otherwise, where $\bar{u} = \bar{h}^{-1}\{\mathbb{E}_s\mathbb{E}_{r,s}[U(R^\star(\bar{h}, \tilde{s})|s)] - \mathbb{E}_r[U(R^\star(0, \tilde{s}))]\}$. Risk aversion is sufficient for more human capital to increase utility notwithstanding its acquisition cost, $\mathbb{E}_s\mathbb{E}_{r,s}[U(R^\star(\bar{h}, \tilde{s})|s)] > \mathbb{E}_r[U(R^\star(0, \tilde{s}))]$. In contrast, the

[5] There is no labor market return to human capital in order to focus only on the main channel of interest, how human capital affects savings. The model could easily be extended to include a labor market return, and this issue receives full consideration in the empirical analysis.

effect of human capital acquisition on savings is unclear. The uncertainty over r_0 means that agents are less willing to be exposed to the possibility of loss, which favors initial consumption through a substitution effect. However, this uncertainty also induces agents to protect themselves against very low levels of future consumption, which favors savings through an income effect. The net effect of the uncertainty is ambiguous without additional assumptions on agents' preferences (Sandmo, 1970).

This model delivers a role for policies that help people save, to overcome the investment uncertainty and knowledge barrier, u, of information acquisition—although whether such policies should raise or lower average savings is not *a priori* clear. This role arises without distinction between agents' decision utility and rational (or 'long-term') utility, as in Carroll et al. (2009). One such policy is a nudge, which is defined herein as:

Definition 1 *A savings nudge is an automatic change in pension wealth, ΔP, which occurs as agents make their second-period consumption decisions. An agent may 'opt out' of this nudge via savings in other accounts, S or R.*[6]

For interior savers, the optimal level of retirement savings in the presence of a nudge is $R^\star(h, \tilde{s}, \Delta P) = (R^\star(h, \tilde{s}) - \Delta P)$, since R and P are treated identically under the tax code (Chetty et al., 2014). Yet some agents may not respond to a nudge if the costs outweigh the associated benefits. The utility gain of responding to a nudge is:

$$G(h, \tilde{s}, \Delta P) = \mathbb{E}_{\tilde{s}}\mathbb{E}_{r|\tilde{s}}[U(R^\star(h, \tilde{s}, \Delta P)|\Delta P)|\tilde{s}] - \mathbb{E}_{\tilde{s}}\mathbb{E}_{r|\tilde{s}}[U(R^\star(h, \tilde{s})|\Delta P)|\tilde{s}] \tag{8}$$

Agents also incur adjustment costs: $k(h, q; \zeta) = \zeta f(h) + (1 - \zeta)q$, where $\zeta \in [0, 1]$ and is homogeneous across agents. This cost comprises: (1) a behavioral element, reflecting costs of attentiveness and learning about savings strategies to understand how to respond optimally, which depends on h according to $f(\bar{h}) < f(0)$; and (2) other costs q, distributed according to the density function $p_q(q)$ where $q > 0$ and $q \perp u$. An agent will crowd out the nudge if and only if $G(h, \tilde{s}, \Delta P) \geq k(h, q; \zeta)$. This discussion leads to the following definition of active versus passive choice:

Definition 2 *Given a nudge ΔP, an agent with human capital $h \in \{0, \bar{h}\}$ and signal $\tilde{s}$ is an active saver if $R^\star = R^\star(h, \tilde{s}, \Delta P)$, and is a passive saver if $R^\star = R^\star(h, \tilde{s})$.*

[6] This definition differs slightly from the conventional definition of a nudge, such as default options (Madrian and Shea, 2001; Choi et al., 2003), where the margin of opt-out is within the pension plan rather than across accounts. Nevertheless, the basic principles of opt-out and active versus passive choice are preserved. This terminology fits the current study as the empirical analysis will exploit variation in P that resembles a true nudging mechanism, as mentioned above.

Existing studies treat passive choice as fixed (Chetty et al., 2014) or assume that adjustment costs are exogenous (Bernheim et al., 2015), whereas this model permits costs to depend on human capital, offering a testable explanation for why nudges work.

2.3 Welfare Implications

Given the information asymmetry over r_0 that exists between the social planner and agents, there is room for social policies to raise aggregate consumer welfare. For an agent with human capital cost u, the *ex post* indirect utility from consuming $\{h, S, R\}$ is:

$$V(h, S, R; u) = -uh + U\left(E - (1+\phi)S - R - P, r_0(S + R + P)\right) \tag{9}$$

To a social planner, agents are heterogeneous along two relevant dimensions: (1) whether human capital is acquired; and (2) whether there is an active or passive savings response to a nudge. A planner has perfect information about r_0 but still forms expectations about realized signals s in designing social policy.

Consider an optimal savings nudge $\Delta P^\star$ that maximizes the weighted average of all agents' *ex post* expected utilities (see the Technical Appendix for a derivation). Figure 1 shows what effect this nudge would have on the expected utility of each type of agent. In particular, the nudge would strictly decrease welfare for all active agents, as they are unaffected by the intervention aside from the crowd-out cost $k(h, q; u)$ incurred. The expected effect for passive agents depends on their human capital. For passive agents with $\{h, \tilde{s}\} = \{0, \emptyset\}$, the nudge is expected to increase consumer welfare by reducing characterization failure. In contrast, for passive agents with $\{h, \tilde{s}\} = \{\bar{h}, s\}$, the expected welfare effect is ambiguous since the nudge may actually overcompensate for their remaining uncertainty about r_0.

The welfare gains from a nudge increase as adjustment costs depend relatively more on human capital, $\zeta \to 1$. For example, if $\zeta = 1$ then all active savers have low adjustment costs relative to passive savers, and the expected welfare effect of a nudge for passive savers is strictly positive, in which case a nudge likely increases consumer welfare in practice.

Interestingly, a consequence of this model is that a nudge is not the only policy instrument available to a social planner to influence savings outcomes. A reform that targets human capital is an imperfect substitute for a nudge, defined as follows:

Definition 3 *A human capital reform is a policy that changes the compulsory level of human capital that agents must obtain from $h^c = 0$ to $h^c = \bar{h}$.*

This constraint would simply be added to agents' optimization problem, inducing them all to choose $h^\star = \bar{h}$ irrespective of the cost. For agents who already choose $h^\star = \bar{h}$, the reform

is non-binding, suggesting that there is a potentially large segment of the population for whom the human capital reform has no welfare effect. For agents who would otherwise choose $h^\star = 0$, the reform is *ex ante* strictly welfare-diminishing but may still raise consumers' *ex post* expected welfare, as shown in Figure 1 and derived in the Technical Appendix. This occurs because uncertainty about the benefit from obtaining human capital results in a sub-optimal aggregate level of acquisition, from a social planner's perspective. Although a nudge and human capital reform both have the capacity to help people save, the winners and losers from each intervention may be different, which ultimately affects each program's political feasibility to the extent that policy-makers design reforms taking into account implementation costs and voter preferences.

2.4 Testable Predictions

The model posits a relation between human capital, savings adjustment, and savings behavior that can be tested empirically via the three steps outlined in Section 1, which I restate here in the context of the model. First, whereas rational agency predicts $\frac{\Delta R^\star}{\Delta P} = -1$ for all interior savers, this response should be smaller than unity in aggregate due to adjustment costs, which is recoverable if ΔP is well-identified. Second, the crowd-out response should be dichotomous across individuals based at least partly on human capital (unless $\zeta = 0$), which can be tested if exogenous variation in human capital is observed. Third, savings inadequacy can be assessed by estimating $\frac{\Delta R^\star}{\Delta h}$ provided that Δh is well-identified and indirect channels through which human capital affect savings are controlled for. If the first two test predictions hold and $\frac{\Delta R^\star}{\Delta h} > 0$, there is a role for a nudge $\Delta P > 0$ to increase welfare but a human capital reform imperfectly substitutes.

3 Institutional Details

This section briefly reviews the institutional features of Canada's retirement income system, focusing on the details most relevant to the empirical analysis. The system can be described as comprising three tiers (Baker et al., 2007).

3.1 Public Pensions

The first two tiers constitute Canada's public pension system. First, there is a citizenship-based public pension available to most Canadians aged 65 or older who meet status and residency requirements, funded out of tax revenues. As of December 2010, the maximum monthly entitlement was $522. Benefits are linked to inflation and are taxable as income.

For low-income recipients, a supplement can be available of $658 and $435 per person for single and married individuals, respectively, as of December 2010.

The second tier is an income-tested pension designed to replace approximately 25 percent of workers' mean lifetime earnings up to the average industrial wage. In 2010, the maximum monthly benefit payment was $934. Pension benefits are determined according to a complex function of workers' earnings histories, lengths of time spent contributing, and ages at which benefits start to be collected. Importantly, this program is funded out of matching employer and employee payroll deductions. Every person in Canada between the ages of 18 and 70 who earns a salary above a basic exemption is required to contribute into this plan. In 2010, workers and employers equally contributed 4.95 percent of earnings between the basic exemption ($3,500) and the average industrial wage ($47,200), and self-employed workers were required to pay both shares. However, the marginal contribution rate on earnings above the average industrial wage for both employees and employers falls to zero, a fact that I exploit in the crowd-out analysis.

3.2 Private Pensions and Retirement Savings Plans

Private pensions and retirement savings plans constitute the third tier of Canada's retirement income system. These plans are especially important for middle-income and high-income households to prevent significant drops in their living standard at retirement (Ostrovsky and Schellenberg, 2009). Employer-sponsored pensions have historically played an important role in this process. These plans can be defined benefit, defined contribution, or hybrid arrangements to provide pension income to retired employees in the form of periodic payments. Employers offering workplace pensions must contribute at least 1 percent of earnings into these plans annually, and employees also make positive contributions approximately 75 percent of the time. Contributions are non-taxable, investment income accumulates on a tax-free basis, and assets typically lock in after a short vesting period. Income from these plans are subject to normal income taxes at retirement. For a sample of frequent tax filers used in this study, Table 1 gives a comparison of demographic, earnings, and savings characteristics of workers according to whether or not they have workplace pension coverage. For example, workers with pension coverage tend to earn more and have higher savings than individuals who do not belong to such plans.

A unique characteristic of workplace pensions in Canada that I exploit in the empirical analysis is that most plans are *integrated* with the contribution schedule of the second-tier public pension. Specifically, employers usually set different marginal contribution and benefit rates for earnings below or above the average industrial wage. This occupational

pension structure exists because firms recognize the added costs imposed on them by the public pension; the majority of employers have designed their plans in this way since the means-tested public pension first came into effect, in 1966. For example, over the period 1991-2000, approximately 85 percent of workplace pension plans had integrated contributions and/or benefit formulas (Statistics Canada, 2003). Ultimately, integration means that a typical employer contribution schedule *kinks* upward at the average industrial wage since employers contribute more generously when they are no longer charged at the margin by the public pension system. This change in employer pension contributions is a plausibly exogenous source of variation with which to identify crowd-out in other savings plans given that integration is decided at the firm level.

