

SeaTE: Subjective *ex ante* Treatment Effect of Health on Retirement

Pamela Giustinelli

Bocconi University, IGER, and LEAP

Matthew D. Shapiro

University of Michigan and NBER

First draft: December 30, 2017

This draft: August 31, 2018

The research reported herein was pursuant to a grant from the U.S. Social Security Administration (SSA), funded as part of the Retirement Research Consortium (RRC). The findings and conclusions expressed are solely those of the author(s) and do not represent the views of SSA, any agency of the federal government, Bocconi University, the University of Michigan, the Michigan Retirement Research Center at the University of Michigan, or the Vanguard Group Inc. We thank Feiya Shao and Ann Rodgers for excellent research assistantship. The research uses data from the Vanguard Research Initiative (VRI) what was developed by a research team under a program project grant from the National Institute on Aging P01-AG026571. The Vanguard Group Inc. supported the data collection of the VRI. Vanguard's Client Insight Group and IPSOS SA were responsible for implementing the VRI survey and provided substantial input into its design. John Ameriks, Andrew Caplin, and Matthew D. Shapiro are co-principal investigator of the VRI. The design of the VRI benefited from the collaboration and assistance of Joseph Briggs, Wandi Bruine de Bruin, Alycia Chin, Mi Luo, Minjoon Lee, Brooke Helppie McFall, Ann Rodgers, and Christopher Tonetti as part of the program project, from Annette Bonner (Vanguard), and Wendy O'Connell (IPSOS SA). This project uses Survey 4 of the VRI that was designed by Ameriks, Briggs, Caplin, Lee, Shapiro, and Tonetti. For documentation of the VRI, including a dynamic link to the survey instrument, see <http://ebp-projects.isr.umich.edu/VRI/>. A preliminary version of this research was presented at the Retirement Research Consortium conference in August 2017 under the title "Using Subjective Conditional Expectations to Estimate the Effect of Health on Retirement." We are grateful to Gabor Kezdi for very helpful suggestions, to Wilbert van der Klaauw for discussing the paper at the AEA 2018 meetings and for providing detailed feedback along with many useful comments as well as to all other participants in the AEA 2018 session on "Subjective Expectations, Belief Formation, and Economic Behavior."

Contact: Pamela Giustinelli, Bocconi University and IGER. E-mail: pamela.giustinelli@unibocconi.it.

“I was thinking about my health, the way things are going in the economy, I don’t know if it’s going to really pick up [...] There might be a chance of me working and there might be a chance that there won’t be much work when I’m that age. If I’m in good health [...] Well, I have no retirement, so if I am working, it’s going to have to be later than 65 if my health is good where I can work.” (57 years old working man explaining his answer to a question asking the percent chance that he will be working past age 62.)

How does a treatment, such as poor health, affect a behavior, such as retirement? In behavioral data, by which we mean realized outcomes or decisions of economic agents, the econometrician observes only the behavior conditional on the treatment, but not the counterfactual behavior absent the treatment. When there is unobserved heterogeneity across individuals, this inherent feature of behavioral data makes difficult inferences about causal effects. Rather than relying on behavioral data, the paper uses subjective expectations of a behavior or outcomes conditional on values of the treatment to elicit *ex ante* treatment effects. This strategy allows each individual to be both treatment and control, thereby obviating the unobserved counterfactual problem and allowing for unrestricted heterogeneity across individuals. This approach yields an individual-level estimate of the treatment effect that we call the Subjective *ex ante* Treatment Effect, or *SeaTE* for short. Thus, this paper provides a strategy for quantifying person-specific treatment effects and for characterizing the distribution of causal effects across the population.

We implement our approach by asking older workers participating in the Vanguard Research Initiative (VRI) to report the conditional likelihood, on a 0-100 percent chance scale that they will be working to specified horizons under alternative health scenarios. They also report their unconditional likelihoods of working to those horizons and of experiencing those health states. Using these data we generate individual and aggregate level estimates of the Subjective *ex ante* Treatment Effect (*SeaTE*) of health on retirement age, given by the difference between respondents’ likelihoods of working in low versus high health. We interpret the *SeaTE* in a dynamic programming framework, and estimate a simple structural econometric model of health and retirement combining the conditional probability measures with this theoretical framework.

The Subjective *ex ante* Treatment Effect (*SeaTE*) of health on retirement is zero for almost 30% of working respondents aged 57 and higher at both 2 and 4 year horizons. The remaining 70% reports have a strictly negative *SeaTE* (median of -40% and standard deviation of 24% percent for the 2-year horizon; median of -30% and standard deviation of 25% percent for the 4-year horizon). The structural model implies that moving from high to low health has a large, negative effect on the mean value of continued work, but the within-person correlation of this value across states is high.

Literature on Health and Labor Supply. The determinants of retirement have been widely studied in economics and elsewhere (e.g., see recent reviews by Coile (2015), Fisher et al. (2016), and French and Jones (2017)). The role of health has been subject to much debate due to the difficulties of unpacking the health-work nexus.

First, the sign of the relationship is theoretically ambiguous. Health might operate through a variety of mechanisms such as preferences, productivity, financial incentives, horizon, expectations (e.g., see Rust and Phelan (1997), Blau and Gilleskie (2008), van der Klaauw and Wolpin (2008), Bound et al. (2010), French (2005), French and Jones (2011), and Garcia-Gomez et al. (2017)), and in several forms such as expected trajectory vs. unexpected shocks, earlier vs. later changes, types of conditions (e.g., see Grossman (1972), Bound et al. (1999), Lumsdaine and Mitchell (1999), McGarry (2004), and Blundell et al. (2016)).

Second, the magnitude of the relationship is hard to quantify empirically as health and work are jointly determined and tend to feed dynamically into each other. This has prompted researchers to investigate the potential effect of retirement on health (e.g., see Rohwedder and Willis (2010), Coe and Zamarro (2011), and Behncke (2012)).

Finally, both retirement and health are subject to several measurement issues which exacerbate the challenges of studying their relationship (e.g., see Bound (1991), Dwyer and Mitchell (1999), Lindeboom and Kerkhofs (2009), and Kapteyn and Meijer (2014) on health measurement, and Gustman et al. (1995, 2010) and Maestas (2010) on concepts and measures of retirement).

Literature on Subjective Expectations. Since the early 1990s economists have increasingly measured individuals' subjective expectations in surveys, using a 0-100 scale of percent chance. This endeavor was stimulated by the importance of subjective expectations in economic models of lifecycle behavior (e.g., Hamermesh (1985) builds the case for measurement of perceived horizons or longevity expectations) and by earlier empirical evidence and theoretical arguments demonstrating the greater informativeness of elicited probabilities for binary events over more commonly used "yes/no" intention measures (see Juster (1966) and Manski (1990)).

Manski (2004, 2017), Attanasio (2009), Hurd (2009), van der Klaauw (2012), Armantier *et al.* (2013), Delavande (2014), Schotter and Trevino (2014), Carroll (2017), and Giustinelli and Manski (2018) trace the development of the subjective expectations literature from various perspectives. Manski (2004) provides an historical account and discusses the main issues arising when measuring expectations. Hurd (2009) focuses on measurement and analyses of

expectations of older Americans in the Health and Retirement Study (HRS). Attanasio (2009) and Delavande (2014) deal with measurement and analysis of subjective expectations in developing countries. van der Klaauw (2012) discusses use of expectations data in structural models of intertemporal choice. Schotter and Trevino (2014) extend the discussion to decision-making under uncertainty by experimental subjects in the lab and Giustinelli and Manski (2018) to family processes of human capital investment and schooling decisions. Armantier et al. (2013) deal with inflation expectations and, more generally, Manski (2017) and Carroll (2017) summarize the progress and discuss the promise of measurement of macroeconomic expectations. Papers by Arrondel et al. (2017), Bordalo et al. (2017), and Fuster et al. (2017) included with ours for presentation at the AEA 2018 session on “Subjective Expectations, Belief Formation, and Economic Behavior,” are recent advances in this literature.

Using Subjective Expectations To Study Labor Supply and Its Relation with Health.

Advancing and combining earlier ideas on use of probabilistic expectations data to predict choice behavior in incomplete scenarios posed by the researcher (Manski (1999); see also Dominitz (1997), Wolpin (1999), and Blass et al. (2010)), and to estimate dynamic programming models (van der Klaauw (2012); see also Pantano and Zheng (2013)), we use choice expectations conditional on a specified future state, in addition to unconditional choice and state expectations, to quantify the causal effect of health on retirement and to provide an interpretation of the effect within two mainstream frameworks of econometric causality, the potential outcome framework (POF) and dynamic programming (DP).

Methodologically our paper is closest to recent works by Arcidiacono et al. (2017) and Wiswall and Zafar (2016). These papers generate estimates of individual and aggregate level effects of college major choice on post-graduation outcomes of college undergraduates enrolled in two top U.S. universities, by comparing students’ subjective conditional expectations about post-graduation outcomes across alternative scenarios about their graduation major. Hence, these papers consider a modeling framework à la Roy (Roy, 1951), where potential treatments are alternative human capital investments (college majors) and potential outcomes are the monetary and/or non-monetary consequences of making alternative investments (occupation, earnings, marriage market outcomes, etc.).¹ We, instead, focus on a class of models of intertemporal decision-making amenable to dynamic programming treatment, where potential treatments are

¹ These studies build on earlier work by Dominitz and Manski (1996) who elicit and analyze subjective earnings distributions of a sample of Wisconsin high school students and college undergraduates under alternative scenarios for future schooling.

alternative states of nature (individuals' health) and potential outcomes are feasible actions that agents can take after learning the realized states (working vs. not).

Within the labor supply literature, only a few studies to date have employed survey measures of subjective working and/or health expectations to study retirement behavior and its relationship with individuals' health. In all cases, *unconditional* working probabilities have been used as an outcome variable *in place of or in combination with* actual labor supply data to estimate *ceteris paribus* effects (McGarry, 2004) or structural parameters (van der Klaauw and Wolpin, 2008). In turn, health and longevity expectations have been used to generate moment conditions for structural estimation (van der Klaauw and Wolpin, 2008), or to construct expectations-based health shocks (Shao, 2016).

McGarry (2004) investigates the effect of health on labor supply expectations of working respondents in the Health and Retirement Study (HRS). Using a simple regression analysis, the paper explores the roles of a variety of health measures (e.g., contemporaneous, lagged, and changes in self-reported health, diagnosed health conditions, and subjective longevity expectations), on respondents' subjective probability of working past age 62 and its changes, finding large health effects. The key innovation of the analysis is to replace actual labor supply with unconditional working expectations as a dependent variable in order to focus on working respondents and avoid "justification bias" in self-reported health among retirees. Construction of counterfactuals, identification, and estimation are standard ones for mean linear regression.²

van der Klaauw and Wolpin (2008) develop and estimate a rich dynamic programming model of household retirement and savings, using multiple waves of the HRS. Innovating on earlier structural models of retirement, the authors combine respondents' unconditional working and longevity expectations with observed realizations of respondents' labor supply, health, and the other state variables in order to increase estimates precision rather than to aid parameters identification.³ Construction of counterfactuals, identification conditions, and estimation are otherwise standard for dynamic programming models.⁴

Extending the reduced-form literature where health effects on retirement are identified off health shocks, Shao (2016) employs subjective health expectations elicited in the 2006 HRS and

² Based on conditioning on the relevant set of regressors to the end of performing *ceteris paribus* comparisons.

³ In a related paper, van der Klaauw (2012) estimates a structural dynamic model of teacher career decisions under uncertainty where information on expected future occupation of NLS-72 respondents and their realized occupations are combined to improve estimates precision.

⁴ For example, see Rust (1992), Magnac and Thesmar (2002), and Abbring (2010) for main identification results and discussions. Pantano and Zheng (2013) propose using measures of unconditional subjective expectations for future choice or state variables to identify time-invariant unobserved heterogeneity in the first step of Hotz and Miller (1993)'s estimation method. Kalouptsidei et al. (2017) study identification of counterfactuals.

the corresponding health realizations observed in the 2010 HRS to construct an expectations-based measure of health shock. She then uses this measure in standard regression analysis to investigate the effect of a health shock on observed labor supply and related outcomes.⁵

In general, data on unconditional choice expectations, (or on choice expectations elicited under a single scenario), can be used to perform unconditional (or conditional) prediction of population behavior (see Manski (1999)). Combined with data on choice realizations, it can be further used to improve estimation efficiency (as in van der Klaauw and Wolpin (2008)). Moreover, its greater variation over realized choices can sometimes help address specific issues (as in McGarry (2004)). However, identification of treatment effects or structural parameters and extrapolation to alternative scenarios requires elicitation of conditional choice expectations under multiple alternative scenarios, or a combination of unconditional and conditional choice expectations (see Wolpin (1999)). This is the line we pursue in this paper.

A small set of studies has investigated preferences for work and retirement arrangements using stated preference methods. For instance, Kapteyn et al. (2007) and van Soest and Vonkova (2014) study preferences for full and partial retirement in the Netherlands using hypothetical choices. These papers recognize that uncertainty about future health, employment, and other factors might play a role when respondents make their choice evaluations, nonetheless they do not incorporate these dimensions in the choice scenarios, nor they allow respondents to express uncertainty in the form of choice probabilities. Our approach explicitly addresses these aspects.

The paper is organized as follows. Section I deals with interpretation of the *SeaTE* as causal effects within the potential outcomes framework. Section II deals with interpretation of the *SeaTE* within the dynamic programming framework and discusses econometric implementation. Section III presents the VRI study and describes the analytic sample and the main survey measures. Section IV presents individual and aggregate level estimates of the *SeaTE* directly obtained from the raw conditional probability data, investigates their predictors, and presents parameter estimates of a simple structural model of taste for work, health, and labor supply. Section V concludes by discussing work in progress and next steps.

⁵ In fact, observed health shocks have been exploited in both structural models and more reduced-form analyses; e.g., see Bound et al. (1999), Cai (2010), Disney et al. (2006), Maurer et al. (2011), McGeary (2009), van der Klaauw and Wolpin (2008), Garcia-Gomez (2011), Blundell et al. (2016), and Jones et al. (2016).