Individuals may also save for retirement in voluntary, defined-contribution accounts that they set up and maintain through financial institutions. The maximum amount that may be contributed annually into voluntary plans is the lesser of 18 percent of income and a pre-determined threshold ($22,000 in 2010), although any unused contribution room has carried forward indefinitely since 1991. As a result, the majority of Canadians do not contribute close to their contribution limits. For workers with workplace pensions or deferred profit sharing plans, their annual contribution room in voluntary accounts is reduced dollar-for-dollar by the amounts saved in these occupational plans.

4 Data and Sample Selections

I use several related administrative tax databases from Statistics Canada throughout this study. These data provide rich, longitudinal information on Canadian tax filers' savings behavior in both workplace and voluntary retirement accounts. The data sets and sample restrictions are described in this section.

First, the Longitudinal Administrative Databank (LAD) is a panel file comprising a 20-percent representative sample of Canadians' tax records, which originates from central income-tax registers. The file contains detailed information about demographics, earnings, income, taxes, credits and allowances, transfers, and savings characteristics for the individuals represented and their Census families. Although information on taxable savings or other wealth accumulation such as home equity are not available as this information is not reported on tax forms, the LAD remains one of the most comprehensive database for studying retirement savings in Canada.

The following sample restrictions are imposed. I restrict to the years 1991-2010, which coincides with the first year that data on workplace pension coverage became available and which runs up to the last year of data availability when this study began. I condition on

individuals who were between the ages of 25 and 49 in 1991 (equivalently, 44 and 68 in 2010). The upper age limit (68 in 2010) takes into account the fact that individuals are required to start collecting from their workplace pensions by the time they turn 69 years of age. I focus the analysis on individuals satisfying these criteria who are observed filing their taxes on time in at least 18 out of the 20 years. While relaxing this assumption does not change the main results, this restriction allows me to exploit the longitudinal component of the data in later stages of the analysis. Finally, I omit observations where individuals are observed collecting public or private pension income in order to focus the analysis on pre-retirement savings behavior. Taken together, the sample comprises approximately 34 million observations on 1.8 million tax filers.

Unfortunately, the LAD does not have education variables since this information is not collected by tax authorities. For this, I turn to newly-created data sets that link individuals' 1991 or 2006 Census of Canada responses to panels of their tax records (hereafter, the 'Census-tax linkages' file). The linkage file is limited in that it does not contain the breadth of information or sample size needed to carry out all robustness checks and tests of model extensions that I perform in the crowd-out analysis on the LAD. However, replicating the baseline crowd-out estimates is possible and, in addition, the Census has respondents reported highest level of educational attainment. I impose the same sample restrictions in the linkage file as before, except that I condition on individuals aged 25-68 at any time over the period 1991-2010 and do not specify a minimum number of repeat occurrences, so as to increase sample size for the RKD analysis. I also restrict to individuals who were born in Canada in order to assign assign compulsory schooling laws, which vary by province and cohort, to individuals based on their province of birth as reported in the Census. In all, the Census-tax linkages comprise observations on nearly 800,000 individuals.

5 The Crowd-out Effect of a Nudge

In this section, I estimate the extent to which a savings nudge in workplace pensions raises private retirement savings or induces workers to reduce contributions in voluntary accounts. I begin by describing the empirical strategy and identifying assumptions. Then the primary results, robustness checks, and various tests of model extensions are presented. While the LAD data set is primarily used, I also show that the main results hold in the Census-tax linkage file.

5.1 Empirical Strategy

The identification exploits an automatic change in employer pension contributions arising from the integration feature of these plans, as described in Section 3. The local average change in workplace pension contributions at the kink point (i.e. the average industrial wage), and the resulting displacement effect on voluntary retirement savings, are estimated using a two-stage RKD approach:

$$P_{it} = \alpha_1 + \beta_1(E_{it} - T_t) + \gamma_1(E_{it} - T_t)D_{it} + X'_{it}\delta + \upsilon_{it} \tag{10}$$

$$R_{it} = \alpha_2 + \beta_2(E_{it} - T_t) + \gamma_2(E_{it} - T_t)D_{it} + X'_{it}\theta + \xi_{it} \tag{11}$$

conditional on $(E_{it} - T_t) \in [-B, B]$, where B is the estimation bandwidth. The variables E_{it}, P_{it}, and R_{it}, are labor earnings, workplace pension contributions, and voluntary retirement savings for individual i at time t, respectively. Denote X_{it} as a vector of observable covariates and $D_{it} = \mathbb{1}(E_{it} \geq T_t)$ as an indicator of whether earnings exceeds the kink point T_t in the reference year. In the primary analysis, I rule out corner solutions by conditioning on interior savers, $R_{it}, P_{it} > 0$ and $R_{it} + P_{it} < L_t$, to ensure the results are not confounded by individuals not saving enough to respond or by mechanical effects of the tax code.

Equation (10) estimates a 'first-stage' effect of workplace pensions being integrated with the public pension, whereas equation (11) estimates the 'second-stage' effect on voluntary savings. The crowd-out effect is $\gamma = \frac{\gamma_2}{\gamma_1}$. I estimate both equations jointly as seemingly unrelated regressions (SUR) and obtain standard errors for crowd-out using the Delta method. The standard errors are clustered at the individual level to account for unit-specific correlations of the residuals (Lee and Lemieux, 2010).

The estimator makes two key identifying assumptions. The first assumption is that the change in voluntary retirement savings comes only from the effect of integration on workplace pensions and not from a direct response on the part of workers to the kink in the public pension's contribution schedule.[7] I conduct a placebo test of how voluntary savings behaves at the kink for savers not covered by workplace pensions in order to test this assumption. Second, the estimator provides a direct test of substitution between workplace and non-workplace retirement savings accounts holding constant total compensation given that the

[7]There are two reasons why this assumption is reasonable in practice. First, contributions made into the public pension do not directly tie into retirement benefits. Rather, benefits are determined by a complex function of lifetime earnings and other personal characteristics (see Baker and Benjamin (1999) for more information). As a result, changes in contributions to the public pension—holding everything else constant—do not lead to changes in personal wealth accumulation. Second, to the extent that workers do respond directly to the public pension's contribution rate schedule, the expectation would be for workers to save more in voluntary accounts as they move beyond the kink point, an effect that works in the opposite direction as the actual response detected.

running variable is employment income. However, the method requires that workers' earnings are randomly assigned around the kink point. To test this assumption, Figure 2 plots the distributions of employment incomes relative to the kink point for various groups. Chart (a) shows that a sorting response is detected among the full sample, but charts (b) to (d) indicate that the response is driven by workers who do not belong to unions. This may arise because the collective bargaining process for unionized workers makes it more difficult to precisely control earnings at the individual level. This leads me to condition this analysis on workers who belong to unions (two-thirds of workplace pension plan members, as shown in Table 1). In addition, Figure 3 shows that covariates are generally smooth through the kink, which adds further support for the RKD methodology.

Note that only approximately 50 percent of individuals with workplace pensions also contribute to a voluntary account in a given year. These workers have higher total income ($54,750 versus $43,850), are less likely to be unionized (64.1 percent versus 71.9 percent), to collect Employment Insurance (8.1 percent versus 14.8 percent), and to make employee pension contributions (72.1 percent versus 79.1 percent) compared to workers who are not saving in a voluntary account. Therefore, the analysis necessarily centers on individuals who save comparatively well in retirement accounts. However, the heterogeneity that I document in later sections of this paper sheds some light on what might be expected if this analysis could be extended to a broader class of tax filers.

5.2 Salience of the Treatment Effect

I consider two issues concerning the salience of this retirement savings nudge: (1) how to interpret changes in workplace pension contributions within defined-benefit pension plans; (2) and how information about workplace savings is disseminated to workers.

First, the variable for workplace pension contributions available in Canadian tax records is called a 'pension adjustment' (PA). The PA is a measure of occupational savings for workers belonging to employer-sponsored pension plans, deferred profit-sharing plans (DPSPs), and some unregistered accounts.[8] For every one-dollar increase in workplace savings, the amount that can be saved in voluntary accounts is 'adjusted' downward on a dollar-for-dollar basis. Hence, this system requires an accurate measure of how much people are saving in their workplace pension accounts in order to operate properly. For defined-contribution plans, the PA is simply the sum of employer and employee contributions in a given year. For

[8]The inclusion of DPSPs in the PA results in a modest over-estimation of pension coverage. Morissette and Ostrovsky (2006) discuss concerns about using the PA to proxy for workplace pension coverage but ultimately recommend using this variable. In the context of this study, *changes* in the PA are expected to be driven entirely by the effect of interest.

defined-benefit plans, the PA translates the accrued pension compensation from the past year of service into a dollar equivalent. In other words, it converts the annual change in eventual retirement wealth into a net present value that can be roughly interpreted as a pension contribution. This feature is important especially given that a large share of workplace pensions in Canada are defined-benefit plans.

Second, information about workers' savings within their occupational pensions is disseminated to them annually in two ways. Workers receive a statement of renumeration from their employers each year during tax season, which clearly states their PA. In addition, tax filers receive a deduction limit statement annually from the central tax authority, which indicates their PA from the past year and the amount they are permitted to contribute in voluntary retirement accounts without exceeding their contribution limits. Therefore, people need not have a deep understanding of how their pension plans operate (Luchak and Gunderson (2000) show that this is often the case in Canada) in order to know roughly how much they saved in their workplace plans in a given year.

5.3 Results

The primary graphical analysis and estimation results are shown in Figure 4. First, employer contributions rise on average by 2.5¢ per $1 of employment income beyond the kink point, a significant rise of 36.2 percent from the base savings rate of 6.9¢ per $1. The magnitude of this effect is consistent with expectations. More precisely, since the average tax on employers from the public pension on earnings below the kink point was 3.2¢ per $1 over the time period 1991-2010, and given that approximately 85 percent of workplace pensions are integrated, then $1 taxed on employers displaces workplace pension contributions by $\$0.025 \div (\$0.032 \times 0.85) = 91.9¢$, approximately a unit-elastic response.

Whether to include employee contributions to workplace pensions in the first or second stage of this analysis is not initially clear. If these contributions are also integrated, as Frenken (1996) states may be the case, this variable belongs in the first stage. Yet employees sometimes have control over their workplace contributions beyond the base rates set out in their plan contracts. However, while some plans permit additional voluntary contributions, there is generally no equivalent allowance for reductions. This feature is a problem given that the second-stage prediction is for savings to fall beyond the kink point, which leads me to add employee contributions in the first stage. In practice, the results are robust to either case since there is no discernible change in employee contributions at the kink point.