I. Subjective *ex ante* Treatment Effects (*SeaTE*)

In this section, we present our econometric approach based on measures of labor supply and health expectations and we compare it to more traditional approaches using data on realized health and labor supply. Our approach employs survey data on subjective working probabilities contingent on alternative hypothetical states of the person's future health to construct measures of person-specific *Subjective ex ante Treatment Effect (SaeTE)* of health on labor supply and to characterize the implied *distribution of ex ante* treatment effects across the population.

A major advantage of our expectations-based approach over the traditional approach based on realizations data is that it enables us to circumvent the logical impossibility of observing individuals' counterfactual labor supply behavior corresponding to the health states that individuals do not experience. This is achieved by asking respondents to predict their labor supply behavior under *all* health states that they might experience. We additionally elicit respondents' subjective expectations of experiencing those health states. Hence, for each respondent our data encompass probabilistic *ex ante* predictions of the respondent's health and labor supply outcomes which *ex post* will be either realized (actual) or counterfactual.

An important related advantage of this data structure is that it naturally allows for within-person comparisons of (expected) labor supply behavior across alternative health states. These comparisons yield person-specific effect of health on labor supply and allow for unrestricted effect heterogeneity across individuals. This effect may be interpreted as a causal partial effect (as opposed to a non-causal or total one), under specific conditions (spelled out below). These conditions involve the relationship (or lack thereof) between the state variables whose values are set by the researcher in the elicitation scenario and, thus, are kept fixed by the respondent when reporting their labor supply expectations and those variables whose values are not specified by the researcher and, thus, might not be kept fixed by the respondent.⁶

The *ex ante* approach and measures that we advance in this paper can be easily employed to study the causal relationship between health and retirement – or any other state and behavior – within two mainstream microeconomic frameworks, potential outcomes (POF)⁷ and dynamic programming (DP).⁸ Both frameworks involve the three main tasks characterizing any causal

⁶ See Heckman (2003) for a discussion about the distinction between conditioning and fixing. In short, conditioning refers to having a variable on the RHS of a regression. Fixing involves an exogenous manipulation of that variable that can have a causal interpretation.

⁷ This usually denotes the Rubin (1974, 1976)-Holland (1986)'s model of causality developed in statistics. However, we refer here to its interpretation and developments in econometrics, as discussed for instance in Heckman (2001, 2005, 2008, 2010) and Manski (1995, 2007).

⁸ Eckstein and Wolpin (1989), Rust (1992, 1994), Keane and Wolpin (2009), Abbring (2010), Aguirribeira and Mira (2010), and Arcidiacono and Ellickson (2011) review the approach from various perspectives.

analysis (see Heckman (2001, 2005, 2008, 2010)): (1) laying out a set of hypotheticals or counterfactuals alongside a mechanism of manipulation or selection for choosing them; (2) studying identification of the causal parameters under ideal population data;⁹ (3) estimating the parameters from actual data.

Whereas empirical causal analyses have been overwhelmingly implemented using data on realizations of the relevant variables,¹⁰ the outlined three-task structure contains no restrictions limiting inference to parameters identifiable from *ex post* measures of realized hypotheticals. In this paper, we demonstrate how to use *ex ante* measures of predicted hypotheticals to identify and estimate causal parameters of interest, which can be given precise meaning both within the context of a DP model of retirement and within the context of a POF or regression-based analysis of the effect of health on retirement.

While we view the *ex post* and *ex ante* approaches as complements rather than substitutes, the latter approach naturally delivers person-specific effects which, in turn, enable one to recover and analyze the implied population distributions of effects in a straightforward manner.¹¹

A. Treatment Effects in Realizations

We consider a simple setting where labor supply is modelled as a binary variable. In period t , after observing the realized value of the state vector, s_{it} , the decision maker decides whether to work or not: $d_{it} \in \{1,0\}$, where 1 denotes working and 0 not working. The state vector, s_{it} , includes the decision-maker's health and other variables discussed below. Health is also modelled as a binary variable, $h_{it} \in \{0,1\}$, where 1 denotes low health and 0 high health.

Figure 1 illustrates the setting by means of a decision tree with three time periods. Within the context of a formal DP model of labor supply, the decision tree is the extensive-form representation of the decision maker's problem as a game against nature, where nodes are information sets and arcs are alternating decisions by nature and the agent (see Rust (1992)). Here we use the tree as a unifying tool to help us illustrate the connections between the POF and DP settings as well as between the *ex post* and *ex ante* approaches.

N-labeled nodes (or N-nodes for short) denote nature's decision points and A-labeled nodes (or A-nodes for short) denote the agent's decision points. For simplicity we drop subscript i and

⁹ These are called structural parameters in DP and treatment effects in the POF.

¹⁰ These are called realized states and choices in DP and realized treatments and outcomes in the POF.

¹¹ The VRI also has similar elicitation tasks for those who have stopped working. Its job separation module asks conditional probability question analogous to the ones discussed in this paper, e.g., what is the percent chance that the respondent would have continued to work conditional on high or low health.

display the simple case where health is the only state variable ($s_{it} \equiv h_{it}$). At each N-node nature assigns a health level to the agent from the set of feasible health levels (high or low), represented as arcs exiting each N-node. In the figure, high-health arcs are labeled as H and low-health arcs as L . At each A-node the agent optimally chooses between working and not working after learning whether their health is high or low; thus, retirement is not necessarily an absorbing state. In the figure, working arcs are labeled as W and non-working arcs as $\sim W$. At each terminal node the agent obtains a payoff, corresponding to a separate path through the tree.

Arches exiting from N-nodes can be easily reinterpreted as *feasible treatments* and arches exiting from A-nodes as *potential outcomes*. At any given t , individual i is characterized by a response function, $d_{it}(h_{it})$, which maps mutually exclusive and exhaustive health treatments into labor supply outcomes. Hence, $d_{it}(h_{it})$ is the potential outcome of person i at time t associated with treatment h_{it} .

Within-person differences in potential outcomes across pairs of hypothetical treatments yield *individual-level treatment effects* of the form,

$$\Delta_{it} = d_{it}(1) - d_{it}(0), \quad (1.1)$$

with $\Delta_{it} \in \{-1, 0, 1\}$. Here being treated corresponds to experiencing a negative health shock contemporaneous to the time of the decision. Recovering this effect entails the evaluation and comparison of the labor supply decisions that person i would make in two mutually exclusive and alternative states of the world at time t , described respectively by $h_{it} = 1$ and $h_{it} = 0$.

For example, consider an agent who is initially in high health and working, represented by the thickened solid path in Figure 1. Conditional on being alive at time $(t+1)$ (assumed in the figure), this agent will be “treated” with either high health or low health in that period. The individual-level treatment effect of health on labor supply at time $(t+1)$ for this particular agent is given by the difference in the agent’s labor supply decisions across the two health states at $(t+1)$, corresponding to equation (1.1) with t replaced by $(t+1)$, 1 replaced by L , and 0 replaced by H .

Let now $z_{it} \in \{0, 1\}$ denote the *actual* health state of person i at time t , where 0 means low health and 1 high health as before. Then, $d_{it} \equiv d_{it}(z_{it}) \in \{0, 1\}$ is the person’s *realized outcome*, that is the labor supply alternative i actually selects given that they experience health state z_{it} ; whereas, $d_{it}(1 - z_{it}) \in \{0, 1\}$ is the person’s *counterfactual outcome*, or the labor supply alternative that i would have selected had they experienced health state $(1 - z_{it})$.

Continuing with the example in Figure 1, let us assume that the agent happens to experience low health at time $(t+1)$ and decides not to work in that period. Then, $z_{i,t+1} = L$ is the actual treatment and $d_{i,t+1} = \sim W$ is the realized outcome, represented by the thickened dashed path. The high health-working combination, (H, W) , represented by the thickened dotted path is instead counterfactual. The counterfactual outcome corresponding to the unrealized treatment is logically unobservable.¹² Hence, any attempt to causal analysis using data on realized treatment and outcomes requires construction of the relevant counterfactuals.

Different econometric approaches deal with this issue differently (e.g., see Heckman et al. (1999)). Most regression-based approaches within the POF compare outcomes *across groups* of suitably similar persons rather than within person. A popular parameter choice is the Average Treatment Effect (ATE),

$$\begin{aligned} ATE_t(1-0) &= E[d_{it}(1) - d_{it}(0)] \\ &= E[d_{it}(1)] - E[d_{it}(0)], \end{aligned} \tag{1.2}$$

where $E[d_{it}(h)]$ denotes the mean of the population labor supply distribution at time t if everyone were to experience health state h , covariates are omitted for simplicity, and $E[d_{it}(h)] \equiv P[d_{it}(h) = 1]$ as d is binary. Decomposition of the two expectations into realized and counterfactual components yields,

$$\begin{aligned} ATE_t(1-0) &= \{E[d_{it}(1) | z_{it} = 1]P(z_{it} = 1) + E[d_{it}(1) | z_{it} = 0]P(z_{it} = 0)\} \\ &\quad - \{E[d_{it}(0) | z_{it} = 0]P(z_{it} = 0) + E[d_{it}(0) | z_{it} = 1]P(z_{it} = 1)\}. \end{aligned} \tag{1.3}$$

Random sampling of $\{z_{it}, d_{it}\}$ from the population distribution of realized health and labor supply asymptotically reveals all components of (1.3) but the counterfactual distributions,

$E[d_{it}(1) | z_{it} = 0]$ and $E[d_{it}(0) | z_{it} = 1]$. Thus, without additional assumptions on these counterfactual distributions, the ATE parameter is not (point) identified (e.g., see Manski (1995, 2003)).

Studies based on randomized control trials (RCT) address the identification problem by randomly assigning subjects to treatments. For example, suppose that individuals are randomly assigned to either a *Control* group or to a *Treatment* group. Control individuals receive treatment $z_{it} = 0$; whereas, treated individuals receive treatment $z_{it} = 1$. In the ideal case of perfect randomization, treatment and control individuals are identical with respect to all observable and

¹² Since Holland (1986) this impossibility has been referred to as the “*fundamental problem of causal inference.*”

unobservable characteristics but their treatment status. Hence, the distribution of realized outcomes in each group can be used as a measure of the counterfactual outcome distribution for the other group.¹³ Formally,

$$\begin{aligned} E[d_{it} | i \in Treatment] &= E[d_{it}(1) | i \in Control] = E[d_{it}(1)] \\ E[d_{it} | i \in Control] &= E[d_{it}(0) | i \in Treatment] = E[d_{it}(0)]. \end{aligned} \quad (1.4)$$

These conditions yield point identification of the ATE parameter in (1.3),

$$ATE_i(1-0) = E[d_{it}(1)] - E[d_{it}(0)] = E(d_{it} | i \in Treatment) - E(d_{it} | i \in Control). \quad (1.5)$$

Obviously, randomization of good and bad health across individuals is not a viable strategy to study the causal relationship between health and labor supply of interest to this paper. The bulk of the literature to date has relied on data on health and labor supply realizations, thus having to grapple with non-random selection of individuals into health states, formally,

$$\begin{aligned} E[d_{it}(0) | z_{it} = 1] &\neq E[d_{it}(0) | z_{it} = 0] \\ E[d_{it}(1) | z_{it} = 0] &\neq E[d_{it}(1) | z_{it} = 1]. \end{aligned} \quad (1.6)$$

Equation (1.6) says that the observed labor supply distribution among high-health (respectively low-health) individuals is likely to differ from the unobserved labor supply distribution that those individuals would have featured had they experienced low (respectively high) health.

Whenever longitudinal data on health and labor supply realizations are available, a popular strategy has been to identify the effect of interest off health shocks. This approach, however, requires observation of a sufficient number of shocks (“switching states”) in the data and makes it challenging to extrapolate to individuals who have yet to experience any shock and/or might be unlikely to experience any. More importantly, it requires that the health change be uncorrelated with respect to other factors affecting retirement in order to identify a casual effect.¹⁴

B. Treatment Effects in Expectations

The approach we advance in this paper circumvents the impossibility of simultaneously observing $d_{it}(1)$ and $d_{it}(0)$ *ex post* by measuring them *ex ante*. This is done by directly asking individuals to predict their outcome (decision in this case) at specific horizons under specified

¹³ Unfortunately, perfectly implemented social experiments or randomized control trials are hard to find in practice (e.g., see Heckman and Smith (1995), Heckman (1996)).

¹⁴ Another approach exploits the availability of some instrument. For instance, Rohwedder and Willis (2010) use cross-country variation in retirement policies to estimate the causal effect of retirement timing on cognition. In general, finding sources of random variation in health is challenging.

scenarios about their treatment (health state in this case) at those horizons.¹⁵ We define the *Subjective ex ante Treatment Effect*

$$\begin{aligned}
SeaTE(i, t, \tau) &= E_{i, t-\tau}(\Delta_{it}) \\
&= E_{i, t-\tau}[d_{it}(1)] - E_{i, t-\tau}[d_{it}(0)] \\
&= P_{i, t-\tau}[d_{it}(1) = 1] - P_{i, t-\tau}[d_{it}(0) = 1],
\end{aligned} \tag{1.7}$$

as the individual-level expectation at $t - \tau$ of the individual level treatment effect at t , Δ_{it} . We measure the *SeaTE* by eliciting the probability of the decision in a survey τ periods in advance. Specifically, we elicit the percent chance of decision $d_{it}(h_{it}) = 1$ given h_{it} equal to 0 and 1.

These individual-level effects can be aggregated across individuals to generate subjective *ex ante* versions of popular group-level parameters (see Cornelissen et al. (2016) for a recent review); for example, the Average Subjective *ex ante* Treatment Effect (A*SeaTE*), the Average Subjective *ex ante* Treatment Effects on the Treated (A*SeaTT*), and the Average Subjective *ex ante* Treatment Effects on the Untreated (A*SeaTU*), defined as following:

$$\begin{aligned}
ASeaTE(t, \tau) &= E[SeaTE(i, t, \tau)], \\
ASeaTT(t, \tau) &= E\left[\frac{\pi_{i, t-\tau}}{E[\pi_{i, t-\tau}]} \cdot SeaTE(i, t, \tau)\right], \\
ASeaTU(t, \tau) &= E\left[\frac{1 - \pi_{i, t-\tau}}{E[1 - \pi_{i, t-\tau}]} \cdot SeaTE(i, t, \tau)\right],
\end{aligned} \tag{1.8}$$

where $\pi_{i, t-\tau} = \text{Prob}_{i, t-\tau}(h_{it} = 1)$ is agent's i subjective probability at time $(t - \tau)$ of entering health state 1 (low health) in the period t and the expectations are taken across individuals. Thus, the latter two parameters are weighted averages of the *SeaTE* (i, t, τ) . The A*SeaTT* gives more weight to individuals with higher subjective probability of entering low health (=being treated), whereas the A*SeaTU* gives more weight to individuals with higher probability of remaining in high health (=being untreated). The A*SeaTT* and A*SeaTU* parameters in (1.8) are functions of individuals' expectations about their future health in addition to their working expectations under alternative health scenarios, which we also elicit in the survey. These parameters can be consistently estimated using their sample analogs.