The estimated decrease in voluntary savings rates for workers who belong to employer-sponsored plans resulting from the nudge is 1.5¢ per $1, or 33.3 percent of the base savings

rate of 4.5¢ per $1. To check that the response is driven entirely by the effect of interest, and is not confounded in some way by a direct response to the public pension, the procedure is reapplied on workers who are observed saving in voluntary accounts but who do not have employer-sponsored pension coverage. As Figure 4 shows, this 'placebo test' does not detect a behavioral savings response at the kink point, consistent with expectations. To explore this issue further, I augment the baseline statistical model (equations (10) and (11)) to a difference-in-differences (DD) framework using savers from the placebo test as a control group. Specifically, the second-stage estimating equation becomes:

$$R_{it} = \tilde{\alpha}_2 + \chi I_{it} + \tilde{\beta}_2(E_{it} - T_t)I_{it} + \tilde{\gamma}_2(E_{it} - T_t)D_{it}I_{it} + \\ \lambda(E_{it} - T_t) + \psi(E_{it} - T_t)D_{it} + I_{it}X'_{it}\tilde{\theta} + X'_{it}\omega + \tilde{\xi}_{it} \quad (12)$$

where $I_{it} = \mathbb{1}(P_{it} > 0)$ is an indicator of workplace pension coverage, and the sample is extended to include all unionized workers who meet the sample selection criteria. Hence, $\tilde{\gamma}_2$ measures the change in the marginal savings rate in voluntary accounts net of any detectable response from the control group as measured by ψ. The results from equation (12) are shown in Table 2, which demonstrates that the crowd-out results are robust to the DD estimator. Moreover, as the table shows, the results continue to hold when individual-specific fixed effects are included, forcing the estimator to exploit dynamic variation in savings within individuals who move around the (non-stationary) kink point over time.[9] Finally, Table 2 indicates that the results are very robust to different choices in the bandwidth (B) and polynomial order in the running variable.

Overall, a $1 automatic increase in workplace pension contributions crowds out other retirement savings by approximately 50¢, depending on the exact model specifications. This finding raises the question of why dollar-for-dollar substitution was not detected as predicted by the model, $\frac{\Delta R^\star}{\Delta P} = -1$. There are several reasons why workplace and voluntary plans may be imperfect substitutes in practice. Due to the lock-in provisions of workplace pensions, assets held in these plans are almost exclusively intended for retirement whereas the tax code does not prevent or explicitly discourage pre-retirement withdrawals from voluntary accounts.[10] Thus, assets held in voluntary accounts may reflect savings for a variety of reasons, including income-smoothing and precautionary purposes (Mawani and Paquette, 2011), home

[9] Lee and Lemieux (2010) note that, since the assignment of treatment in a Regression Discontinuity design is assumed to be random, then including fixed effects is not necessary for identification. However, their inclusion should not affect the results.

[10] For example, Individual Retirement Accounts (IRAs) in the United States typically impose a 10 percent penalty on early withdrawals. No equivalent penalty exists in Canada. Moreover, assets held in voluntary plans in Canada may be withdrawn to finance the purchase of a first home (Steele, 2007) or higher education, the former being used frequently.

equity accumulation, and retirement savings. I expect crowd-out to be larger for individuals who use these plans predominantly for retirement. While administrative tax records do not offer information on how savers *intend* to use voluntary accounts, the longitudinal data provide some insight into how these plans are *actually* being used. With this in mind, Table 3 investigates how the crowd-out estimates change as the analysis centers on different types of savers. First, columns (1) and (2) show that crowd-out is twice as large among savers who contribute at their tax-deduction limits, $R_{it} + P_{it} \geq L_t$ a mechanical result of the tax treatment of these plans. Second, columns (3) to (5) account for pre-retirement withdrawal behavior either directly—by using *net* contributions as the second-stage dependent variable—or indirectly—by conditioning on whether individuals were observed making pre-retirement withdrawals in the past. In all cases, the RKD estimates consistently show that a nudge partially crowds out other retirement savings, indicating that non-retirement savings motives likely do not affect the substitutability of these plans. Third, to consider whether adjustments along other savings margins occur, column (6) uses investment income as the second-stage dependent variable to proxy for taxable savings. Consistent with the theoretical prediction that $\frac{\Delta S^\star}{\Delta P} = 0$, no change at the kink point is detected in this case.

The results ultimately indicate that workers are fairly responsive to a savings nudge on average, more so than would have been expected based on the results of previous research (Madrian and Shea, 2001; Choi et al., 2004; Chetty et al., 2014). However, the fact that crowd-out is less than unity may still mean that some savers are under-responding to the distortion, which leads me to explore heterogeneity in more detail.

6 Education and Active versus Passive Choice

In this section, I investigate the extent to which exogenous variation in educational attainment affects active versus passive choice. This analysis is predicated on the model, which posited that the costs of responding to a nudge depend at least in part on human capital.

As broad motivation, Figure 5 begins by showing that the primary crowd-out result of approximately 50¢ per $1 can be decomposed based on workers' propensity to save for retirement. Specifically, using the longitudinal structure of the data, Panel (a) indicates that workers who have contributed regularly into voluntary retirement savings in the past are more likely to crowd out a savings nudge than their counterparts who save irregularly. In addition, Panel (b) shows that crowd-out approaches zero as the sample conditions progressively on workers with larger unused contribution room. As discussed in Section 3, unused contribution room carries forward such that, holding everything else constant, individuals with more room have saved less in these plans over their lifetimes. Therefore, workers who save more regularly

appear to understand how to substitute adequately between the two savings plans, which could occur due to gradual learning from repeated interactions.

Figure 6 decomposes individuals' responsiveness to a nudge by their reported highest level of educational attainment. Specifically, individuals are sorted according to whether they have at least a high school diploma ('high education') or not ('low education').[11] The results indicate that, first, a savings nudge within workplace pensions exists for both groups. Second, workers with low education under-respond to a nudge as compared to workers with comparatively high education, a finding that suggests human capital at least partly determines active versus passive choice, consistent with the theoretical prediction.

To control for the possibility that individuals' responsiveness to a nudge is in some way correlated with their education, Figure 7 conditions on a *predicted* measure of education using compulsory schooling reforms in the identification. In particular, I estimate the incidence of completing high school, given by:

$$\mathbb{1}(HS_i = 1) = \varpi + \vartheta M_{pc} + \bar{X}'_{ipc}\Lambda + \bar{Z}'_{ipc}\Omega + K'_{pc}\Phi + \varepsilon_{ipc} \tag{13}$$

where HS_i is an indicator of high school completion, and M_{pc} is the mandatory number of years of schooling required for individuals from province-of-birth p and cohort c.[12] The estimation results of this model are shown in Table 4. The regression controls for observable factors X_{ipc} described above, as well as several factors Z_{ipc} from the Census-tax linkage data. Note that the education variables are only available for individuals cross-sectionally from their 1991 or 2006 Census responses, whereas individuals' tax records are observed longitudinally. Hence, I construct $\bar{X}_{ipc}$ and $\bar{Z}_{ipc}$ as inflation-adjusted averages of X_{ipc} and Z_{ipc} for each individual across all observable years; see the notes in Table 4 for more information. Lastly, to control for endogeneity between compulsory schooling reforms and other institutional factors affecting education (e.g., Stephens Jr. and Yang (2014)), I include additional education policy covariates K_{pc} (listed in Table 4), province-of-birth and cohort fixed effects, and

[11] The Census-tax linkage file does not provide a measure of respondents' total number of years of schooling. However, the data provide comparable measures of highest reported level of attainment, with which constructing variables for high school completion, some post-secondary attainment, and university graduate is possible. Since this analysis exploits exogenous variation in education brought about by compulsory schooling reforms, which affect attainment predominantly at lower levels of attainment, the groups were separated by high school completion. The extent to which higher levels of attainment affect active versus passive choice is an interesting avenue for future research but is outside the scope of this study given the policy variation being exploited empirically.

[12] I match individuals to the entry age legislation that was in effect in their province of birth at 6 years old and to the exit age legislation in effect at 14 years old. The Census-tax linkage sample consists of individuals who were 14 years of age between 1938 and 1995 inclusive, hence the compulsory schooling reforms exploited in this paper span this time period. The compulsory schooling and other education variables are described further in Oreopoulos (2006a).

province-of-birth specific cohort trends in the regressions. Based on the estimation results of equation (13), I separate individuals into low-education or high-education groups to keep the sample sizes in the RKD similar to the sample sizes in Figure 6, which yields approximately an 85 percent success rate. Overall, these results continue to support the finding that workers with high education are more active savers than those with low education, which suggests that active versus passive choice to some extent results from knowledge-based traits that are amenable to change through education policy.

To explore this finding in more detail, I assess whether underlying differences in the observable characteristics or savings behavior or low-education and high-education groups likely drive the results. First, Figure 8 shows how the distributions of earnings, workplace pension contributions, and voluntary retirement savings differ across the two groups. Whereas high-education workers earn slightly more and save more in workplace pensions, the distributions of voluntary savings of both groups nearly perfectly overlap. This addresses the concern that low-education workers may under-respond to a nudge because they save less and are simply not able to cut back their contributions as much in voluntary accounts. Second, the basic characteristics of workers in the RKD regressions by education are comparable; for example, low-education and high-education workers are approximately the same age (45.0 versus 44.5, respectively), are equally likely to be married (68.2 versus 66.5 percent), and they have similar average incomes from non-employment ($2,750 versus $3,000). These similarities likely arise because the RKD conditions on workers who are localized in their earnings around the average industrial wage, thus imposing a degree of homogeneity between the two groups irrespective of education.

7 Education and Savings Behavior

The previous section found that education is a strong determinant of whether individuals actively adjust savings in other accounts in response to a nudge, or are passive to the intervention. The findings indicate that a nudge raises the wealth accumulation of individuals with low education, but that those with high education respond by crowding-out contributions in voluntary retirement accounts (suggesting that the adjustment costs for the high-education group are comparatively low). Therefore, a necessary condition for a nudge to increase private welfare is that workers with low education also benefit the most from help saving for retirement, as illustrated in Figure 1. In this section, I provide empirical evidence of whether this condition likely holds in practice by testing whether education directly affects how workers save. In an instrumental variables (IV) framework that uses historical reforms to compulsory schooling laws across Canadian provinces in the identification, I estimate the

causal effect of high school completion on individuals' (observed) lifetime average savings rates in tax-preferred accounts, $\bar{sr}_{ipc}$. The first-stage regression is given by equation (13), above, whereas the second-stage effect of high school completion on savings is:

$$\bar{sr}_{ipc} = \iota + \eta \mathbb{1}(HS_i = 1) + \bar{X}'_{ipc}\Gamma + \bar{Z}'_{ipc}\Pi + K'_{pc}\Psi + \pi_{ipc} \tag{14}$$

Note that the additional variables $\bar{Z}_{ipc}$ are included in the IV regressions so as to control for as many factors affecting savings behavior as possible. This set of variables includes a 7th-order polynomial in total income, sector of employment indicators, and the value of home equity reported on the 1991 or 2006 Censuses; see Table 4 for a complete listing of the covariates used. Using observed life-cycle averages of savings rates, income, and other control variables helps to mitigate bias in the estimation arising from permanent-income effects of education on life-cycle savings.