¹⁵ Following Manski (1999), a scenario can be formalized a function assigning a potential choice set and environment to each member of the population. Hence, it is interpretable as a treatment policy or program. In our application, we focus on specification or fixing of specific features of the choice environment (a state variable) and leave the choice set unspecified. We assume that the latter consists of the two alternative options of working vs. not working.

Measurement and interpretation of the individual-level *SeaTE* does not rely on specific assumptions about the nature of respondents' expectations. However, as noted by Arcidiacono et al. (2017), “*If individuals form rational expectations over their future outcomes, and in the absence of unanticipated aggregate shocks, this [the ASeaTE] parameter coincides with the mean (ex post) effect of treatment on outcome.*” Thus, the aggregate population parameters and their estimates are more directly interpretable under the assumption that individuals form rational expectations about their future health. For example, under this assumption, the expressions for *ASeaTT* and *ASeaTU* still hold after replacing the subjective health probability π with the health indicator h .¹⁶

II. Dynamic Programming Interpretation of *SeaTE*

In this section, we relate the *SeaTE* to the individuals' decision problem. Doing so allows us to have an interpretation of the *SeaTE* in terms of the individuals' optimization problem or, to be more explicit, about the role of the health state in the decision and, hence, about the meaning of conditional expectations with respect to the choice environment.

Individual agents are represented by primitives, $u_{it}(s_{it}, d_{it})$ and $\pi_{it}(s_{i,t+1} | s_{it}, d_{it})$. As before, i indexes individuals and t time periods, with $i = 1, \dots, N$ and $t = 0, 1, \dots, T < \infty$. $u_{it}(s_{it}, d_{it})$ is the utility that agent i derives in period t from choosing labor supply $d_{it} \in \{0, 1\}$, given the realizations of the state variables collected in s_{it} (the state vector), including health and other variables. Because health and the other state variables are generated by a Markovian stochastic process governing their evolution over time, their future values are uncertain from the viewpoint of the decision-maker. Specifically, $\pi_{it}(s_{i,t+1} | s_{it}, d_{it})$ is agent i 's subjective probability over the states' realizations in the next period ($t+1$), conditional on the agent's information set in the current period. The latter is summarized by the realized state and decision at t (as opposed to the whole history of states and choices since the first period), as implied by the Markov-process assumption.

With additively time-separable utility, the agent's utility functional at t is given by

¹⁶As discussed in Manski (1999), some weakened forms of rational expectations (e.g., respondents' subjective choice probabilities for a certain action are unbiased estimates of their objective probabilities of choosing that action), and of statistical independence of the realized treatments across the population (needed for applicability of the law of large numbers), would leave the above conclusion intact. However, aggregate shocks making treatments dependent across the population and systematic deviations from rational expectations in the form of biased expectations would generally invalidate it.

$$U_{it} = \sum_{j=0}^T \beta^j u_{i,t+j} (s_{i,t+j}, d_{i,t+j}) \quad (1.9)$$

where β is the discount factor. The agent behaves optimally according to the expected-utility maximizing decision rule, $\delta_{it}^*(s_{it})$, which satisfies the Bellman (1957)'s optimality principle.

That is, at any time t and state s_{it} , δ_{it}^* is optimal also for the continuation process featuring the current state as starting point,

$$\delta_{it}^*(s_{it}) = \arg \max_{d_{it} \in \{0,1\}} \left\{ u_{it}(s_{it}, d_{it}) + \beta \sum_{s_{i,t+1}} V_{i,t+1}^* [s_{i,t+1}, \delta_{i,t+1}^*(s_{i,t+1})] \cdot \pi_{it}(s_{i,t+1} | s_{it}, d_{it}) \right\}, \quad (1.10)$$

where $V_{i,t+1}^*$ is the value function representing the expected present discounted value of lifetime utility from following δ_{it}^* . This expression makes transparent that $\delta_{it}^*(s_{it})$ is a deterministic function of s_{it} , given the primitives.¹⁷

A. Dynamic Programming Fixing Health

Our approach elicits respondents' expectations about their labor supply in τ periods, both unconditionally and upon fixing the level of future health. Because health is just one element of the state vector, s_{it} , interpretation of respondents' answers within the context of the DP framework requires that the state vector be partitioned into components that are specified or fixed by the researcher in the elicitation task and those that are not specified.

In general, interpretation of respondents' answers and of the derived *SeaTE* parameters depends on the relationship (or lack thereof) between the specified and unspecified components of the state vector. We therefore partition the state vector into variables fixed in the elicitation task and variables not fixed in the elicitation task. We further partition the latter into variables that the researcher could reasonably fix in the elicitation task, if they decided to do so, and variables capturing any residual uncertainty of the agent at the time of elicitation about aspects of the choice environment that might affect their future decision.

Formally, $s_{it} = (x_{it}, y_{it}, \varepsilon_{it})$, where x_{it} denotes the *specified* component of s_{it} , y_{it} denotes the *unspecified* component of s_{it} , and ε_{it} denotes the *residual* component of s_{it} . Under this partition, the expression for the agent's utility in equation (1.9) becomes

¹⁷ Following Rust (1992) and the traditions of the DP literature, at this point we specify this dynamic program at a high level of abstraction including leaving constraints implicit. When we give an example below, we will introduce the relevant constraints.

$$U_{it} = \sum_{j=0}^T \beta^j u_{i,t+j} \left[(x_{i,t+j}, y_{i,t+j}, \varepsilon_{i,t+j}), d_{i,t+j} \right], \quad (1.11)$$

and the related expression for the agent's optimal solution in equation (1.10) becomes

$$\begin{aligned} \delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it}) = & \arg \max_{d_{it} \in \{0,1\}} u_{it} \left[(x_{it}, y_{it}, \varepsilon_{it}), d_{it} \right] + \\ & \beta \sum_{(x_{i,t+1}, y_{i,t+1}, \varepsilon_{i,t+1})} V_{i,t+1}^* \left[(x_{i,t+1}, y_{i,t+1}, \varepsilon_{i,t+1}), \delta_{i,t+1}^* (x_{i,t+1}, y_{i,t+1}, \varepsilon_{i,t+1}) \right] \cdot \pi_{it}^{xy\varepsilon} \left[(x_{i,t+1}, y_{i,t+1}, \varepsilon_{i,t+1}) \mid (x_{it}, y_{it}, \varepsilon_{it}), d_{it} \right] \end{aligned} \quad (1.12)$$

where summations are implicitly taken as many times as required by the dimension of the state vector.

In our application, we specify the individual's health state so $x_{it} = h_{it}$, while leaving unspecified in y_{it} additional factors typically assumed to affect retirement decision (e.g., family and financial conditions, income, and so on).¹⁸

Because x_{it} is fixed in the elicitation task, it is no longer stochastic to the respondent at the time of elicitation. We maintain that the respondent treats it as fixed when reporting their subjective labor supply expectations given health. In this context, the variation in health should be regarded as experimental, so even if health is endogenous, because we are fixing health, the estimates can have a causal interpretation.¹⁹

On the other hand, y_{it} and ε_{it} are stochastic from the perspective of time of elicitation. We assume that respondents hold subjective distributions for the unspecified components of the choice environment at time t and allow them to express any uncertainty they might have about future decision, $\delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it})$, due to the uncertainty they might perceive about y_{it} and ε_{it} .

¹⁸ Because it is generally impractical—if not at all infeasible—for a researcher to fully describe relevant scenarios to respondents, specified scenarios are generally *incomplete* (Manski, 1999). An incomplete scenario can be thought of and formalized as a collection of scenarios, each sharing some common feature (the specified components). In our application, scenarios have in common a specified health level and a horizon length. Because the agent's age at the time of choice is simply equal to their age at the time of elicitation plus the horizon, we handle this by including the individual's age in the state vector. Furthermore, it is important to notice that the elicitation tasks implicitly condition on being alive. Likewise, the Markov transitions are implicitly conditional on being alive. Using the standard convention that utility when dead is normalized to be zero, conditioning on being alive is natural and has no effect on the optimization problem. A full model would, of course, need to account for mortality risk.

¹⁹ In particular, we assume that agents place themselves in the hypothetical situation defined by the scenario, without trying to infer why one or the other scenario might be realized. Dominitz and Manski (1996) discuss this issue within the context of elicitation of students' expectations for future earnings under hypothetical scenarios about their schooling. (See also Dominitz (1997) for additional discussion and examples.) In the setting considered by Dominitz and Manski (1996), and more recently by Arcidiacono et al. (2017) and Wiswall and Zafar (2016), the most subtle aspect of elicitation is to avoid that respondents condition on endogenously choosing the schooling levels posed by the scenario.

Absent this uncertainty about factors driving choices in the future, respondents would give either a zero or one response to the elicitation task because labor supply in the future would be a deterministic function of health.²⁰

We will now discuss two main cases concerning the stochastic structure of x_{it} and y_{it} , whose distinction bears consequence for interpretation of respondents' answers and the nature of *SeaTE*. The first and more general case allows y_{it} to depend on x_{it} either deterministically or stochastically. In our example, y_{it} would include wage. In virtually all microeconomic models of retirement, bad health affects agents' labor supply by increasing their disutility of working *and* by lowering their productivity (e.g., see French and Jones (2011, 2017)). The second case deals with the simpler situation where x_{it} and y_{it} are independent.

Without loss of generality because y_{it} embodies all omitted factors that could be specified, we maintain that ε_{it} is orthogonal to x_{it} and y_{it} , which implies

$$\pi_{it}^{xy\varepsilon} \left[\left(x_{i,t+1}, y_{i,t+1}, \varepsilon_{i,t+1} \right) \mid \left(x_{it}, y_{it}, \varepsilon_{it} \right), d_{it} \right] = \pi_{it}^{xy} \left[\left(x_{i,t+1}, y_{i,t+1} \right) \mid \left(x_{it}, y_{it} \right), d_{it} \right] \pi_{it}^{\varepsilon} \left(\varepsilon_{i,t+1} \mid \varepsilon_{it}, d_{it} \right).^{21}$$

Note that ε here is unknown to both the econometrician *and* the individual at the time of elicitation. This contrasts to the more typical setting for modeling outcome data where the respondent knows a component that is unobserved to the econometrician.

Given this partitioning, equation (1.12) becomes

$$\begin{aligned} \delta_{it}^* \left(x_{it}, y_{it}, \varepsilon_{it} \right) = & \arg \max_{d_{it} \in \{1,0\}} u_{it} \left[\left(x_{it}, y_{it}, \varepsilon_{it} \right), d_{it} \right] + \\ & \beta \sum_{x_{i,t+1}} \left(\int_{\varepsilon_{i,t+1}} \int_{y_{i,t+1}} V_{i,t+1}^* \left[\left(x_{i,t+1}, y_{i,t+1}, \varepsilon_{i,t+1} \right), \delta_{i,t+1}^* \left(x_{i,t+1}, y_{i,t+1}, \varepsilon_{i,t+1} \right) \right] \cdot \pi_{it}^y \left(y_{i,t+1} \mid x_{i,t+1}, y_{it}, d_{it} \right) dy_{i,t+1} \cdot \pi_{it}^{\varepsilon} \left(\varepsilon_{i,t+1} \mid \varepsilon_{it}, d_{it} \right) d\varepsilon_{i,t+1} \right) \\ & \pi_{it}^x \left(x_{i,t+1} \mid x_{it}, y_{it}, d_{it} \right) \end{aligned} \quad (1.13)$$

where we replace summation with integral to allow for the possibility that y and ε are continuous.

The expectation of the optimal decision $\delta_{it}^* \left(x_{it}, y_{it}, \varepsilon_{it} \right)$ as of the time of elicitation $t-1$ is

²⁰ That is, using the terminology of Blass et al. (2010), these subjective distributions express *resolvable uncertainty*, as they refer to utility components whose values may be unknown to the respondent (and the researcher) at the time of elicitation, but will be known by the respondent at the time of choice. A respondent may also face *unresolvable uncertainty*, concerning future utility components that will still be unknown to them at the time of choice. The latter components do not generate uncertainty in the elicitation task, as they are subsumed by the continuation value in the dynamic programming problem.

²¹ Here we treat ε as a scalar. Each process of the choice environment could feature its own residual component, e.g., one in the agent's utility, one in the wage process, and so on.

$$\begin{aligned}
& P_{i,t-1} \left[\delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it}) = 1 \right] \\
&= \sum_{x_{it}} \left[\int_{\varepsilon_{it}} \int_{y_{it}} \delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it}) \cdot \pi_{i,t-1}^y (y_{it} | x_{it}, x_{i,t-1}, y_{i,t-1}, d_{i,t-1}) dy_{it} \cdot \pi_{i,t-1}^\varepsilon (\varepsilon_{it} | \varepsilon_{i,t-1}, d_{i,t-1}) d\varepsilon_{it} \right] \cdot \pi_{i,t-1}^x (x_{it} | x_{i,t-1}, y_{i,t-1}, d_{i,t-1}) \\
&= \sum_{x_{it}} P_{i,t-1} \left[\delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it}) = 1 | x_{it} \right] \cdot \pi_{i,t-1}^x (x_{it} | x_{i,t-1}, y_{i,t-1}, d_{i,t-1}),
\end{aligned} \tag{1.14}$$

where the expression for $\delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it})$ is given in (1.13) and $P_{i,t-1} \left[\delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it}) = 1 | x_{it} \right]$ in the last line of the expression is the individual's expected optimal decision in period t , obtained by *integrating* $\delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it})$ with respect to the distributions of ε_{it} and y_{it} and by *evaluating* the resulting function at a particular realization of x_{it} specified by the elicitation task.

That is, from the viewpoint of a respondent at time $(t-1)$, their optimal choice at time t , $\delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it})$, is a random variable, as it is a function of random variables x_{it} , y_{it} , and ε_{it} . As the elicitation scenario fixes the value of x_{it} , the uncertainty associated to the stochastic process for x_{it} gets partialled out into the transition probabilities, $\pi_{it}^x (x_{i,t+1} | x_{it}, y_{it}, d_{it})$. On the other hand, uncertainty may remain about y_{it} and ε_{it} . For this reason, we allow respondents to report their expected choice probabilistically, expressed as their *subjective probability of working contingent on the specified value of the state*.

Specifically, we elicit all components of (1.14), as follows:

- (i) On the right-hand side of (1.14), the probability of working *given* fixed values of the specified state component, $P_{i,t-1} \left[\delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it}) = 1 | x_{it} \right]$, and the probability of the specified state, $\pi_{i,t-1}^x (x_{it} | x_{i,t-1}, y_{i,t-1}, d_{i,t-1})$, with $x_{it} \equiv h_{it}$.
- (ii) On the left-hand side of (1.14), the *unconditional* probability of working,

$$P_{i,t-1} \left[\delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it}) = 1 \right].$$

Clearly, health can affect the unconditional probability of working through three channels. The first channel is preference (i.e., agent's utility). The second is the mechanism or mechanisms represented by the unspecified component, y_{it} , (in our example, wage or productivity). The third is uncertainty (i.e., agent's subjective belief about the stochastic process governing health).

On the other hand, health only affects the conditional working probabilities, $P_{i,t-1} \left[\delta_{it}^* = 1 | x_{it} \right] \equiv P_{i,t-1} \left[\delta_{it}^* (x_{it}, y_{it}, \varepsilon_{it}) = 1 | x_{it} \right]$, through the first two channels. This observation is key to interpretation of the *SeaTE*, which is given by the difference in subjective conditional

probabilities of working across values of the specified state component. The 1-period-ahead *SeaTE* of health h on labor supply for individual i at time t is equal to,

$$SeaTE(i, t, 1) = P_{i,t-1}[\delta_{it}^* = 1 | h_{it} = 1] - P_{i,t-1}[\delta_{it}^* = 1 | h_{it} = 0]. \quad (1.15)$$

As long as y may depend on h , in equations (1.13) and (1.14) the agent integrates over the future values of y_{it} that are consistent with the values of h_{it} fixed in the elicitation task. The main implication for interpretation of the *SeaTE* in equation (1.15) is that in this case the measured effect is a *total* effect. That is, it is the effect of health, operating through all of the mechanisms by which health affects labor supply. In our working illustration, it is the effect of health on labor supply through both utility and productivity.²²

B. Econometric Implementation: *Ex ante* and Conditional Value Functions

To implement the model econometrically it is useful to re-write the conditional choice probabilities in (1.14) in terms of the *ex ante* value function and the conditional value function. The DP literature makes specific additivity and orthogonality assumptions that permit estimation.

Following Arcidiacono and Ellickson (2011), the *ex ante* (or integrated) value function at a generic future time t , $\bar{V}_{it}^*(x_{it})$, is the continuation value of being in state x_{it} obtained by integrating $V_{it}(x_{it}, y_{it}, \varepsilon_{it})$ over y_{it} and ε_{it} , that is,

$$\begin{aligned} \bar{V}_{it}^*(x_{it}) &= \int \int_{\varepsilon_{it} y_{it}} V_{it}^*[(x_{it}, y_{it}, \varepsilon_{it}), \delta_{it}^*(x_{it}, y_{it}, \varepsilon_{it})] \cdot \pi_{i,t-1}^y(y_{it} | x_{it}, y_{i,t-1}, d_{i,t-1}) dy_{it} \cdot \pi^\varepsilon(\varepsilon_{it}) d\varepsilon_{it} \\ &= \int \int_{\varepsilon_{it} y_{it}} \left[u_{it}(x_{it}, y_{it}, d_{it}) + \varepsilon_{it}(d_{it}) \right] + \beta \sum_{x_{i,t+1}} \bar{V}_{i,t+1}^*(x_{i,t+1}) \pi_{it}^x(x_{i,t+1} | x_{it}, y_{it}, d_{it}) \cdot \pi_{i,t-1}^y(y_{it} | x_{it}, y_{i,t-1}, d_{i,t-1}) dy_{it} \cdot \pi^\varepsilon(\varepsilon_{it}) d\varepsilon_{it}. \end{aligned} \quad (1.16)$$

This formulation assumes additivity and serial independence of the residual component ε_{it} in order to deliver the standard single-crossing result for a discrete choice problem.²³ In Arcidiacono and Ellickson's setting, the econometrician is doing the integration, while in ours it is the respondent.

The conditional value function $v_{it}(x_{it}, y_{it}, d_{it})$ is the present discounted value net of ε_{it} of choosing d_{it} and behaving optimally from period $(t+1)$ onward, that is,

$$v_{it}(x_{it}, y_{it}, d_{it}) = u_{it}(x_{it}, y_{it}, d_{it}) + \beta \sum_{x_{i,t+1}} \bar{V}_{i,t+1}^*(x_{i,t+1}) \pi_{it}^x(x_{i,t+1} | x_{it}, y_{it}, d_{it}). \quad (1.17)$$

²² One could decompose the effect of health that operates through wages versus other factors by conditioning on wages and health jointly in the elicitation task. We are pursuing this approach in future surveys.

²³ Additionally, its transition probability is normalized to be the same across individuals without loss of generality.

The conditional value function is a key component for forming the conditional choice probabilities that we measure and that use as a basis for estimation of the parameters of the simple structural model that we specify below. Specifically, equation (1.13) can be re-written as

$$\delta_{it}^*(x_{it}, y_{it}, \varepsilon_{it}) = \arg \max_{d_{it} \in \{1,0\}} [v_{it}(x_{it}, y_{it}, d_{it}) + \varepsilon_{it}(d_{it})], \quad (1.18)$$

and the conditional choice probabilities in equation (1.14) can be re-written as

$$\begin{aligned} & P_{i,t-1} \left[\delta_{it}^*(x_{it}, y_{it}, \varepsilon_{it}) = 1 \mid x_{it} \right] \\ &= \int \int_{\varepsilon_{it} y_{it}} \delta_{it}^*(x_{it}, y_{it}, \varepsilon_{it}) \cdot \pi_{i,t-1}^y(y_{it} \mid x_{it}, x_{i,t-1}, y_{i,t-1}, d_{i,t-1}) dy_{it} \cdot \pi^\varepsilon(\varepsilon_{it}) d\varepsilon_{it} \\ &= \int \int_{\varepsilon_{it} y_{it}} \arg \max_{d_{it} \in \{1,0\}} [v_{it}(x_{it}, y_{it}, d_{it}) + \varepsilon_{it}(d_{it})] \cdot \pi_{i,t-1}^y(y_{it} \mid x_{it}, x_{i,t-1}, y_{i,t-1}, d_{i,t-1}) dy_{it} \cdot \pi^\varepsilon(\varepsilon_{it}) d\varepsilon_{it}. \end{aligned} \quad (1.19)$$

Since we measure $P_{i,t-1} \left[\delta_{it}^*(x_{it}, y_{it}, \varepsilon_{it}) = 1 \mid x_{it} \right]$ directly, equation (1.19) links our data to the primitives of the model.

C. Interpreting Conditional Probability Responses in DP Framework

We now present a simple model to illustrate how our dynamic programming framework can be used to elicit information about individual-specific valuations using the standard assumption that underlie econometric implement of the DP framework. As above, we treat work and health as binary.

First consider the case where there is no unspecified variable. As in the discussion of decision tree (Figure 1), let $d_{it} \in \{W, \sim W\}$ be work and not work and health state $x_{it} = h_{it} \in \{H, L\}$ be high or low.²⁴ Because there are only four combinations of health states and labor decisions at time t , it is easy to write out the problem. Define $V_{it}(h_{it}, d_{it})$ to be the value of individual i of being in state h and making choice d at time t given expectation and optimization from $t+1$ onward. Then

$$\begin{aligned} V_{it}(h_{it}, d_{it}) &= (1-h_{it}) \left\{ d_{it} (v_{it}(H, \cdot, W) + \varepsilon_{it}(W)) + (1-d_{it}) (v_{it}(H, \cdot, \sim W) + \varepsilon_{it}(\sim W)) \right\} \\ &\quad + h_{it} \left\{ d_{it} (v_{it}(L, \cdot, W) + \varepsilon_{it}(W)) + (1-d_{it}) (v_{it}(L, \cdot, \sim W) + \varepsilon_{it}(\sim W)) \right\} \end{aligned} \quad (1.20)$$

where the first row refers to actions in high health and the second row in low health.

²⁴ Recall that the indicators $d_{it}=1$ corresponds to working and $h_{it}=1$ corresponds to low health.

Given this standard dynamic programming formulation, maximizing $V_{it}(h_{it}, d_{it})$ yields the standard single crossing conditions for the specified health states as follows.

$$\begin{aligned}
&\text{When } H \\
&\quad W \text{ if } 0 \leq v_{it}(H, \cdot, W) - v_{it}(H, \cdot, \sim W) + \varepsilon_{it}(W) - \varepsilon_{it}(\sim W) \\
&\quad \sim W \text{ otherwise} \\
&\text{When } L \\
&\quad W \text{ if } 0 \leq v_{it}(L, \cdot, W) - v_{it}(L, \cdot, \sim W) + \varepsilon_{it}(W) - \varepsilon_{it}(\sim W) \\
&\quad \sim W \text{ otherwise}
\end{aligned} \tag{1.21}$$

Again, as traditional in dynamic programming applications, we define objects that are the *differenced conditional value functions* and of the residual components. Let

$$\begin{aligned}
\tilde{v}_{it}^H &= v_{it}(H, \cdot, W) - v_{it}(H, \cdot, \sim W) \\
\tilde{v}_{it}^L &= v_{it}(L, \cdot, W) - v_{it}(L, \cdot, \sim W) \\
\tilde{\varepsilon}_{it} &= \varepsilon_{it}(W) - \varepsilon_{it}(\sim W)
\end{aligned} \tag{1.22}$$

where again the differencing is across the working and not working decision. Note that because the residual components across decisions are independent of elements of the state vector, then their difference ε_{it} is also independent. In terms of these variables, the single cross conditions for working given health state become

$$\begin{aligned}
&\text{When } H, \delta_{it}^* = 1 \text{ if } 0 \leq \tilde{v}_{it}^H + \tilde{\varepsilon}_{it} \\
&\quad = 0 \text{ otherwise} \\
&\text{When } L, \delta_{it}^* = 1 \text{ if } 0 \leq \tilde{v}_{it}^L + \tilde{\varepsilon}_{it} \\
&\quad = 0 \text{ otherwise}
\end{aligned} \tag{1.23}$$

where we are using the 0/1 notation for working to be transparently analogous to discrete choice dynamic programming econometrics.

We now show how conditional probability estimates can measure the differenced conditional value functions with arbitrary heterogeneity across individuals. At the time of elicitation, the survey respondents need to integrate out the residual component. Therefore, to analyze the conditional probabilities, we make a distributional specification for the residual uncertainty $\tilde{\varepsilon}_{it}$. Denote the cumulative distribution function of $\tilde{\varepsilon}_{it}$ as Φ . Since the differenced conditional value function (and the underlying utility functions) are only defined up to scale and location, we can take the Φ to be zero mean and unit variance without loss of generality as in the standard discrete choice model.

The survey's elicitation task maps precisely into the discrete choice problem in (1.23). Specifically, the question

If your health is excellent/very good/good two years from now, what are the chances that you will be working for pay?

yields

$$P_{i,t-1}^H \equiv P_{i,t-1} \left[\delta_{it}^* = 1 \mid h_{it} = H \right] \quad (1.24)$$

and

If your health is fair/poor two years from now, what are the chances that you will be working for pay?

yields

$$P_{i,t-1}^L \equiv P_{i,t-1} \left[\delta_{it}^* = 1 \mid h_{it} = L \right] \quad (1.25)$$

Then (1.23) given the distributional assumption for $\tilde{\varepsilon}_{it}$ implies

$$\begin{aligned} P_{i,t-1}^H &= \Phi(\tilde{v}_{it}^H) \\ P_{i,t-1}^L &= \Phi(\tilde{v}_{it}^L). \end{aligned} \quad (1.26)$$

We invert these expressions to yield

$$\begin{aligned} \tilde{v}_{it}^H &= \Phi^{-1}(P_{i,t-1}^H) \\ \tilde{v}_{it}^L &= \Phi^{-1}(P_{i,t-1}^L). \end{aligned} \quad (1.27)$$

Given a functional form for Φ , the individually-elicited conditional probabilities yield individual-level measures of the conditional value of working versus not working in high and low health.²⁵ To implement (1.27), one needs to specify a distribution such as Φ . In what follows, we will specify the distribution as normal.

Note that this case, where there is no unspecified component of the state vector y , covers many cases of interest. The value function may shift for multiple reasons given health. Preferences for work versus leisure may be a function of health; wages may be a function of health; medical costs may be a function of health. If these are all deterministic functions of health for an individual then the value functions given in (1.27) will completely characterize decision making. Hence, they can be used to infer the causal effect of health. One might, however, want separate measurements based on y . There would be two reasons for these separate measurements. The first would be to separate the channels for health affecting work, e.g., separating the role of preferences and wages. The second would be if the nonspecified state

²⁵ Note that there are two level of conditioning. First, the netting of the residual uncertainty defines the conditional value function. On top of this, our elicitation task is done fixing health.

y were related to health, though not perfectly. The appendix shows how the conditional probabilities can be interpreted in this second case.

III. Eliciting Conditional Probabilities: Survey and Basic Results

A. The Vanguard Research Initiative (VRI)

The VRI is a survey-administrative linked dataset on older wealthholders. Survey respondents are account holders at the Vanguard Group who are aged 55 and above, are web-survey eligible, and have at least \$10,000 in financial assets at Vanguard.