The second-stage results of this regression are also shown in Table 4. In particular, the analysis suggests that high school completion induces individuals to raise their savings in tax-preferred accounts by 3-6 percentage points over the course of the life-cycle. This finding is robust across a variety of model specifications. Most importantly, the inclusion of additional income or demographic controls does not appear to influence the magnitude of the regression estimates, suggesting that income is the primary channel through which education indirectly affects savings and that other omitted variables are not a significant concern. In Figure 9, the effect of high school completion on savings is well-illustrated. In particular, high-education workers appear to save more in tax-preferred accounts consistently over the life-cycle up to retirement when savings rates fall significantly for both groups. The finding that education raises individuals' retirement savings even at early ages indicates that there are large cumulative returns to schooling on wealth in retirement.

Overall, these findings show that high school completion is associated with an increase in savings within tax-preferred accounts over the life-cycle. This means workers with low education save less for retirement than their counterparts with high education, even controlling for income and other observable factors affecting such behavior. Since workers with low education are also passive to a nudge, whereas those with high education costlessly adjust savings across accounts in response to such an intervention, a nudge is expected to increase consumer welfare in practice.

8 Conclusion

This paper provided new insight into whether a savings nudge in workplace pensions raises private wealth accumulation or simply induces workers to reduce how much they contribute in other retirement accounts. Using rich administrative tax records and linked Census data from Canada, and exploiting unique institutional characteristics of the Canadian retirement income system, I showed that a nudge in employer pension contributions partially crowds other savings out, by approximately \$0.50 per \$1. Hence, this type of policy intervention is a viable approach for increasing net savings on average, but many individuals also actively respond to it.

These findings raise the question of why I detect a larger crowd-out response to a nudge than what would have been expected based on the evidence from related behavioral public finance research. For example, Madrian and Shea (2001), Choi et al. (2002), and Chetty et al. (2014) show that automatic changes in pension contributions are very effective at raising individuals' net savings precisely because of how non-responsive people are to such distortions. The differences in results may partly arise due to the comparatively high degree of salience about how much employers contribute annually into workers' pensions in Canadian for institutional reasons, as discussed above. In addition, whereas previous research has mostly focused on younger, lower-income, or less-educated workers in the identification of crowd-out, this study has centered on workers earning close to the average industrial wage, meaning there is greater diversity in individuals' characteristics and behavior, which likely also explains why some individuals were found to be more responsive to the nudge than others.

Further, in this paper I showed that a nudge raises the savings of individuals who are the most at-risk of under-preparing for retirement without significantly disadvantaging those who save adequately on their own. Less-educated workers save less than highly-educated workers but benefit from a nudge by remaining passive, whereas highly-educated workers are very responsive to the intervention. While many factors affect active versus passive choice, these findings indicate that the behavior to some extent results from knowledge-based characteristics. An implication of this finding is that nudges are expected to increase private welfare in practice. However, these findings also imply that education and nudges are imperfect substitutes, a concept that was modeled in detail in the conceptual framework. More schooling reduces both the extent to which nudges influence savings outcomes *and* the role for soft-paternalism as individuals learn to save on their own. This finding is consistent with past research that finds targeted interventions to increase financial literacy or reduce the complexity of investing have meaningful effects on savings behavior (Bernheim and Garrett,

2009; Bayer et al., 2009; Beshears et al., 2013). Although nudging mechanisms have become popular recently as viable alternative measures to tax incentives for encouraging people to save, this study ultimately indicates that other mechanisms should also be considered when deciding on socially efficient strategies. Future research on the intersection of soft-paternalism and financial literacy training would likely have implications for the optimal design of retirement income systems.

References

Acemoglu, D. and Angrist, J. (2000). How large are human-capital externalities? evidence from compulsory schooling laws. *NBER Macroeconomics Annual*, 15:9–59.

Ambuehl, S., Bernheim, B. D., and Lusardi, A. (2014). Financial education, financial competence, and consumer welfare. Working Paper Series 20618, National Bureau of Economic Research.

Angrist, J. D. and Krueger, A. B. (1991). Does compulsory school attendance affect schooling and earnings? *The Quarterly Journal of Economics*, 106(4):979–1014.

Arnberg, S. and Barslund, M. (2013). The crowding-out effect of mandatory labour market pension schemes on private savings: Evidence from renters in Denmark. Discussion Papers 11/2013-054, Network for Studies on Pensions, Aging and Retirement.

Attanasio, O. P., Banks, J., and Wakefield, M. (2004). Effectiveness of tax incentives to boost (retirement) saving: Theoretical motivation and empirical evidence. *OECD Economic Studies*, 39(2):145–172.

Baker, M. and Benjamin, D. (1999). Early retirement provisions and the labor force behavior of older men: Evidence from Canada. *Journal of Labor Economics*, 17(4):724–756.

Baker, M., Gruber, J., and Milligan, K. (2007). Simulating the response to reform of Canada's income security program. In Gruber, J. and Wise, D. A., editors, *Social Security Programs and Retirement Around the World: Fiscal Implications of Reform*, pages 83–118. University of Chicago Press.

Bayer, P. J., Bernheim, B. D., and Scholz, J. K. (2009). The effects of financial education in the workplace: Evidence from a survey of employers. *Economic Inquiry*, 47(4):605–624.

Benjamin, D. J. (2003). Does 401(k) eligibility increase saving? evidence from propensity score subclassification. *Journal of Public Economics*, 87(56):1259–1290.

Bernheim, B. D. (2002). Taxation and saving. In Auerbach, A. J. and Feldstein, M., editors, *Handbook of Public Economics*, volume 3, pages 1173–1249. Elsevier.

Bernheim, B. D., Fradkin, A., and Popov, I. (2013). The welfare economics of default options in 401(k) plans. Working Paper Series 17587, National Bureau of Economic Research.

Bernheim, B. D., Fradkin, A., and Popov, I. (2015). The welfare economics of default options in 401(k) plans. Unpublished Manuscript.

Bernheim, B. D. and Garrett, D. M. (2009). The effects of financial education in the workplace: Evidence from a survey of households. *Journal of Public Economics*, 124(1):51–104.

Bernheim, B. D. and Rangel, A. (2009). Beyond revealed preference: Choice-theoretic foundations for behavioral welfare econnomics. *The Quarterly Journal of Economics*, 87(7-8):1487–1519.

Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2013). Simplification and saving. *Journal of Economic Behavior & Organization*, 95(C):130–145.

Carroll, G. D., Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2009). Optimal defaults and active decisions. *The Quarterly Journal of Economics*, 124(4):1639–1674.

Chetty, R., Friedman, J. N., Leth-Petersen, S., Nielsen, T., and Olsen, T. (2014). Active vs. passive decisions and crowdout in retirement savings accounts: Evidence from Denmark. *The Quarterly Journal of Economics*, 129(3):1141–1219.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177.

Choi, J. J., Laibson, D., and Madrian, B. C. (2011). $100 bills on the sidewalk: Suboptimal investment in 401(k) plans. *Review of Economics and Statistics*, 93(3):748–763.

Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2002). Defined contribution pensions: Plan rules, participant choices, and the path of least resistance. In Poterba, J. M., editor, *Tax Policy and the Economy*, volume 16, pages 67–114. MIT Press.

Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2003). Optimal defaults. *American Economic Review*, 93(2):180–185.

Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2004). For better or for worse: Default effects and 401(k) savings behavior. In Wise, D. A., editor, *Perspectives on the Economics of Aging*, pages 81–126. University of Chicago Press.

Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2005). Passive decisions and potent defaults. In Wise, D. A., editor, *Analyses in the Economics of Aging*, pages 59–78. University of Chicago Press.

de Serres, A. and Pelgrin, F. (2003). The decline in private saving rates in the 1990s in OECD countries: How much can be explained by non-wealth determinants? Economic Studies 36, Organisation for Economic Co-operation and Development.

Engelhardt, G. V. and Kumar, A. (2011). Pensions and household wealth accumulation. *Journal of Human Resources*, 46(1):203–236.

Engen, E. M. and Gale, W. G. (1994). Do saving incentives work? *Brookings Papers on Economic Activity*, 1994(1):85–180.

Engen, E. M. and Gale, W. G. (2000). The effects of 401(k) plans on household wealth: Differences across earnings groups. Working Paper Series 8032, National Bureau of Economic Research.

Engen, E. M., Gale, W. G., and Scholz, J. K. (1996). The illusory effects of saving incentives on saving. *Journal of Economic Perspectives*, 10(4):113–138.

Euwals, R. (2000). Do mandatory pensions decrease household saving? evidence from the Netherlands. *De Economist*, 148(5):643–670.

Frenken, H. (1996). The impact of changes in the Canada Pension Plan on private pensions. *Canadian Business Economics*, 4(4):65–72.

Gale, W. G. (1998). The effects of pensions on household wealth: A reevaluation of theory and evidence. *Journal of Political Economy*, 106(4):706–723.

Gale, W. G. and Scholz, J. K. (1994). IRAs and household saving. *American Economic Review*, 84(5):1233–1260.

Gelber, A. M. (2011). How do 401(k)s affect saving? evidence from changes in 401(k) eligibility. *American Economic Journal: Economic Policy*, 3(4):103–122.

Harvey, R. (2004). Comparison of household savings ratios: Euro area/United States/Japan. Statistics Brief 8, Organisation for Economic Co-operation and Development.

Hubbard, G. R. and Skinner, J. S. (1996). Assessing the effectiveness of saving incentives. *Journal of Economic Perspectives*, 10(4):73–90.

Ippolito, R. A. (1997). *Pension Plans and Employee Performance: Evidence, analysis, and policy*. University of Chicago Press.

Iwry, M. J. and John, D. C. (2009). Pursuing universal retirement security through automatic IRAs. The Retirement Security Project 2009-3, The Brookings Institute.

Lee, D. S. and Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2):281–355.

Lochner, L. and Moretti, E. (2004). The effect of education on crime: Evidence from prison inmates, arrests, and self-reports. *American Economic Review*, 94(1):155–189.

Luchak, A. A. and Gunderson, M. (2000). What do employees know about their pension plan? *Industrial Relations: A Journal of Economy and Society*, 39(4):646–670.

Lusardi, A. and Mitchell, O. S. (2010). How ordinary consumers make complex economic decisions: Financial literacy and retirement readiness. Working Paper 2010/11, Center for Financial Studies.

Lusardi, A. and Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, 52(1):5–44.

Madrian, B. C. and Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *The Quarterly Journal of Economics*, 116(4):1149–1187.

Mawani, A. and Paquette, S. (2011). Pre-retirement RRSP withdrawals. *Canadian Tax Journal*, 59(2):183–219.

Milligan, K. (2002). Tax-preferred savings accounts and marginal tax rates: Evidence from RRSP participation. *Canadian Journal of Economics*, 35(3):436–456.

Milligan, K., Moretti, E., and Oreopoulos, P. (2004). Does education improve citizenship? evidence from the United States and the United Kingdom. *Journal of Public Economics*, 88(9-10):1667–1695.

Morissette, R. and Ostrovsky, Y. (2006). Pension coverage and retirement savings of Canadian families, 1986 to 2003. Analytical Studies Branch Research Paper Series Catalogue No. 11F0019MIE - No. 286, Statistics Canada.

Mullock, K. and Turcotte, J. (2012). Financial literacy and retirement savings. Working Paper 2012-01, Department of Finance.