As of December 2015, four surveys were completed by a panel of about 3,000+ respondents, with each survey focusing on a different aspect of retirement decision-making. Our analysis is based mainly on Survey 4 (Labor), while Survey 1 (Wealth), Survey 2 (Long-term Care), and Survey 3 (Transfers) provide relevant covariates.

Survey 4 begins by asking whether or not an individual is working. If so, it gets facts about the current job and establishes if it is the career job (Current job battery). If yes, it gets information about whether the individual is searching for another job (On-the-job search battery). If not, it gets information about the career job, separation from it, and subsequent search (Career job, Separation, and Career-to-bridge search batteries). If not working, there is a similar sequence starting with information about last job. This sequence establishes information about career job, bridge job (if relevant), and the transitions and search.

Respondents who were working in either a career job or bridge job at the time of Survey 4 were asked a series of questions regarding their labor supply and health expectations (described below), which constitute the main measures on which the *SeaTE* and related aggregate parameters are based.

B. Selection and Description of the Analytic Samples

We select our sample from respondents who meet the following criteria: (i) who have taken the first 4 surveys of the VRI; (ii) who were working at the time of Survey 4 and, thus, eligible to answer the labor supply and health expectations battery;²⁶ (iv) who gave complete and consistent responses to the latter battery; and (iv) who reported being in high health in Survey 4.²⁷ Table 1 summarizes the selection process and provides the relevant figures at each step.

²⁶ Some of these individuals had already retired from their career job and were working in a bridge job at the time of the survey. These individuals, too, were asked the expectations questions just described with reference to their bridge job.

²⁷ As fewer than 3% of respondents reported being in low health (fair or poor), we decided to exclude this small group and focus on the majority of respondents who reported being in high health (excellent, very good, or good).

The final sample amounts to 970 respondents aged 57 to 81, currently in high health and working. Sample size decreases to 839 respondents for the analysis of expectations with a 4 years horizon, which applies to individuals who reported a positive probability of working in 2 years. Table 2 summarizes the main characteristics of the two analytic samples which we use for the 2 and 4 years analyses.

As easily seen in Table 2, the VRI respondents tend to be wealthier, more educated, and healthier than the general population. However, conditional on the sample screens (age, positive financial wealth, internet access) used to select the sample, they are broadly similar to those from the HRS and the Survey of Consumer Finances (SCF) (Ameriks et al., 2015). Our main analytic sample is further selected as it is made of working respondents in good or better health. This sample therefore represents a subpopulation of particular interest for our analysis, as it is made of individuals who in principle have the capacity to work longer but for whom assessing the causal link between health and retirement is particularly challenging as they have not experienced the health shocks used for identification in standard approaches using realizations data.

C. Elicited Probabilities: Summary Statistics and Internal Validity

Do currently healthy older workers expect to work in 2 (4) years? Do they expect to have the health capacity to work at those horizons? Were they (not) to have the capacity to work in 2 (4) years, how would their labor supply expectations change?

To answer these questions we begin by analyzing respondents' labor supply and health expectations at 2 and 4 years. Specifically, eligible VRI respondents were first asked for the percent chance out of 100 that they will be working in 2 years from the point of the survey. Next, they were asked for their self-rated health on a 5-point scale (Excellent, Very Good, Good, Fair, and Poor) and for the percent chance out of 100 that their health will be some particular state (discussed next) in 2 years.²⁸ Finally, respondents were asked about their probability of working in the next 2 years, conditional on different health states. These questions were then repeated for the 4 years horizon.

²⁸ The VRI uses the standard five point scale for self-assessed health. To economize on the number of questions, conditional expectations module used only partitions that group health states. The partition of future health states used in the questions depends on the current level of health reported by the respondent. For example, respondents in excellent health were asked about their likelihood of being in very good/good and fair/poor health in 2 and 4 years. The partitions used map uniquely to the high (excellent/very good/good) and low (fair/poor) dichotomous classification used in this paper. See Table 3 (VRI Survey 4, health_fill table) for the partitions.

The partition of future health states used in the questions eliciting respondents' health expectations and working expectations given health depends on the current level of health reported by the respondent at Survey 4. For example, consider a respondent who reported being in good health. This respondent was asked the following sequence of questions for the 2 (and 4) years horizon:

- 1) What are the chances that you will be working 2 years from now? [fill-in box]%
- 2) What are the chances that your health will be fair or poor 2 years from now? [fill-in box]%
- 3) What are the chances that your health will be very good or excellent 2 years from now? [fill-in box]%
- 4) If your health is very good or excellent 2 years from now, what are the chances that you will be working for pay? [fill-in box]%
- 5) If your health is good 2 years from now, what are the chances that you will be working for pay? [fill-in box]%
- 6) If your health is fair or poor 2 years from now, what are the chances that you will be working for pay? [fill-in box]%

Table 3 shows the exact partition of future health states used in the questions as a function of self-reported health at Survey 4. The health fills 1 and 2 refer to the health states used in the health expectations questions, whereas the health fills 3-5 refer to the health states used in the questions measuring the respondent's expectations of working given health.

Unconditional working expectations. We analyze 2 and 4 years working expectations of these individuals elicited on a 0-100 scale of chance, where 0 means "no chance of working 2 (4) years from now" and 100 means "will work for sure 2 (4) years from now." Sample distributions (mean, standard deviation, and main quantiles) of survey answers are shown in Table 4.

Respondents' working expectations at 2 and 4 years are heterogeneous and span the whole support of 0-100 percent chance scale. In column 1, over a third of the respondents (37.42%) expects that they will work for sure in 2 years, as opposed to the almost 10% of those who expect

not to work for sure. The remaining majority of respondents (approx. 53%) views working in 2 years as an uncertain event. Nearly 12% of respondents report that they have 50 chances out of 100 of working in 2 years;²⁹ a similar fraction gives a percent chance below 50 percent; whereas almost 30% gives a percent chance above 50 percent (the second most frequent response category after 100 percent). Indeed, the median belief of 90 percent is quite high, indicating that half of the respondents expect to work in 2 years with a likelihood of 90 percent or higher.

Column 2 shows main sample statistics for “calculated” unconditional working expectations, obtained by combining respondents’ working expectations conditional on health and their health expectations according to the law of total probability. In this case, the fraction of respondents who think that they will work for sure decreases to 17%, and the fraction of unsure respondents increases to 73%. However, the median probability is still very high and equal to 80 percent. Moreover, the medians, means, and standard deviations are very similar across the two distributions.

As the time horizon increases from 2 to 4 years in column 3, the fraction of respondents who think that they will not work for sure and the fraction of uncertain respondents increase respectively to 12 and 77 percent, whereas the fraction of respondents who think that they will work for sure decreases to 11 percent. While respondents’ subjective probabilities of working in 4 years are still fairly high overall (e.g., the 75th percentile is equal to 90 percent), these probabilities tend to be lower than those with a 2 years horizon (e.g., the median is down to 50 percent and the mean to 52.7 percent).

Health expectations. Working VRI respondents are healthy; over 97% of them reported being in excellent, very good, or good health in Survey 4. Do these healthy working individuals expect to stay healthy or do they anticipate health declines?

To answer this question we analyze respondents’ expectations of being in high versus low health in 2 and 4 years elicited on the usual 0-100 scale of chance.

Features of the sample distributions (mean, standard deviation, and main quantiles) of percent-chance expectations are shown in Table 5 for all respondents in our analytic sample. Expectations are constructed from questions asking respondents to report the percent chance that they will be in a specific health state (or set of states) in 2 and 4 years, where the possible states are excellent, very good, good, fair, and poor and where excellent, very good, and good were

²⁹ Some of these respondents might be using 50 percent as an expression of “epistemic uncertainty” (e.g., see Fischhoff and Bruine de Bruin (1999) and Bruine de Bruin et al. (2002)), thereby conveying that they don’t know or are unsure about their chances of working in two years.

subsequently consolidated into “high health” and fair and poor were consolidated in “low health.”

Respondents’ expectations of being in high health in 2 and 4 years range between 20 and 100 percent. Three fourths of the respondents report subjective likelihoods of being in high health in 2 years equal to or above 75 percent and of being in high health in 4 years equal to or above 70 percent. In fact, 10% of respondents expect to be in high health for sure in 2 years and 5% of them think that they will be in high health for sure in 4 years. Large fractions of healthy respondents express some uncertainty about their future health by reporting expectations strictly between 0 and 100 percent; about 90% for the 2 years horizon and 95% for the 4 years horizon.

Overall, the majority of respondents appear fairly optimistic about their chances of remaining in high health – thus maintaining their capacity to work – as opposed to experiencing a health decline. The means of the distributions of respondents’ 2 years subjective probabilities of high and low health imply a forecast of the proportion of high health individuals equal to 83.4 percent and a forecast of the proportion of low health individuals equal to 16.6 percent. Similarly, the means of the 4 years distributions imply $P(H) = 76.5$ percent and $P(L) = 23.5$ percent respectively.

Conditional working expectations given health. We additionally elicited respondents’ 2- and 4-year ahead working expectations *conditional* on alternative health scenarios by means of the familiar 0-100 scale of chance.

Table 6 displays the mean, standard deviation, and main quantiles of the sample distributions of percent-chance conditional working expectations for both health scenarios and prediction horizons.

The sample distribution of working expectations conditional on experiencing high health in 2 years, summarized in Column 1, is similar to its unconditional counterpart. In particular, about 40% of respondents expect that they will work for sure in 2 years if their health is high; this proportion is 3 percentage points higher than the 37% of respondents who report an unconditional working probability of 100 percent. The remaining respondents are split between a group close to 10% who gives a 0 percent chance of working in 2 years if in high health and the remaining 50% who gives a percent chance strictly between 0 and 100 percent. As the first proportion is virtually identical to the proportion of 0s observed for the unconditional working probability, the second proportion is 3 percentage points lower than the corresponding proportion

of responses to the unconditional question, mechanically compensating for the 3% increase in the 100 percent answers.³⁰

The mean of the distribution of conditional working expectations given high health is equal to 70.5, higher than both the mean of the distribution of unconditional working probabilities as directly reported by respondents (equal to 69.8 percent) and the mean of the distribution of calculated unconditional working probabilities (equal to 65.9 percent). The median working probability conditional on high health is equal to 90 percent and identical to the median probability of the distribution of self-reported unconditional working probabilities.

The close similarity between reports of unconditional working probabilities and conditional working probabilities given high health at 2 years is consistent with the observed high probability that the majority of respondents assigns to the event of experiencing high health in 2 years time. On the other hand, the distributions of the 4 years working probabilities, unconditional and conditional on high health, are less similar to each other than the 2 years distributions. This is also not surprising, as respondents' expectations of experiencing high health in 4 years are lower than the 2 years expectations, while featuring greater cross-sectional variation.

As the time horizon increases from 2 to 4 years between column 1 and column 3, the fraction of respondents who think that they will not work for sure and the fraction of uncertain respondents increase respectively to 12 and 57 percent, whereas the fraction of respondents who think that they will work for sure decreases to 31 percent. While the proportion of 0s is virtually identical to that observed for the distribution of the 4 years unconditional working probabilities in Table 4, the proportion of 100s is dramatically higher (31% versus 11%). In fact, both the mean and median probabilities of working in 4 years when experiencing high health (equal to 58.7 and 68 percent respectively) are visibly higher than the corresponding mean and median probabilities of the unconditional distribution (equal to 52.7 and 50 percent).

Having respondents entertain a scenario of low health – a negative and unlikely shock for the majority of them – lowers substantially their self-reported working expectations at both 2 years (in column 2) and 4 years (in column 4). For example, in the 2 years case the median of the distribution of the conditional working probabilities drops from 90 to 40 percent between high and low health states. Similarly, the mean drops from 71 to 42 percent. In the 4 years case, the median drops from 68 to 20 percent and the mean from 59 to 33 percent.

³⁰ The fraction of 50 percent decreases slightly from 12% to 9%. A possible interpretation of this finding is that conditioning on future health reduces respondents' perceived uncertainty about their chances of working in the future.

Entertaining a low health scenario appears to additionally increase respondents' perceived uncertainty relative to the high health scenario; the fraction of respondents giving a response strictly between 0 and 100 percent increases from 50% to 64% under the 2 years horizon and from 58% to 62% under the 4 years horizon.

Validation. Before moving to the analysis of the *SeaTE*, we perform some additional analysis of the expectations responses in order to provide immediate evidence supporting the validity of our measures.

Heterogeneity by age. As conditioning on health removes a potentially important source of uncertainty that respondents would otherwise need to factor in when answering the unconditional question, we hypothesize that subjective conditional probabilities might be even better predictors of conditional outcomes than unconditional probabilities are of the unconditional outcomes.

A direct comparison between survey reports of the subjective conditional probability of working in 2-4 years with the conditional outcomes will become possible when the future waves of the VRI are fielded and participants' labor supply observed. We do know, however, that labor supply features distinctive age peaks. Do working expectations track such peaks? To investigate this question, in Figures 2A-2C we create box-and-whiskers plots of the working expectations (on the y-axis) by current age of the respondent (on the x-axis). Age bins 60-61, 63-64, and 65 in Figures 2A and 2B are of particular interest, as a 2-year horizon from those ages implies the crossing of the early, normal, and full SS retirement ages (i.e., 62, 65, and 67), where actual labor supply displays well-known peaks. (Similarly, for age turns ≤ 59 , 60-61 & 62, and 63-64 at the 4 years horizon in Figure 2C.)

In Figures 2A and 2B, the mean and median working expectations at 2 years feature sharp declines among the 60-61 years old (corresponding to the 62 peak), among the 63-64 years old (corresponding to the 65 peak), and among the 65 years old (corresponding to the 67 peak). Notice, however, that the mean and median working expectations do not decrease monotonically across groups of increasing age. This is consistent with increasing selectivity of the working and (high) health requirements applying to older respondents. An interesting additional feature of these age-specific distributions is that the cross-sectional variability of the subjective unconditional working expectations tends to increase with age, especially between respondents aged 62 or younger and respondents older than 62.

Moving to the 4-year-ahead horizon, inspection of Figure 2C reveals that the age-specific mean and median decrease sharply and steadily from the ≤ 59 and the 63-64 groups and level off (or tend to increase slightly) thereafter, again consistent with increasing selectivity of older sub-groups. The cross-sectional variance of working expectations is now fairly high in all age groups and appears higher than the cross-sectional variance of the 2 years working expectations. This is consistent with a bigger role of heterogeneity as the forecasting horizon increases.

Figures 3A-D show equivalent plots using the 2- and 4-year-ahead working probabilities given health.

Law of Total Probability (LTP). Since we ask the probability of work given health, the probability of health, and the unconditional probability of work, we are able to evaluate how well survey respondents obey the law of total probability in their responses. The respondents are quite good at applying the LTP. Figure 4 gives plots of the reported unconditional probability of work versus that implied by the LTP for the 2-year horizon using box and whiskers plots for various bins.³¹ A large majority of the observations lies very close to the 45-degree line, corresponding to the case in which the self-reported probability and the calculated one are equal to each other. The correlation between the two measures is approximately 0.95. For those responses that deviate, the deviations are relatively small. Therefore, we conclude that overall respondents appear to understand the logic of probabilities. This finding allows us to counter potential concerns that the probabilistic reasoning needed to answer the questions is too difficult for our respondents.

The deviations from consistency of the LTP is most pronounced for respondents with self-reported unconditional working probabilities equal to 0, 50, or 100 percent. This finding is consistent with the suggestion in the literature that some respondents who give corner or 50/50 responses are less good at probabilistic thinking. Note, however, that there are significant mass points of respondents who get it exactly right at 0 and 100. Though not readily apparent because of the heaping at the corners, this vast majority of corner respondents are getting the LTP exactly right.

IV. Empirical Analysis of *SeaTE*

A. Distributions of Individual-Specific *SeaTE*

³¹ The survey did not ask the unconditional probability of work for the 4-year horizon.

How does working longer depend on health? What is the distribution of these causal effects of health on work for these workers? To answer these questions in Table 7 we report the sample distributions of the individual-level *SeaTE* for the 2-year horizon (column 1) and the 4-year horizon (column 2). Approximately 70% of the respondents gives a conditional subjective probability of working given low health which is strictly smaller than their subjective probability of working if in high health regardless of the horizon (i.e., $SeaTE < 0$). The vast majority of the remaining respondents (28-29% depending on the horizon) gives the same probability of working under the two health scenarios (implying a $SeaTE = 0$). Only a negligible fraction of respondents report a positive *SeaTE*. Note that a positive *SeaTE* is a theoretical possibility, e.g., if the value of leisure falls or need to work increases in low health.

Conditional on a negative *SeaTE* (columns 3 and 4), the distribution of the (absolute values of the) *SeaTE* ranges between 2-4 percent and 100 percent regardless of the horizon, thus displaying substantial heterogeneity across respondents. The absolute values of both the median and mean *SeaTE* are somewhat higher in the 2-year horizon (40 and 41 percent) than in the 4-year horizon (30 and 37 percent), possibly because a negative health shock is a stronger bite at younger ages than at older ages, when additional motivations for stopping working might kick in.

B. Estimates of *ASeaTE*, *ASeaTT*, and *ASeaTU*

Aggregating up the individual-level *SeaTE*s yields subjective *ex ante* versions of population parameters well-known in the treatment-effect literature. Specifically, similar to Arcidiacono et al. (2017) and Wiswall and Zafar (2016), we compute estimates of the average subjective *ex ante* treatment effect (*ASeaTE*), the average subjective *ex ante* treatment effect on the treated (*ASeaTT*), and the average subjective *ex ante* treatment effect on the untreated (*ASeaTU*) using equation (1.9) The calculated values of the three parameters for the two horizons are shown in Table 8.

ASeaTT and *ASeaTU* are very close to each other, providing little support for selection. That is, individuals who are more likely to experience low health in the future expect similar impacts of a negative health shock on their future work-retirement decisions to individuals who are less likely to enter low health. We return to this point in Section E where we simulate behavioral data based on our DP framework and elicited probabilities.

C. Observed Heterogeneity and Predictors of *SeaTE*

In Table 9 we investigate whether the individual-level *SeaTEs* for the 2- and 4-year horizon vary by respondents' characteristics. We report estimates from best linear predictions of 2- and 4-year ahead *SeaTEs* on the covariates listed in Table 1. The reference group corresponds to male respondents, aged 59 or younger, who have attained a high school diploma or a lower degree, currently working in their career job within the Management & Professional sector, who are not partnered, and who are in the highest quintile of the distributions of: total household wealth, current salary, and replacement rate.

At both horizons, the *SeaTE* tends to be larger (in absolute value) among older respondents, more educated respondents, partnered respondents, respondents with a working spouse, and respondents at lower quintiles of the household wealth, replacement rate, and salary distributions. The *SeaTE* tends to be smaller (in absolute value) among female respondents and respondents working in a bridge job or in lower-level sectors. These associations, however, are statistically significant only for respondent's age and for selected quintiles of the replacement rate (in the 2-year case) and of the total household wealth (in the 4-year case). This suggests that most of the heterogeneity featured by respondents' *SeaTEs* cannot be explained by heterogeneity in respondents' observable characteristics. Hence, econometric approaches that use controls such these will still leave unobserved very substantial heterogeneity.

D. Estimates using Dynamic Programming Framework

We now proceed with analysis derived from the dynamic programming specification of Section II. The elicited conditional probabilities yield individual-level values of working versus not working given specified health. Hence, the approach introduced in this paper can deliver individual-level evidence on how the work/not-work choice shifts from high to low health. We now summarize what the survey responses imply about how the value of working versus not working shifts with health.

Table 10 shows summary statistics for these values \tilde{v}^H and \tilde{v}^L from equation (1.27) for the 2-year and 4-year-ahead conditional probabilities.³² As expected, the mean value in high health is substantially greater than that in low health reflecting the lower value of working in low health. The conditional probability of working is reflected in the last row of the table showing the

³² Note for the purpose of the analysis in this paper, we are treating both the horizons as different "one-period ahead" expectations. That is, we are not modeling the transition from 2 to 4 year, but instead presenting them as separate (though obviously related) measurements.

fraction who expect to work in the specified health state. For the 4-year ahead horizon, there is a substantial shift down in the willingness to work in both the high-health and low-health states.

Figure 5A shows a scatter plot of the values \tilde{v}^H and \tilde{v}^L for the 2-year horizon. Figure 5B shows an analogous plot for the 4-year horizon. These plots illustrate many features of the value of work conditional on health across respondents. In each of the two figures, the upper right quadrant contains the individuals who value work more than not work in both health states (where of course value is net of the residual uncertainty that will be realized at the time of the decision). The lower left quadrant has those who value work less in both states. The vast majority of individuals lie below the 45 degree line corresponding to having a lower value of work relative to not work when in low health than in high health. It is not surprising that values shift in this direction. Lower health likely decreases taste for work and the return to work. Yet, shifting in the other direction is perfectly consistent with optimization. For those above the 45 degree line, the relative attractiveness of work increases in low health. This valuation could result from need for insurance, lower value of leisure in low health, or need for income in low health. Indeed, there are a few observations in the upper left quadrant where the value of working is higher in low health than in high health. The opposite—in the lower right quadrant—is not surprisingly much more common. These represent the individual who would quit working after a negative health shock.

There is a strong correlation between the value of work across the health states. A simple framework of summarizing is that there is a value of work in high health that is positively, but imperfectly correlated with that in low health. Consider the model

$$\begin{aligned}\tilde{v}_i^H &= \alpha^H + v_i^H \\ \tilde{v}_i^L &= \alpha^L + \gamma v_i^H + v_i^L,\end{aligned}\tag{1.28}$$

where α^H and α^L are the mean across individuals of the values and v_i^H is the mean-zero heterogeneity across individuals in the value in the high health state. The heterogeneity in the low health states has two components: a component correlated with that in high health γv_i^H and an orthogonal component v_i^L .

Again, from the point of view of the respondent, these components are nonstochastic. Our procedure gives a direct measurement of the LHS of equation (1.28). The orthogonal decomposition is a convenient way to summarize the observed heterogeneity.

Table 11 presents estimates of the parameters specified in equation (1.28).³³ The estimates obtained are sensible. Consider first the estimates for the 2-year ahead horizon.

- The mean utility from work shifts substantially downward when health changes from high to low. The mean is 0.97 in high health and -1.04 in low health.
- The correlation within individual of the willingness to work across health states is fairly high, but far from unity. The coefficient γ that controls this correlation is 0.71. Hence, there is persistence within individuals of valuation of work across health states, implying that those with high value of work in high health carry that over into low health, but in a damped way.

For the 4-year ahead horizon, there is a substantial shift down in the willingness to work in the high-health state—from 0.97 to 0.48. In contrast, no shift in the willingness to work in the low-health state (-1.04 for both horizons). For interpreting these result, it is important to bear in mind that the estimate is based on a single cross section. The willingness to work declines sharply with age in the age range of the VRI sample, and this decrease is much greater in low health. The estimate of the coefficient γ , controlling the correlation within individual of the willingness to work across health states, decrease slightly to 0.66.

In Table 12, we incorporate covariates by modelling the mean utility from work as a linear index of respondents' observable characteristics. In particular, we include all variables described in Table 2 and used as predictors in Table 9.

Inspection of these estimates reveals that the mean utility from work tends to decreases with age in both health states. Similarly, the mean utility from work also decreases with being partnered in both health states, but the estimated coefficient is strongly significant in high health and for the 2-year horizon only.

On the other hand, the mean utility from work tends to increase with having a working spouse or partner, but the estimated coefficient is significant only in high health and for the 2-year horizon. The mean utility from work also tends to increase across quantiles of the wealth distribution, but the estimated coefficients are significant in high health only.

In this specification with covariates, the estimates of the coefficient γ , controlling the correlation within individual of the willingness to work across health states, decrease slightly but are still substantial, 0.64 and 0.61 at 2 and 4 years respectively.

³³ These estimates are from an OLS regression where the first equation just has a constant and the second equation has a constant plus the residual from the first equation. Note that this is a random coefficient model, though we do not have to specify a distribution since the values are observable.

E. Relating Conditional Probabilities to Realizations using Dynamic Programming

We can use the empirical results from the dynamic programming framework to illustrate the benefit of having data on the heterogeneity of values as opposed to data on realized behavior. Consider the estimate of a linear regression model using data on working d_i and health h_i realizations,

$$d_i = b_0 + b_1 h_i + e_i. \quad (1.29)$$

As discussed in Section I, the least squares estimate of b_1 will be an unbiased estimate of the ATE only when health is exogenous. In the terms of our dynamic programming model, exogeneity will fail when realized health is correlated across individuals with the values \tilde{v}_i^H or \tilde{v}_i^L .

Our elicitation approach is designed to render this heterogeneity observable *ex ante*. Specifically, the average *SeaTE* defined in (1.8) will be an unbiased estimate the causal effect of health b_1 even if there is heterogeneity which would be unobserved with conventional data on realizations of health and work.

To demonstrate how biased estimates of causal effects can emerge in data on realizations, we use our framework to construct simulated realizations of work decisions and health states. Using the dynamic programming model of Section II, equation (1.23) implies that the realized decision to work is

$$d_i = (1 - h_i) \mathbf{I}[\tilde{v}_i^H + \tilde{\varepsilon}_i] + h_i \mathbf{I}[\tilde{v}_i^L + \tilde{\varepsilon}_i], \quad (1.30)$$

where $\mathbf{I}[\cdot]$ is the indicator function, equal to 1 if the argument is positive and zero otherwise, and $d_i = 1$ if work and 0 otherwise. To simulate realizations that reflect the observed heterogeneity, we use the measured conditional value functions (\tilde{v}_i^H or \tilde{v}_i^L) and simulated realizations of health (h_i), and the residual component ($\tilde{\varepsilon}_i$) to calculate simulated decisions according to equation (1.30). Health is modeled as (0,1), so it is simulated using Bernoulli draws based on the health transition probability π_i^h . Consistent with the implement of the DP formulation, $\tilde{\varepsilon}_i$ is simulated as standard normal.

We consider three cases for correlation of the health transition probability π_i^h :

1. π_i^h is fixed at the sample mean, so health transitions are uncorrelated with the value of work.

2. π_i^h is the individual-specific probabilities, so health transitions have the empirical correlation with the value of work.
3. π_i^h adjusts the individual-specific probabilities to induce an higher correlation between health and the value of work than is present in the VRI data.³⁴

The first case implies health is exogenous, so the OLS estimate of (1.29) will yield an unbiased estimate of the average treatment effect equal to the average *SeaTE*. The second case will illustrate the extent of the bias that would be present in the VRI data. The third case magnifies the bias.³⁵

Table 13 shows estimates of the regression for simulated realization for the 2- and 4-year horizons simulated over 1000 replications for the three cases. The uncorrelated cases where the health transition probability is set to be equal to its mean are quasi-experimental, so the estimated coefficient of health is unbiased and therefore equals the average *SeaTE* shown in Tables 7 and 8.

The empirical cases yield a biased estimate because of the positive correlated heterogeneity in value of work and health transitions in the VRI. There is a slight, positive correlation between the value of work and the probability of being in high health. The sign of this correlation is not surprising because individuals in situations with attractive jobs (e.g., high SES) are also likely to have better health. The estimated coefficient of health is larger in absolute value than the causal effect because those who get bad health shocks disproportionately have lower value of work. The VRI respondents do not have that much heterogeneity in health (because most are quite healthy), so the magnitudes are small. Even so, the bias is nontrivial implying almost 10% more health-related job transitions than the causal effect.

In other samples with more heterogeneity in health, this bias would be even more important as illustrated by the higher correlation case.

Finally, recall that Table 12 shows that there is substantial heterogeneity in the values even after conditioning a rich set of covariates. Therefore, conditioning on such covariates in

³⁴ Specifically, the simulations are based on adjusting the π_i^h by subtracting 0.1 from individuals in the bottom quintile of \tilde{V}_i^H , subtracting 0.05 from those in the second quintile of \tilde{V}_i^H , leaving the middle quintile unadjusted, and adding 0.075 to the top two quintiles of \tilde{V}_i^H . (The top two quintiles are combined because they have a common \tilde{V}_i^H corresponding to individual who gave a 100% change of working when in high health.)

³⁵ Note that this simulation is carried out under the assumption that unspecified states are ignorable. See appendix for discussion of deviations from this case.

econometric applications using data on realized decisions and states, though helpful, is not likely to eliminate bias from uncontrolled heterogeneity.