Oreopoulos, P. (2006a). The compelling effects of compulsory schooling: Evidence from Canada. *Canadian Journal of Economics*, 39(1):22–52.

Oreopoulos, P. (2006b). Estimating average and local average treatment effects of education when compulsory schooling laws really matter. *American Economic Review*, 96(1):152–175.

Oreopoulos, P. (2007). Do dropouts drop out too soon? wealth, health and happiness from compulsory schooling. *Journal of Public Economics*, 91(11-12):2213–2229.

Ostrovsky, Y. and Schellenberg, G. (2009). Pension coverage, retirement status, and earnings replacement rates among a cohort of Canadian seniors. Analytical Studies Branch Research Paper Series Catalogue No. 11F0019M - No. 321, Statistics Canada.

Poterba, J. M., Venti, S. F., and Wise, D. A. (1994). 401(k) plans and tax-deferred saving. In Wise, D. A., editor, *Studies in the Economics of Aging*, pages 105–142. University of Chicago Press.

Poterba, J. M., Venti, S. F., and Wise, D. A. (1995). Do 401(k) contributions crowd out other personal saving? *Journal of Public Economics*, 58(1):1–32.

Sandmo, A. (1970). The effect of uncertainty on saving decisions. *The Review of Economic Studies*, 37(3):353–360.

Sethi-Iyengar, S., Huberman, G., and Jiang, G. (2004). How much choice is too much? contributions to 401(k) retirement plans. In Mitchell, O. S. and Utkus, S. P., editors, *Pension Design and Structure: New lessons from behavioral finance*, pages 157–184. Oxford University Press.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118.

Sims, C. A. (2006). Rational inattention: Beyond the linear-quadratic case. *Journal of Money, Credit and Banking*, 96(2):158–163.

Statistics Canada (2003). Canada's Retirement Income Programs: A statistical overview (1990-2000) Catalogue No. 75-507-XIE, Statistics Canada.

Steele, M. (2007). The Canadian Home Buyers' Plan: Tax benefits, tax expenditure, and policy assessment. *Canadian Tax Journal*, 55(1):1–30.

Stephens Jr., M. and Yang, D.-Y. (2014). Compulsory education and the benefits of schooling. *American Economic Review*, 104(6):1777–1792.

Thaler, R. H. and Benartzi, S. (2004). Save More Tomorrow^{TM}: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112(S1):S164–S187.

Thaler, R. H. and Sunstein, C. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

Veall, M. R. (2001). Did tax flattening affect RRSP contributions? *Canadian Journal of Economics*, 34(1):120–131.

Venti, S. F. and Wise, D. A. (1996). The wealth of cohorts: Retirement saving and the changing assets of older Americans. Working Paper Series 5609, National Bureau of Economic Research.

Venti, S. F. and Wise, D. A. (2014). The long reach of education: Early retirement. Working Paper Series 20740, National Bureau of Economic Research.

Whitehouse, E. (2013). Policies to encourage private pension savings: Evidence from OECD countries. In Hinz, R., Holzmann, R., Tuesta, D., and Takayama, N., editors, *Matching Contributions for Pensions: A review of international experience*, pages 27–50. International Bank for Reconstruction and Development/The World Bank.

9 Appendix

Table 1: Summary statistics, by workplace pension coverage (LAD)

	Has workplace pension		No workplace pension	
	Average	Median	Average	Median
	(1)	(2)	(3)	(4)
Demographics				
Age	44.5		44.0	
Female	48.6		56.1	
Married	79.1		77.2	
Employment (%)				
Employed	99.7		63.3	
Self-employed	2.9		15.1	
Unionized	67.7		12.5	
Employment Insurance receipt	11.2		16.0	
Conditional income ($)				
Employment income, gross	47,250	44,300	26,850	22,050
Self-employment income, net	1,700	-400	19,550	8,300
Other income	3,600	1,000	12,450	5,150
Total income	49,700	45,950	28,600	20,750
Savings participation (%)				
Has workplace pension	100.0		0.0	
Has employee pension contributions	75.3		0.8	
Contributes to voluntary accounts	53.6		29.0	
Withdraws from voluntary accounts	6.4		5.4	
Has unused room in voluntary accounts	87.6		82.6	
Conditional savings ($)				
Total workplace pension contributions	4,700	4,050	0	0
Employee pension contributions	2,350	1,950	1,400	450
Contributions to voluntary accounts	3,150	2,400	4,350	3,000
Withdrawals from voluntary accounts	2,000	1,300	2,200	1,350
Unused room in voluntary accounts	21,350	15,250	21,150	12,550
Conditional total savings				
Savings ($)	6,250	5,400	3,500	2,250
Savings rate (%)	15.9	11.7	8.4	6.9

Notes: The 'conditional' statistics refer to the average or median value conditional on that value being non-zero. The savings rate is calculated as net tax-deductible savings (in both workplace and voluntary plans) relative to total income. The values shown here (rounded to the nearest $50) correspond to the cohort used in the upcoming empirical analysis from the LAD data and are in nominal dollars for the years 1991-2010 inclusive.

Table 2: RKD robustness checks for interior savers (LAD)

	Estimator		Bandwidth and polynomial order			
			Linear			Cubic
	Difference-in differences	Fixed Effects	B = \$5,000	B = \$7,000	B = \$9,000	B = \$9,000
	(1)	(2)	(3)	(4)	(5)	(6)
	PANEL A: Kink in workplace pensions					
Workplace savings	0.026***	0.017***	0.025***	0.027***	0.028***	0.021***
	(0.002)	(0.002)	(0.003)	(0.002)	(0.001)	(0.005)
	PANEL B: Kink in voluntary retirement savings					
Voluntary savings	-0.016**	-0.008**	-0.012***	-0.014***	-0.014***	-0.012*
	(0.008)	(0.004)	(0.004)	(0.002)	(0.002)	(0.006)
	PANEL C: Crowd-out					
Crowd-out	0.639**	0.466***	0.459***	0.494***	0.494***	0.564*
	(0.309)	(0.175)	(0.156)	(0.093)	(0.066)	(0.323)
No. of individuals	283,259	225,247	202,498	246,277	283,891	283,891
No. of observations	949,730	735,620	614,506	855,176	1,083,696	1,083,696

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. The difference-in-differences estimate of the kink in voluntary savings corresponds to the change for workplace pension plan members net of the change in voluntary savings for those not covered by a workplace pension. The fixed effect estimator includes individual and time fixed effects and captures the crowd-out response in a dynamic setting as workers move around the (non-stationary) kink point. The results are robust to different choices of bandwidth and polynomial order in the running variable, albeit a wider bandwidth is needed for precise estimates of crowd-out at the kink point when the cubic polynomial is used. However, the AIC and BIC tests indicate that a linear specification in the running variable is optimal. See Figure 4 for more information on the regression specifications.

Table 3: RKD tests of model extensions (LAD)

	Type of saver		Net contributions	Withdrew from voluntary plans before retirement?		Investment
	Limit saver	Interior saver	in voluntary plans	No	Yes	income
	(1)	(2)	(3)	(4)	(5)	(6)
	PANEL A: Kink in workplace pensions					
Workplace savings	0.022***	0.026***	0.025***	0.027***	0.024***	0.045***
	(0.004)	(0.002)	(0.002)	(0.003)	(0.003)	(0.006)
	PANEL B: Kink in voluntary retirement savings					
Voluntary savings	-0.031***	-0.015***	-0.014***	-0.017***	-0.013***	0.000
	(0.008)	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)
	PANEL C: Crowd-out					
Crowd-out	1.412***	0.602***	0.553***	0.647***	0.547***	-0.003
	(0.419)	(0.124)	(0.139)	(0.163)	(0.193)	(0.081)
No. of individuals	89,249	225,247	239,065	111,791	113,456	202,334
No. of observations	185,837	735,620	792,839	391,327	344,293	593,651

Notes: *** $p < 0.01$. Columns (1) and (2) show the crowd-out estimates for limit and interior savers, respectively. In column (3), the dependent variable in the second-stage regression is net savings in voluntary accounts, i.e. gross contributions less withdrawals. Columns (4) and (5) use the longitudinal nature of the LAD data to condition on individuals who are not, or who are, observed withdrawing from their voluntary accounts at any time before retirement, although the results are not significantly different in either case. The last column uses investment income as the dependent variable in the second-stage regression to assess whether crowd-out likely exists within taxable savings accounts.

Table 4: The returns to education on savings behavior (Census-tax linkages)

	First stage				Second stage			
	Dep. var.: Completed high school				Dep. var.: Savings rate			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Completed high school					0.060*	0.062*	0.061**	0.032**
					(0.034)	(0.033)	(0.030)	(0.016)
Entry age	-0.009***							
	(0.002)							
Dropout age	0.005***							
	(0.002)							
Mandatory years of schooling		0.007***	0.006***	0.006***				
		(0.001)	(0.002)	(0.002)				
Additional controls			✓	✓			✓	✓
Total income polynomial			✓	✓		✓	✓	✓
Earnings and job characteristics			✓	✓	✓	✓	✓	✓
Education covariates	✓	✓	✓	✓	✓	✓	✓	✓
Birth province and cohort effects	✓	✓	✓	✓	✓	✓	✓	✓
Birth province × cohort trends								
Linear	✓	✓	✓		✓	✓	✓	
Quadratic				✓				✓
No. of observations	492,765	492,765	374,625	374,625	370,598	368,570	368,570	368,570

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. The number of mandatory years of schooling is the difference between the legal age at which students may exit high school and the legal age at which they must begin elementary school; these data are described in more detail in (Oreopoulos, 2006a). The individual covariates include the baseline variables X_{ipc} from the RKD regressions (see notes in Figure 4), and additional variables Z_{ipc} to control for as many indirect channels as possible that may affect savings behavior: indicators of employment status and union coverage, sector of employment indicators (at the 1-digit North American Industry Classification System (NAICS) code level), a 7th-order polynomial in total income, and the value of home equity as reported in the 1991 or 2006 Censuses. The education covariates are: child labor age, an indicator of whether a restrictive labor law was in place, and an indicator of whether exemptions existed, as well as the log of school expenditure, and the number of schools and teachers per capita. Province-of-birth linear cohort trends are used to control for unobserved factors across provinces that may have changed over time that affected education.

Figure 1: Theoretical predictions of the effects of nudging and human capital reform

PANEL A: Effect of a nudge on consumer welfare by type of agent

		Type of savings response to a nudge	
		Active saver, $R^\star = R^\star(h, \tilde{s}, \Delta P)$	Passive saver, $R^\star = R^\star(h, \tilde{s})$
Level of human capital investment	Low, $h^\star = 0$	$(-)$	$(+)$
	High, $h^\star = \bar{h}$	$(-)$	$(-)$ or $(+)$

PANEL B: Effect of a human capital reform on consumer welfare by type of agent

		Type of savings response to a nudge	
		Active saver, $R^\star = R^\star(h, \tilde{s}, \Delta P)$	Passive saver, $R^\star = R^\star(h, \tilde{s})$
Level of human capital investment	Low, $h^\star = 0$	$(-)$ or $(+)$	$(-)$ or $(+)$
	High, $h^\star = \bar{h}$	Nil.	Nil.