IV. Conclusion

In this paper, we provide a novel strategy for assessing the causal effect of a negative health change on the labor supply of healthy older workers, based on individuals' own estimates of their working probabilities at specified horizons under alternative scenarios about their future health. Since these effects are subjective and *ex ante* in nature, we have called them the Subjective *ex ante* Treatment Effect (SeaTE).

Since these effects are obtained from respondents' subjective probability of working if they enter low health minus their subjective probability of working if they remain in high health, they are individual-level effects. By aggregating these effects across individuals we have additionally derived estimates of population-level parameters, including the average SeaTE (ASeaTE) and the average subjective *ex ante* treatment effect on the treated and on the untreated (ASeaTT and ASeaTU respectively).

We have shown the elicited probabilities can be interpreted in the standard dynamic programming framework for analyzing labor supply. Transformations of the conditional probabilities of working given health equal the *ex ante* value of working versus not working given health. We use this framework to simulate realization of health and work given the survey responses that fully reflect the heterogeneity in preferences and expectations that go into the decision. There is correlation in the value of work and health that would be hard to disentangle in behavioral data when attempting to estimate the causal effect of health on retirement. The simulations show that the bias can be considerable.

Appendix.

Dynamic Programming Interpretation with Correlated Unspecified States

We consider the interpretation of the conditional probabilities using the DP framework when there is a unspecified state y that is correlated with health. This case is distinct from the residual uncertainty ε that is additive in the value function and orthogonal to health. Extending the case in Section IIC, suppose that the unspecified state is also binary. To model correlation with health, assume it can take on two values y^{+H}, y^{-H} if health is high and potentially two different values y^{+L}, y^{-L} if health is low. Let the probability of y given health to be

$$\begin{aligned} P(Y^{+H} | H) &= \pi^{+H} \\ P(Y^{-H} | H) &= 1 - \pi^{+H} \\ P(Y^{+L} | L) &= \pi^{+L} \\ P(Y^{-L} | L) &= 1 - \pi^{+L} \end{aligned} .$$

Then equation (1.26) becomes

$$\begin{aligned} P_{i,t-1}^H &= \Phi(\tilde{v}_{ii}^{+H})\pi^{+H} + \Phi(\tilde{v}_{ii}^{-H})(1 - \pi^{+H}) \\ P_{i,t-1}^L &= \Phi(\tilde{v}_{ii}^{+L})\pi^{+L} + \Phi(\tilde{v}_{ii}^{-L})(1 - \pi^{+L}) \end{aligned}$$

where

$$\begin{aligned} \tilde{v}_{ii}^{+H} &= v_{ii}(H, Y^{+H}, W) - v_{ii}(H, Y^{+H}, \sim W) \\ \tilde{v}_{ii}^{-H} &= v_{ii}(H, Y^{-H}, W) - v_{ii}(H, Y^{-H}, \sim W) \\ \tilde{v}_{ii}^{+L} &= v_{ii}(L, Y^{+L}, W) - v_{ii}(L, Y^{+L}, \sim W) \\ \tilde{v}_{ii}^{-L} &= v_{ii}(L, Y^{-L}, W) - v_{ii}(L, Y^{-L}, \sim W) \end{aligned} ,$$

that is, the differenced conditional value functions under the four possible combinations of health and the unspecified state. Hence, the conditional probability of working given health $P_{i,t-1}^h$ is the weighted average of the conditional probability of working given health and the unspecified state $(\Phi(\tilde{v}_{ii}^{+h}), \Phi(\tilde{v}_{ii}^{-h}))$ with weights equal to the probabilities of the unspecified state give health (π^{+h}, π^{-h}) .

Note explicitly mentioning a state does not necessarily cause the complication given above. For example, consider the leading case for studying health and retirement that has the wage a

function of health. If wage is the only state affecting retirement that is a function of health then the model in the main text applies. In terms of the notation of the appendix, the probabilities of the “unspecified” state give health (π^{+h}, π^{-h}) are degenerate corners, so the expression above collapses to (1.26).

The complication in interpretation discussed in the appendix would, however, arise if health shifts the utility function independent of wage (e.g., taste heterogeneity) and the probabilities (π^{+h}, π^{-h}) are not corners. The conditional probability approach can still be used in this setting, but one would need to elicit the conditional probabilities of working fixing all combinations of health and wage.³⁶

³⁶ An upcoming VRI survey pursues this approach.

References

- Abbring, Jaap H. 2010. "Identification of Dynamic Discrete Choice Models." *Annual Reviews of Economics* 2: 367-394.
- Aguirregabiria, Victor and Pedro Mira. 2010. "Dynamic Discrete Choice Structural Models: A Survey." *Journal of Econometrics* 156: 38-67.
- Ameriks, John and Andrew Caplin and Minjoon Lee and Matthew D. Shapiro and Christopher Tonetti. 2015. "The Wealth of Wealthholders." NBER Working Paper, 20972.
- Ameriks, John and Joseph Briggs and Andrew Caplin and Minjoon Lee and Matthew D. Shapiro and Christopher Tonetti. 2017. "Older Americans Would Work Longer If Jobs Were Flexible." Vanguard Research Initiative Working Paper.
- Angrist, Joshua D. and Guido Imbens. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62(2): 467-475.
- Angrist, Joshua D. and Guido Imbens and Donald Rubin. 1996 "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91(434): 444-455.
- Arcidiacono, Peter and Paul B. Ellickson. 2011. "Practical Methods for Estimation of Dynamic Discrete Choice Models." *Annual Reviews of Economics* 3: 363-394.
- Arcidiacono, Peter and Joseph V. Hotz and Arnaud Maurel and Teresa Romano. 2017. "Ex Ante Returns and Occupational Choice." Working Paper, Duke University.
- Armantier, O., W. Bruine de Bruin, S. Potter, G. Topa, W. Van der Klaauw, and B. Zafar. 2013. "Measuring Inflation Expectations." *Annual Review of Economics* 5: 273-301.
- Arrondel, L. and H. Calvo-Pardo and C. Giannitsarou and M. Haliassos. 2017. "Informative Social Interactions." Working Paper, University of Southampton.
- Athey, Susan and Guido Imbens and Thai Pham and Stefan Wager. 2017. "Estimating Average Treatment Effects: Supplementary Analyses and Remaining Challenges." *American Economic Review: Papers & Proceedings* 107(5): 278-281.
- Attanasio, O. 2009. "Expectations and Perceptions in Developing Countries: Their Measurement and Their Use." *American Economic Review* 99(2): 87-92.
- Auerbach, Alan J. and Kerwin K. Charles and Courtney C. Coile and William Gale and Dana Goldman and Ronald Lee and Charles M. Lucas and Peter R. Orszag and Louise M. Sheiner and Bryan Tysinger and David N. Weil and Justin Wolfers and Rebecca Wong. 2017. "How the Growth Gap in Life Expectancy May Affect Retirement Benefits and Reforms." NBER Working Paper, 23329.

- Banerjee, Sudipto and David Blau. 2013. "Employment Trends by Age in the United States: Why Are Older Workers Different?" MRRC Working Paper, 2013-285.
- Behncke, Stefanie. 2012. "Does Retirement Trigger Ill Health?" *Health Economics* 21(3): 282-300.
- Bellman, Richard E. 1957. *Dynamic Programming*. Princeton University Press.
- Blau, David M. and Donna B. Gilleskie. 2008. "The Role of Retiree Health Insurance in the Employment Behavior of Older Men." *International Economic Review* 49(2): 475-514.
- Blass, Asher A. and Saul Lach and Charles F. Manski. 2010. "Using Elicited Choice Probabilities to Estimate Random Utility Models: Preferences for Electricity Reliability." *International Economic Review* 51(2): 421-440.
- Blundell, Richard and Jack Britton and Monica Costa Dias and Eric French. 2016. The Dynamic Effects of Health on the Employment of Older Workers. MRRC Working Paper, 2016-348.
- Bordalo, P. and N. Gennaioli and Y. Ma and A. Shleifer. 2017. "Overreaction in Macroeconomic Expectations." Working Paper, Bocconi University.
- Bound, John. 1991. "Self-Reported Versus Objective Measures of Health in Retirement." *Journal of Human Resources* 26(1): 106-138.
- Bound, John and Michael Schoenbaum and Todd Stinebrickner and Timothy Waidmann. 1999. "The Dynamic Effects of Health on the Labor Force Transitions of Older Workers." *Labour Economics* 6(2): 179-202.
- Bound, John and Todd Stinebrickner and Timothy Waidmann. 2010. "Health, Economic Resources and the Work Decisions of Older Men." *Journal of Econometrics* 156: 106-129.
- Bruine de Bruin, Wandí and Paul S. Fischbeck and Neil A. Stiber and Baruch Fischhoff. 2002. "What Number is "Fifty-Fifty?" Redistributing Excessive 50% Responses in Elicited Probabilities." *Risk Analysis* 22(4): 713-723.
- Cai, Lixin. 2010. "The Relationship Between Health and Labour Force Participation: Evidence from a Panel Data Simultaneous Equation Model." *Labour Economics* 17: 77-90.
- Carroll, C. 2017. "Heterogeneity, Macroeconomics, and Reality." Keynote Lecture given at the Sloan-BoE-OFR Conference on Heterogeneous Agent Macroeconomics, U.S. Department of the Treasury, September 22, 2017.
- Case, Anne and Angus Deaton. 2015. "Rising Morbidity and Mortality in Midlife Among White Non-Hispanic Americans in the 21st Century." *PNAS*, 112(49): 15078-15083.
- Coe, Norma B. and Gema Zamarro. 2011. "Retirement Effects on Health in Europe." *Journal of Health Economics* 30(1): 77-86.
- Coile, Courtney C. 2015. "Economic Determinants of Workers' Retirement Decisions." *Journal of Economic Surveys* 29(4): 830-853.

- Coile, Courtney and Kevin Milligan and David A. Wise. 2016. "The Capacity to Work at Older Ages: Evidence from the U.S." NBER Working Paper, 21940.
- Cornelissen, Thomas and Christian Dustmann and Anna Raute and Uta Schonberg. 2016. "From LATE to MTE: Alternative Methods for the Evaluation of Policy Interventions." *Labour Economics* 41: 47-60.
- Delavande, A. (2014). "Probabilistic Expectations in Developing Countries." *Annual Review of Economics* 6: 1-20.
- Disney, Richard and Carl Emmerson and Matthew Wakefield. 2006. "Ill Health and Retirement in Britain: A Panel Data-Based Analysis." *Journal of Health Economics* 25: 621-649.
- Dominitz, Jeff. 1997. "The Role of Conditioning Information in Reports of Subjective Phenomena." Working Paper, California Institute of Technology.
- Dominitz, Jeff and Charles F. Manski. 1996. "Student Expectations of the Returns to Schooling." *Journal of Human Resources* 31(1): 1-26.
- Dwyer, Debra and Olivia Mitchell. 1999. "Health Problems as Determinants of Retirement: Are Self-Rated Measures Endogenous?" *Journal of Health Economics* 18(2): 173-193.
- Eckstein, Zvi and Kenneth I. Wolpin. 1989. "The Specification and Estimation of Dynamic Stochastic Discrete Choice Models: A Survey." *Journal of Human Resources* 24(4): 562-598.
- Fischhoff, Baruch and Wandu Bruine de Bruin. 1999. "Fifty-Fifty = 50%?" *Journal of Behavioral Decision Making* 12: 149-163.
- Fisher, Gwenith G., Dorey S. Chaffee and Amanda Sonnega. 2016. "Retirement Timing: A Review and Recommendations for Future Research." *Work, Aging and Retirement* 2(2): 230-261.
- French, Erich. 2005. "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour." *The Review of Economic Studies* 72: 395-427.
- French, Erich and John B. Jones. 2011. "The Effects of Health Insurance and Self-Insurance on Retirement Behavior." *Econometrica* 79(3): 693-732.
- French, Erich and John B. Jones. 2017. "Health, Health Insurance and Retirement: A Survey." *Annual Reviews of Economics* 9: 383-409.
- Fuster, A. and R. Perez-Truglia and B. Zafar. 2017. "Demand for Macroeconomic Data: Evidence from an Information-Acquisition Experiment." Working Paper, New York Fed.
- Garcia-Gomez, Pilar. 2011. "Institutions, Health Shocks and Labour Market Outcomes Across Europe." *Journal of Health Economics* 30: 200-213.
- Garcia-Gomez, Pilar and Titus J. Galama and Eddy van Doorslaer and Angel Lopez-Nicolas. 2017. "Interactions Between Financial Incentives and Health in the Early Retirement Decision." Working Paper.

- Giustinelli, P. and C.F. Manski (2018). "Survey Measures of Family Decision Processes for Econometric Analysis of Schooling Decisions." *Economic Inquiry* 56(1): 81-99.
- Grossman, Michael. 1972. "On the Concept of Health Capital and the Demand for Health." *Journal of Political Economy* 80: 223-255.
- Gustman, Alan L. and Olivia S. Mitchell and Thomas L. Steinmeier. 1995. "Retirement Measures in the Health and Retirement Study." *Journal of Human Resources* 30: S57-S83.
- Gustman, Alan L. and Thomas L. Steinmeier. 2005. "The Social Security Early Entitlement Age in a Structural Model of Retirement and Wealth." *Journal of Public Economics* 89: 441-463.
- Gustman, Alan L. and Thomas L. Steinmeier and Nahid Tabatabai. 2010. *Pensions in the Health and Retirement Study*. Harvard University Press.
- Hamermesh, D.S. (1985). "Expectations, Life Expectancy, and Economic Behavior." *The Quarterly Journal of Economics*. 100(2): 389-408.
- Heckman, James J. 1996. "Randomization as an Instrumental Variable." *The Review of Economics and Statistics* 78(2): 336-341.
- Heckman, James J. 2001. "Esseys. Econometrics and Empirical Economics." *Journal of Econometrics* 100: 3-5.
- Heckman, James J. 2003. "Commentary. Conditioning, Causality and Policy Analysis." *Journal of Econometrics* 112: 73-78.
- Heckman, James J. 2005. "The Scientific Model of Causality." *Sociological Methodology* 35(1): 1-97.
- Heckman, James J. 2008. "Econometric Causality." *International Statistical Review* 76(1): 1-27.
- Heckman, James J. 2010. "Building Bridges Between Structural and Program Evaluation Approaches to Evaluating Policy." *Journal of Economic Literature* 48(2): 356-398.
- Heckman, James J. and Jeffrey A. Smith. 1995. "Assessing the Case for Social Experiments." *The Journal of Economic Perspectives* 9(2): 85-110.
- Heckman, James J. and Robert J. LaLonde and Jeffrey A. Smith. 1999. "The Economics and Econometrics of Active Labor Market Programs." In *Handbook of Labor Economics* Volume 3, Part A, Chapter 31, pp. 1865-2097.
- Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81(396): 945-960.
- Hotz, Joseph V. and Robert A. Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *The Review of Economic Studies* 60(3): 497-529.