Notes: This figure shows the expected effects of a savings nudge and human capital reform on different types of agents based on their level of human capital investments and responsiveness to a nudge. Panel A shows that a nudge strictly decreases consumer welfare for active savers given that the effect of a nudge is perfectly offset by adjustments to savings in other accounts and adjustment costs are incurred. For passive savers, the effect of a nudge depends on the level of human capital acquisition. A nudge strictly increases consumer welfare for agents who benefit the most from this assistance, i.e. those with low human capital. However, the effect of a nudge is ambiguous for those with high human capital since the intervention could overcompensate for the investment uncertainty for this group. Panel B shows that a human capital reform has no effect on agents who find it optimal to invest into human capital on their own. For those who would otherwise not obtain human capital, the effect of this reform is ambiguous. Agents who find human capital acquisition exceedingly costly are made worse off, whereas those who are only marginally dissuaded from obtaining human capital due to the uncertainty about the benefit of doing so can be made better off from such a reform.

Figure 2: Distributions of employment income around the kink point, by type of worker (LAD)

(a) Full sample

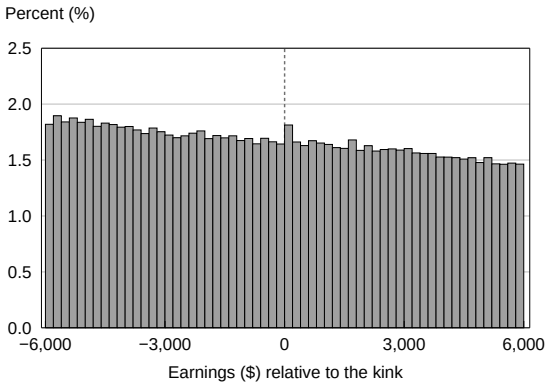

McCrary estimate = 0.049*** (0.010)

(b) Has workplace pension

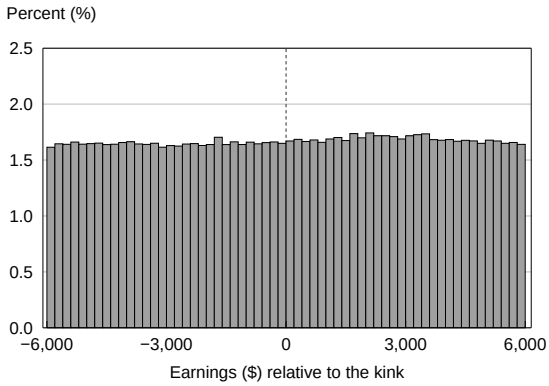

McCrary estimate = 0.013 (0.013)

(c) Unionized, no workplace pension

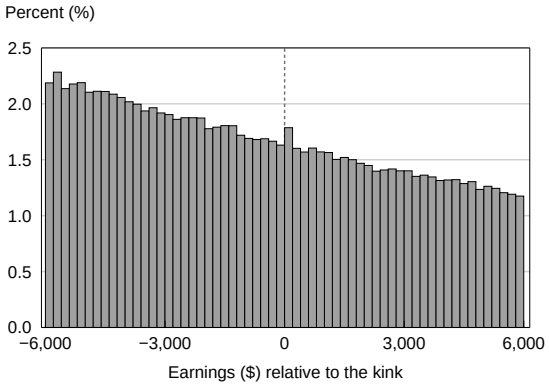

McCrary estimate = 0.028 (0.034)

(d) Non-unionized, no workplace pension

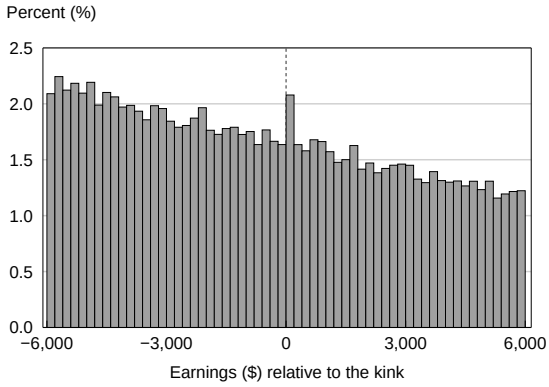

McCrary estimate = 0.118*** (0.017)

Notes: *** $p < 0.01$. Each graph shows the distribution of employment income relative to the average industrial wage (zero corresponds to income equal to the average industrial wage in the reference year) in $200 bins. The McCrary discontinuity test estimates the extent to which a bunching response at the kink point in the public pension contribution schedule is statistically significant using the optimal kernel density bandwidth. Standard errors are in parentheses.

Figure 3: Test of covariates around the kink point for interior savers (LAD)

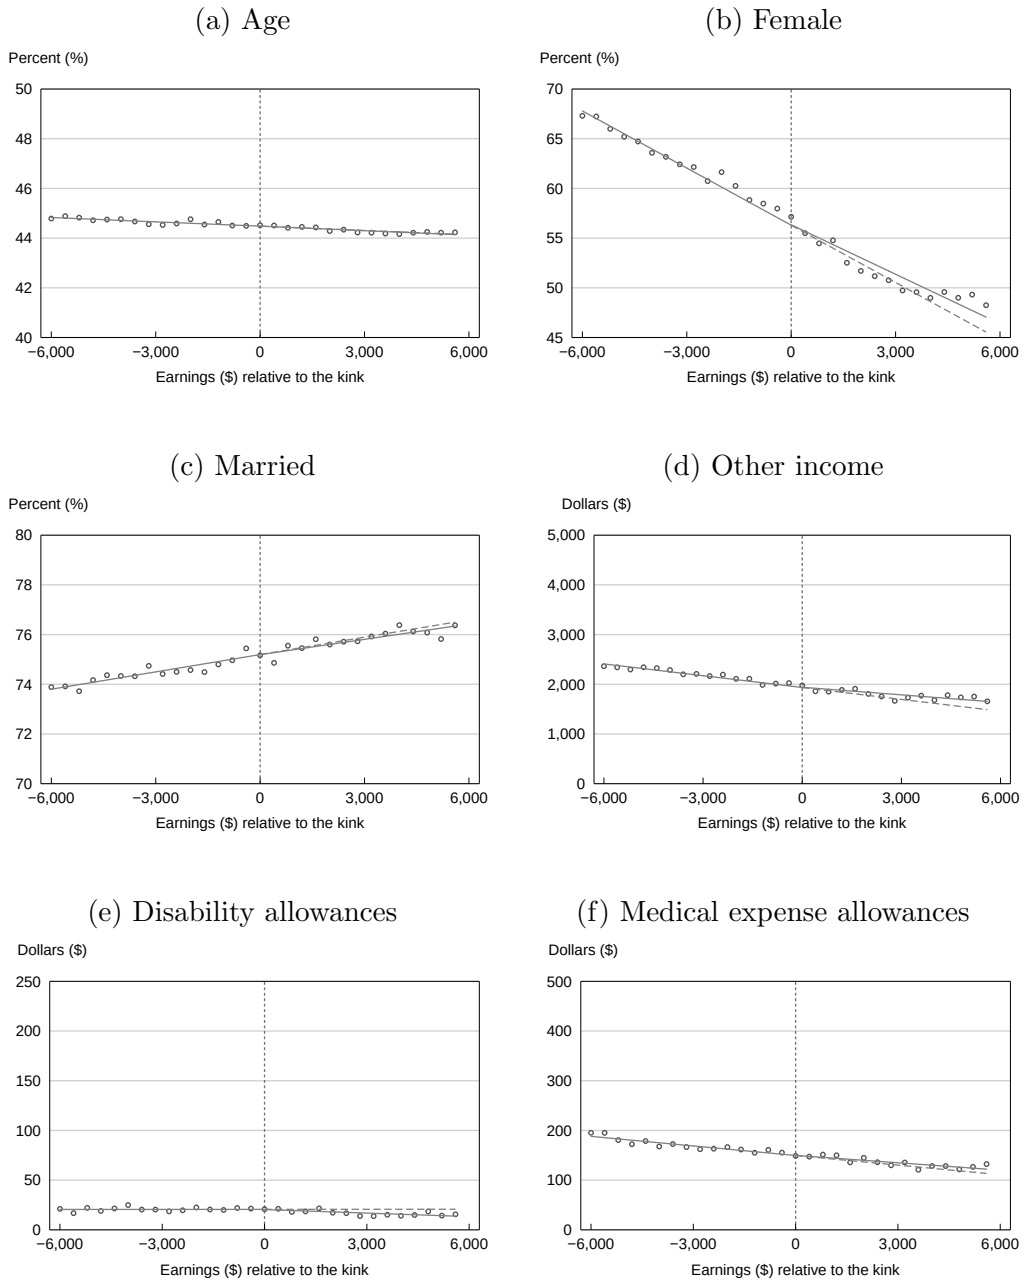

Notes: Each graph plots the average value of the relevant covariate in $400 bin widths for individuals with employment income around the average industrial wage (zero corresponds to income equal to the average industrial wage in the reference year).

Figure 4: RKD primary crowd-out estimates for interior savers (LAD)

(a) Employer contributions

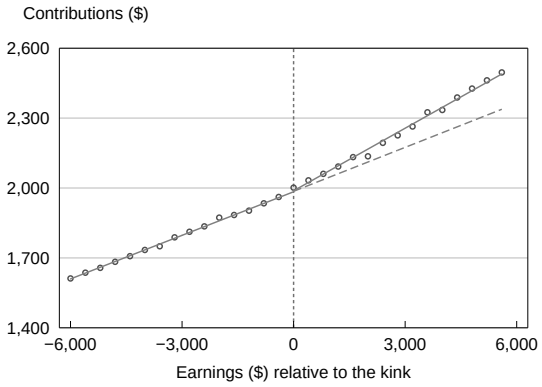

Slope = 0.069*** (0.001)
Kink = 0.025*** (0.002)

(b) Employee contributions

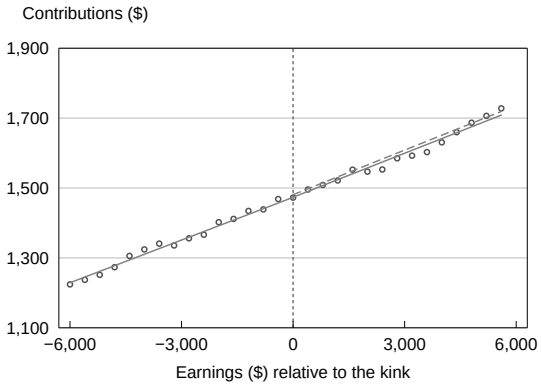

Slope = 0.049*** (0.001)
Kink = 0.001 (0.002)

(c) Voluntary savings, has workplace pension

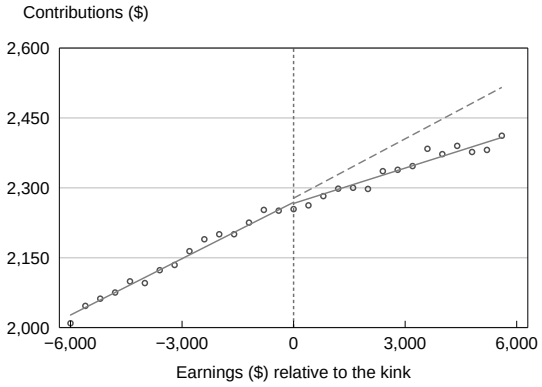

Slope = 0.045*** (0.002)
Kink = -0.015*** (0.003)

(d) Voluntary savings, no workplace pension

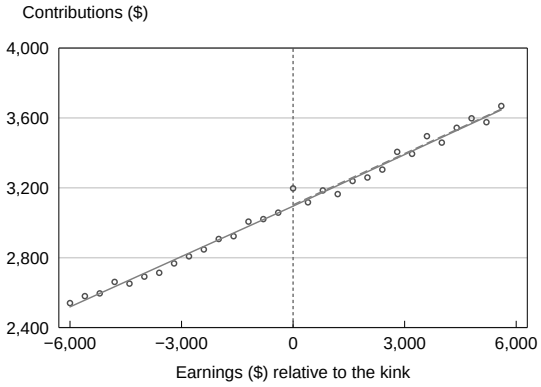

Slope = 0.101*** (0.003)
Kink = 0.001 (0.006)

Crowd-out = 0.602*** (0.124)

Notes: *** $p < 0.01$. Each graph shows contributions to workplace pensions or voluntary savings as a function of employment income relative to the average industrial wage. Each data point represents the average contribution in $400 bin widths. Results are for the years 1991-2010 inclusive. The corresponding linear regression results are also reported, where standard errors (in parentheses) are clustered at the individual level and the bandwidth is set at $6,000. The following control variables are included: indicators of female, married, province of residence, Employment Insurance receipt, has self-employment income, age, age-squared, disability and medical expense tax allowances, other income. The crowd-out estimate is obtained by dividing the (absolute value of the) kink in voluntary savings by the kink in total (employer and employee) workplace pension contributions. I estimate the two equations together in a seemingly unrelated regression (SUR) framework and obtain standard errors (in parentheses) for crowd-out using the Delta method.

Figure 5: RKD crowd-out estimates for interior savers, by contribution history and unused contribution room (LAD)

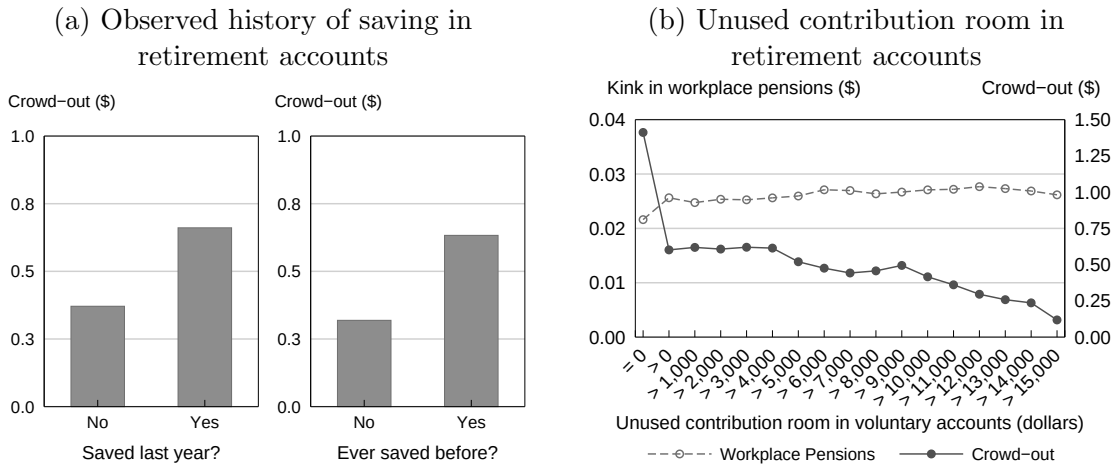

Notes: Chart (a) shows how crowd-out varies according to whether individuals were observed contributing in a voluntary account last year (left) or ever in the past (right). The results indicate that workers who use these plans more frequently also exhibit a higher degree of substitution between the voluntary and workplace accounts. In chart (b), the first-stage effect of integration on workplace pension contributions (left vertical axis) and resulting crowd-out estimates (right vertical axis) are shown, conditional on savers with progressively larger unused contribution room in voluntary retirement savings accounts. Holding everything else constant, individuals with larger unused contribution room have contributed less over their lifetimes into these plans. See the discussion in text for more information about this institutional characteristic. See Figure 4 for more information on the regression specifications.

Figure 6: RKD crowd-out estimates for interior savers, by educational attainment (Census-tax linkages)

PANEL A: Low education

(a) Workplace pension contributions

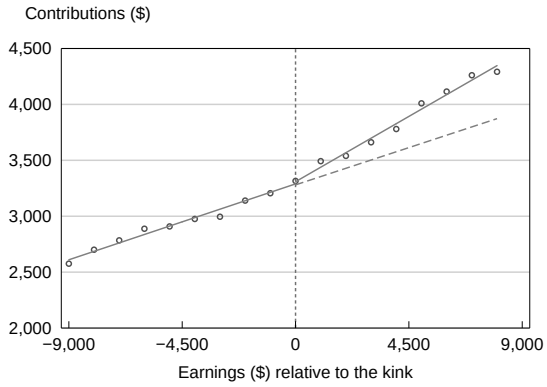

Slope = 0.101*** (0.006)
Kink = 0.047** (0.010)

(b) Voluntary savings

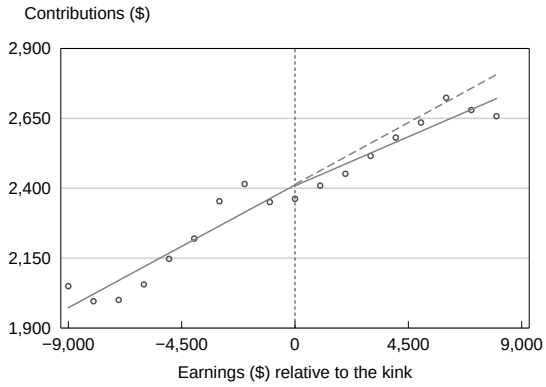

Slope = 0.052*** (0.008)
Kink = -0.009 (0.013)

Crowd-out = 0.184 (0.273)

PANEL B: High education

(a) Workplace pension contributions

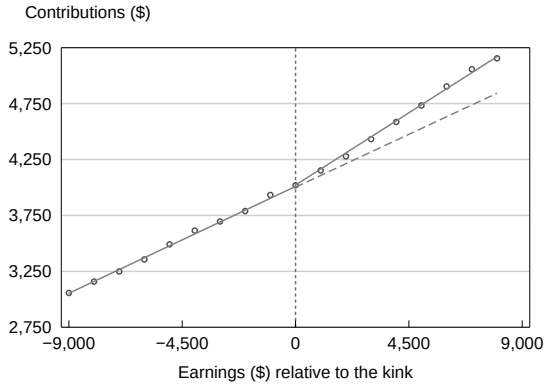

Slope = 0.120*** (0.002)
Kink = 0.032** (0.003)

(b) Voluntary savings

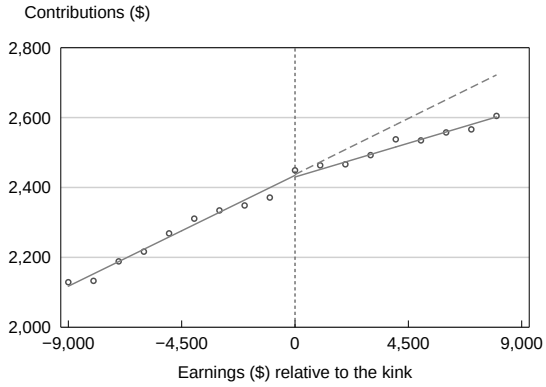

Slope = 0.042*** (0.003)
Kink = -0.018*** (0.004)

Crowd-out = 0.556*** (0.135)

Notes: *** $p < 0.01$; ** $P < 0.05$. Panel A conditions on individuals with low education, defined as less than high school (number of individuals = 6,092; number of observations = 23,819). Panel B conditions on individuals with high education, defined as high school completion or more (number of individuals = 52,943; number of observations = 193,895). The same set of covariates as those used in the LAD data are included here, except for: the variable for medical expense tax allowances is not observed, home equity value from the Census data is included, and sector of employment indicators (at the 1-digit NAICS code) are included; see Figure 4 for more information.

Figure 7: RKD crowd-out estimates for interior savers, by predicted educational attainment (Census-tax linkages)

PANEL A: Low predicted education

(a) Workplace pension contributions

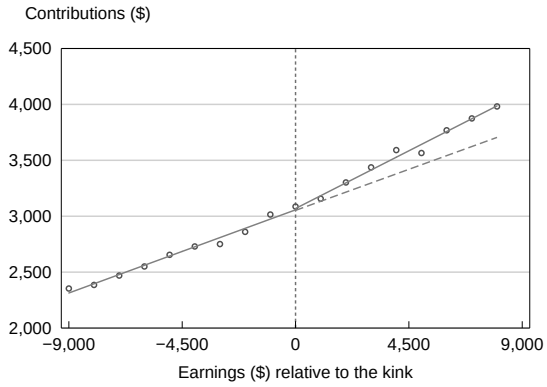

Slope = 0.099*** (0.006)
Kink = 0.023** (0.011)

(b) Voluntary savings

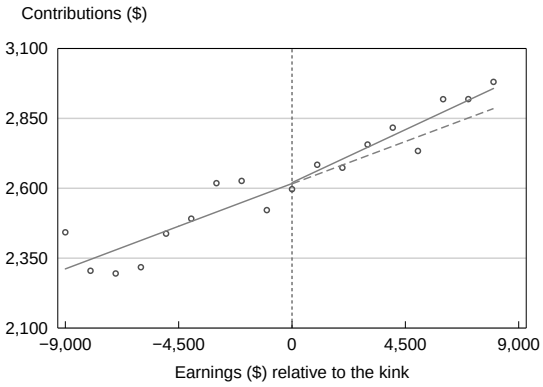

Slope = 0.041*** (0.007)
Kink = 0.001 (0.013)

Crowd-out = -0.033 (0.566)

PANEL B: High predicted education

(a) Workplace pension contributions

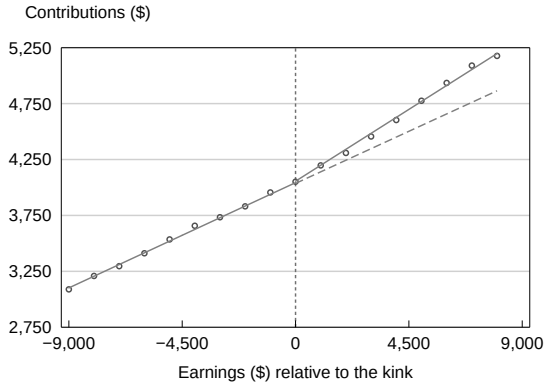

Slope = 0.119*** (0.002)
Kink = 0.034** (0.003)

(b) Voluntary savings

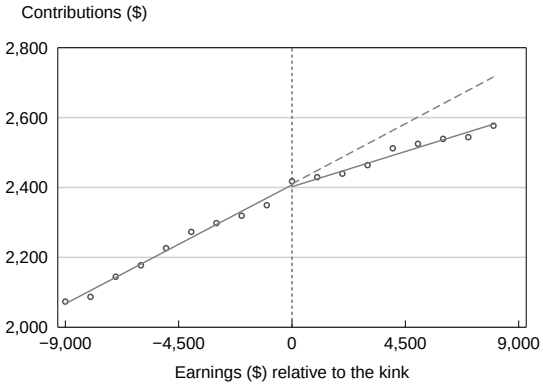

Slope = 0.044*** (0.003)
Kink = -0.019*** (0.004)

Crowd-out = 0.569*** (0.131)

Notes: *** $p < 0.01$; ** $P < 0.05$. Panel A conditions on individuals with low predicted education, defined as less than high school (number of individuals = 6,765; number of observations = 24,113). Panel B conditions on individuals with high education, defined as high school completion or more (number of individuals = 52,297; number of observations = 193,601). See Figure 6 for more information.

Figure 8: Distributions of employment income, workplace pension contributions, and voluntary retirement savings for workers in the RKD sample, by educational attainment (Census-tax linkages)

(a) Employment income

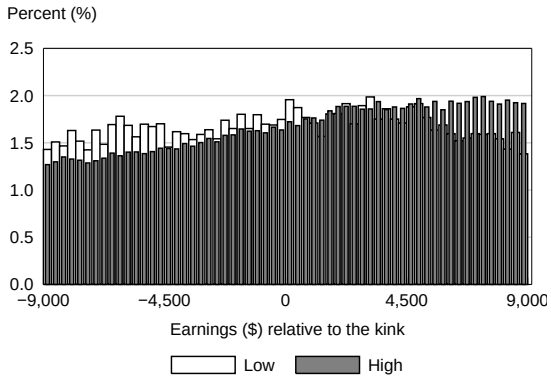

(b) Workplace pension contributions

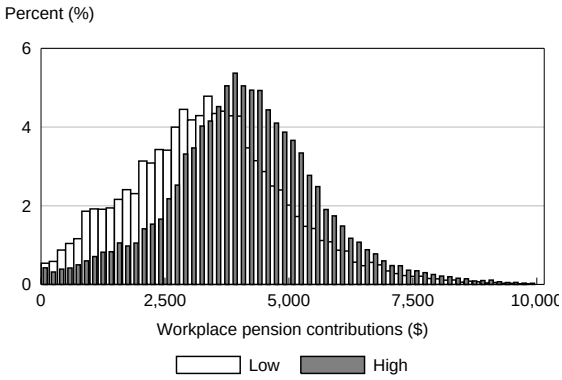

(c) Voluntary savings

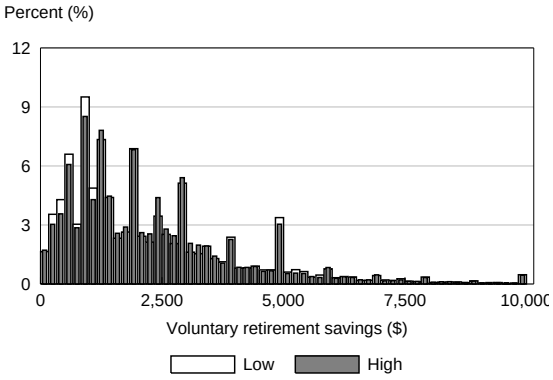

Notes: These graphs plot the distributions of employment income, workplace pension contributions, and voluntary retirement savings by education, conditional on workers who appear in the RKD analysis. Employment income is relative to the average industrial wage in the reference year, as in the RKD graphs, so as to be centered around the kink point. Education levels are based on actual reported highest level of schooling attained.

Figure 9: The effect of education on savings rates over the life-cycle (Census-tax linkages)

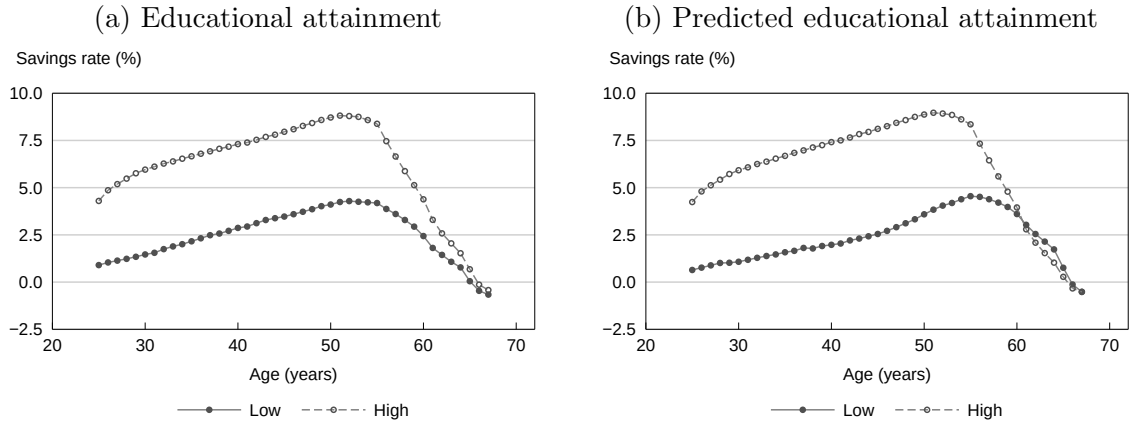

Notes: This figure plots the average savings rate by age for individuals aged 25-67. The savings rate is the total net contributions in tax-preferred accounts (workplace and voluntary) divided by total income. Chart (a) shows the savings rates for individuals with low education versus high education using actual reported values of educational attainment. Chart (b) separates individuals based on their predicted levels of educational attainment from the first-stage instrumental variables regression.

10 Technical Appendix

I derive the conditions described in Section 2.3 concerning the welfare implications of a savings nudge and human capital reform. Recall that, from a social planner's perspective, agents are heterogeneous along two relevant dimensions, as shown in Figure 1.

10.1 Nudge

Consider a nudge $\Delta\tilde{P} \neq 0$ that maximizes the weighted average of agents' realized (or *ex post*) utilities conditional on expectations over signals s, given by:

$$\Delta\tilde{P} \in \arg\max_{\Delta P} \left\{ \sum_{u\leq\bar{u}} \mathbb{E}_s\Big[V(\bar{h}, R^\star(\bar{h},\tilde{s}) + \Delta P; u)p_u(u)\Big] + \sum_{u>\bar{u}} V(0, R^\star(0,\tilde{s}) + \Delta P; u)p_u(u) \right\} \quad (15)$$

Recall that $\hat{R}$ is the optimal retirement savings if r_0 were known. Since $|\hat{R} - R^\star(0,\tilde{s})| > |\hat{R} - \mathbb{E}_s[R^\star(\bar{h},\tilde{s})]|$, i.e. more information about r_0 reduces characterization failure on average, $\Delta\tilde{P}$ increases expected utility for passive agents with $\{h,\tilde{s}\} = \{0,\emptyset\}$. If this were not true, the nudge would make agents worse off on average irrespective of their human capital, but then there would be no role for the intervention and $\Delta\tilde{P} = 0$, a contradiction. In contrast, depending on the relative masses of agents with and without human capital, those with $\{h,\tilde{s}\} = \{\bar{h}, s\}$ could still be made worse off on average from this reform.

Denote $\bar{q}(h) \in \mathbb{E}_{\tilde{s}}[G(h,\tilde{s})] = k(h,\bar{q}(h);\zeta)$ as the maximum cost for which agents with human capital h are active savers, in expectation. The share of agents with $h \in \{0,\bar{h}\}$ who are active savers is $\sum_{q\leq\bar{q}(h)} p_q(q)$ if the number of agents is large. The expected aggregate effect of introducing a nudge $\Delta\tilde{P}$ is:

$$\begin{aligned} W(\Delta\tilde{P}) = &\sum_{u\leq\bar{u}}\sum_{q\leq\bar{q}(\bar{h})} -k(\bar{h},q;\zeta)p_q(q)p_u(u) + \sum_{u>\bar{u}}\sum_{q\leq\bar{q}(0)} -k(0,q;\zeta)p_q(q)p_u(u) + \\ &\sum_{u\leq\bar{u}}\sum_{q>\bar{q}(\bar{h})} \mathbb{E}_s\big[V(\bar{h}, R^\star(\bar{h},\tilde{s},\Delta\tilde{P}); u)\big]p_q(q)p_u(u) + \\ &\sum_{u>\bar{u}}\sum_{q>\bar{q}(0)} V(0, R^\star(0,\tilde{s},\Delta\tilde{P}); u)p_q(q)p_u(u) \end{aligned} \quad (16)$$

Equation (16) comprises both the expected gains and losses for passive agents and the adjustment costs incurred by active agents. Hence, the optimal nudge is $\Delta P^\star = \Delta\tilde{P}$ if

$W(\Delta\tilde{P}) > 0$ and $\Delta P^\star = 0$ otherwise. For a nudge to increase aggregate consumer welfare, there must be sufficiently many passive agents who the planner expects to benefit from this intervention to offset the adjustment costs imposed on active savers and (possibly) the distortions on passive savers who already tend to save adequately.

10.2 Human Capital Reform

Since r_0 is known to a social planner, from this perspective the optimal acquisition rule should be: $h^\star = \bar{h}$ if and only if $u \leq \bar{\bar{u}}$, where $\bar{\bar{u}} = \bar{h}^{-1}\{\mathbb{E}_s[U(R^\star(\bar{h}, \tilde{s}))] - \mathbb{E}_r[U(R^\star(0, \tilde{s}))]\}$. Notice that $\mathbb{E}_s[U(R^\star(\bar{h}, \tilde{s}))] > \mathbb{E}_s\mathbb{E}_{r|s}[U(R^\star(\bar{h}, \tilde{s}))]$ since more information is always better. Hence, the uncertainty about the benefit of obtaining human capital leads to sub-optimal acquisitions: $\bar{u} < \bar{\bar{u}}$. The expected welfare effect of this reform is:

$$W(\Delta h^c) = \sum_{u > \bar{u}} \left(E_s[V(\bar{h}, R^\star(\bar{h}, \tilde{s}); u)] - V(0, R^\star(0, \tilde{s}); u) \right) p_u(u) \tag{17}$$

The expected effect is negative for the mass of agents $\sum_{u \geq \bar{\bar{u}}} p_u(u)$ whose costs of human capital acquisition are exceedingly large, but positive for the mass $\sum_{u \in (\bar{u}, \bar{\bar{u}})} p_u(u)$. Note that agents who obtain human capital on their own, $u < \bar{\bar{u}}$, do not enter into this welfare analysis since the reform is non-binding for this group, which contrasts with a nudge that affects all agents in some way.