- Jones, Andrew M. and Nigel Rice and Francesca Zantomio. 2016. "Acute Health Shocks and Labour Market Outcomes." University of York, HEDG Working Paper, #04.
- Juster, F.T. (1966). "Consumer Buying Intentions and Purchase Probability: An Experiment in Survey Design." *Journal of the American Statistical Association* 61: 658-696.
- Kalouptsi, Myrto and Paul T. Scott and Eduardo Souza-Rodrigues. 2017. "Identification of Counterfactuals in Dynamic Discrete Choice Models." Working Paper.
- Kapteyn, Arie and Erik Meijer. 2014. "A Comparison of Different Measures of Health and their Relation to Labor Force Transitions at Older Ages." In *NBER Book Discoveries in the Economics of Aging*, Chicago University Press, chapter 3, pp. 115-150.
- Kapteyn, Arie and Arthur van Soest and Julie Zissimopoulos. 2007. "Using Stated Preferences Data to Analyze Preferences for Full and Partial Retirement." IZA Discussion Paper, 2785.
- Keane, Michael and Kenneth I. Wolpin. (2009). "Empirical Applications of Discrete Choice Dynamic Programming Models." *Review of Economic Dynamics* 12: 1-22.
- Lindeboom, Maarten and Marcel Kerkhofs. 2009. "Health and Work of the Elderly Subjective Health Measures, Reporting Errors and the Endogenous Relationship between Health and Work." *Journal of Applied Econometrics* 24(6): 1024-1046.
- Lumsdaine, Robin and Olivia S. Mitchell. 1999. "New Developments in the Economic Analysis of Retirement." In *Handbook of Labor Economics* volume 3, chapter 50, pp. 3261-3307.
- Maestas, Nicole. 2010. "Back to Work: Expectations and Realizations of Work After Retirement." *Journal of Human Resources* 45(3): 718-747.
- Maestas, Nicole and Kathleen J. Mullen and David Powell and Till von Wachter and Jeffrey B. Wenger. 2017. "Working Conditions in the United States: Results of the 2015 American Working Conditions Survey." RAND Research Report.
- Magnac, Thierry and David Thesmar. 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica* 70(2): 801-816.
- Manski, C.F. (1990). "The Use of Intentions Data to Predict Behavior: A Best-Case Analysis." *Journal of the American Statistical Association* 85(412): 934-940.
- Manski, Charles F. 1995. *Identification Problems in the Social Science*. Harvard University Press.
- Manski, Charles F. 1999. "Analysis of Choice Expectations in Incomplete Scenarios." *Journal of Risk and Uncertainty* 19(1-3): 49-66.
- Manski, Charles F. 2003. *Partial Identification of Probability Distributions*. New York: Springer.
- Manski, C.F. (2004). "Measuring Expectations." *Econometrica* 72: 1329-1376.
- Manski, Charles F. 2007. *Identification for Prediction and Decision*. Harvard University Press.

- Manski, C.F. (2017). "Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise." Manuscript prepared for the 2017 NBER Macro Annual Conference.
- Maurer, Jurgen and Roger Klein and Francis Vella. 2011. "Subjective Health Assessments and Active Labor Market Participation of Older Men: Evidence from a Semiparametric Binary Choice Model With Nonadditive Correlated Individual-Specific Effects." *The Review of Economics and Statistics* 93(3): 764–774.
- McGarry, Kathleen. 2004. "Health and Retirement: Do Changes in Health Affect Retirement Expectations?" *Journal of Human Resources* 39(3): 624-648.
- McGeary, Kerry A. 2009. "How Do Health Shocks Influence Retirement Decisions?" *Review of Economics of the Household* 7: 307-321.
- McKenzie, David. 2017. "Can Business Owners Form Accurate Counterfactuals? Eliciting Treatment and Control Beliefs about Their Outcomes in the Alternative Treatment Status." Working Paper, World Bank.
- Pantano, Juan and Yu Zheng. 2013. "Using Subjective Expectations Data to Allow for Unobserved Heterogeneity in Hotz-Miller Estimation Strategies." Working Paper, University of Washington in St. Louis.
- Rohwedder, Susan and Robert J. Willis. 2010. "Mental Retirement." *Journal of Economic Perspectives* 24(1): 119-138.
- Roy, A.D. 1951. "Some Thoughts on the Distribution of Earnings." *Oxford Economic Papers* 3(2): 135-146.
- Rubin, Donald. 1974. Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. *Journal of Educational Psychology* 66(5):688-701.
- Rubin, Donald. 1976. Inference with Missing Data. *Biometrika* 63(3): 581-592.
- Rust, John. 1992. "Do People Behave According to Bellman's Principle of Optimality?" Working Paper, E-92-10, The Hoover Institution, Stanford University.
- Rust, John. 1994. "Structural Estimation of Markov Decision Processes." In *Handbook of Econometrics* Volume IV, Chapter 51, pp. 3081-3143.
- Rust, John and Christopher Phelan. 1997. "How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets." *Econometrica* 65(4): 781-831.
- Schotter, A. and I. Trevino (2014). "Belief Elicitation in the Laboratory." *Annual Review of Economics* 6: 103-128.
- Shao, Feiya. 2016. Subjective Health Shocks and Retirement Timing. Job market paper, University of Michigan.
- van der Klaauw, Wilbert. 2012. "On the Use of Expectations Data in Estimating Structural Dynamic Models." *Journal of Labor Economics* 30(3): 521-554.

van der Klaauw, Wilbert and Kenneth I. Wolpin. 2008. "Social Security and the Retirement and Savings Behavior of Low-Income Households." *Journal of Econometrics* 145: 21-42.

van Soest, Arthur and Hana Vonkova. 2014. "How Sensitive Are Retirement Decisions to Financial Incentives? A Stated Preference Analysis." *Journal of Applied Econometrics* 29(2): 246-264.

Wiswall, Matthew and Basit Zafar. 2016. "Human Capital Investments and Expectations about Career and Family" NBER Working Paper, 22543.

Wolpin, Kenneth I. 1999. "Commentary on 'Analysis of Choice Expectations in Incomplete Scenarios'." *Journal of Risk and Uncertainty* 19(1-3): 67-69.

Figure 1. Treatments and Outcomes on a Simple Health-Work Decision Tree




Figure 2A. 2-Year Ahead Working Expectations by Age, Self-Reported




Figure 2B. 2-Year Ahead Working Expectations by Age, Calculated (LTP)




Figure 2C. 4-Year Ahead Working Expectations by Age, Calculated (LTP)




Figure 3A. 2-Year Ahead Working Expectations in High Health by Age




Figure 3B. 2-Year Ahead Working Expectations in Low Health by Age




Figure 3C. 4-Year Ahead Working Expectations in High Health by Age




Figure 3D. 4-Year Ahead Working Expectations in Low Health by Age






Figure 4. Do Respondents Apply the Law of Total Probability?




Note: Figure shows the distribution of responses for the unconditional probability of working in 2 years computed using the law of total probability (on the vertical axis) versus the self-reported unconditional probability of working in 2 years (on the horizontal axis).

Figure 5. Measured Conditional Value Functions

A. 2-Year Ahead



B. 4-Year Ahead



Note: Figures plots the measured differenced conditional value functions \tilde{v}^H and \tilde{v}^L for each respondent at 2- and 4-year horizons.

Table 1. Sample Selection

Selection Stages	Sample Size
Total sample in Survey 4	3314
Not eligible for the 2 years expectations battery	2249
Career salary reported as 0 USD	9
Not in high health	29
Inconsistent answer to 2 years expectations questions	57
2 Years Sample	970
Not eligible for the 4 years expectations battery	87
Inconsistent answer to 4 years expectations questions	44
4 Years Sample	839

Table 2. Sample Characteristics

	2-Year Ahead	4-Year Ahead
Characteristic	Percent	Percent
Respondent's age (at VRI Survey 4)		
≤ 59	22.9	24.4
60-61	14	14.3
62	6.7	7
63-64	13.4	13.3
65	4.9	4.2
66-67	8.5	8.7
68-69	8.6	8.2
70-71	5.2	5.2
≥ 72	15.9	14.5
Respondent's gender		
Female	37.01	36.83
Male	62.99	63.17
Respondent's race/ethnicity		
Non-Hispanic white	94.74	94.87
Asian	2.68	2.86
Other	2.58	2.26
Respondent's marital status (at VRI S4)		
Partnered (married & share financial future)	65.5	64.8
Not partnered	34.5	35.2
Respondent's educational attainment		
High school or less	5.77	5.96
Some college	14.95	13.83
College graduate	38.97	38.38
Other advanced degree	19.59	20.50
MBA	7.94	8.46
JD, PhD, MD	12.78	12.87
Respondent's health status (at VRI S4)		
High (excellent, very good, or good)	100	100
Respondent's employment status (at VRI S4)		
Working (full-time or part-time)	100	100
Respondent's job type (at VRI S4)		
Career	60.62	61.50
Bridge	39.38	38.50
Respondent's occupation (at VRI S4)		
Management and professional	71.75	71.99
Other services	17.32	17.04
Operative	10.96	10.97
Sample size	970	839

Table 2 (Continued). Sample Characteristics

	2-Year Ahead	4-Year Ahead
Characteristic	Percent	Percent
Total household wealth in USD (at VRI S4)		
First quintile	0 – 258,475	0 – 255,584
Second quintile	258,475 – 533,739	255,584 – 537,700
Third quintile	533,739 – 874,867	537,700 – 877,000
Fourth quintile	874,860 – 1,583,538	877,000 – 1,559,059
Fifth quintile	≥ 1,583,538	≥ 1,559,059
Replacement rate (Expected pension & SS; replacement rate, career job wage, at VRI S4)		
First quintile	0 – 24	0 – 24
Second quintile	24 – 39	24 – 39
Third quintile	39 – 58	39 – 58
Fourth quintile	58 – 87	58 – 88
Fifth quintile	87 +	88 +
Respondent's annual salary in USD (at VRI S4)		
First quintile	0 – 12,000	0 – 13,000
Second quintile	12,000 – 45,714	13,000 – 47,000
Third quintile	45,714 – 77,534	47,000 – 80,000
Fourth quintile	77,534 – 117,000	80,000 – 120,000
Fifth quintile	≥ 117,000	≥ 120,000
Spouse's age (at VRI S4)		
≤ 59	28.55	29.94
60-61	15.11	15.90
62	5.50	5.09
63-64	11.91	12.19
65	3.82	4.01
66-67	7.48	6.94
68-69	9.16	8.95
70-71	4.58	4.48
≥ 72	13.89	12.50
Sample size	635	544
Spouse's health status (at VRI S4)		
Excellent	21.54	21.45
Very good	44.40	44.75
Good	25.81	26.08
Fair	6.87	6.79
Poor	1.37	0.93
Sample size	635	544
Spouse's employment status (at VRI S4)		
Working (full-time or part-time)	48.85	49.07
Not working	51.15	50.93
Sample size	635	544
Sample size	970	839

Table 3. Partition of the Health States

Health fill	Health fill 1	Health fill 2	Health fill 3	Health fill 4	Health fill 5
Self-rated health	(Health PC)	(Health PC)	(Working given health PC)	(Working given health PC)	(Working given health PC)
Excellent	worse	fair or poor	excellent	very good or good	fair or poor
Very good	worse	fair or poor	very good or excellent	good	fair or poor
Good (DK or RF)	fair or poor	very good or excellent	very good or excellent	good	fair or poor
Fair	about the same or worse	very good or excellent	very good or excellent	good	fair or poor
Poor	about the same or worse	very good or excellent	very good or excellent	fair or good	poor

Table 4. 2- and 4-Year Ahead Unconditional Working Expectations

	Percent chance of working in 2 years (self-reported)	Percent chance of working in 2 years (calculated)	Percent chance of working in 4 years (calculated)
Mean	69.8	65.9	52.7
Std. Dev.	35.6	35.3	37
Q01	0	0	0
Q05	0	0	0
Q10	5	4.4	0
Q25	50	40	17
Median	90	80.1	50
Q75	100	97.5	90
Q90	100	100	100
Q95	100	100	100
Q99	100	100	100
N of obs.	970	970	839
% 0	9.48	9.18	11.68
% in (0, 100)	53.09	73.51	76.88
% 100	37.42	17.32	11.44
% 0	9.48	9.18	11.68
% in (0, 50)	12.16	21.86	36
% 50	11.96	2.16	2.5
% in (50, 100)	28.97	49.48	38.38
% 100	37.42	17.32	11.44

Table 5. 2-Year and 4-Year Ahead Health Expectations

	Percent chance of high health in 2 years	Percent chance of low health in 2 years	Percent chance of high health in 4 years	Percent chance of low health in 4 years
Mean	83.4	16.6	76.5	23.5
Std. Dev.	16.5	16.5	19.5	19.5
Q01	25	0	20	0
Q05	50	0	40	0
Q10	50	0	50	5
Q25	75	5	70	10
Median	90	10	80	20
Q75	95	25	90	30
Q90	100	50	95	50
Q95	100	50	100	60
Q99	100	75	100	90
N of obs.	970	970	839	839
% 0	0	13.61	0	6.56
% in (0, 100)	86.39	86.39	93.44	93.44
% 100	13.61	0	6.56	0
% 0	0	13.61	0	6.56
% in (0, 50)	1.75	75.15	6.67	74.37
% 50	9.48	9.48	12.40	12.40
% in (50, 100)	75.15	1.75	74.37	6.67
% 100	13.61	0	6.56	0

Table 6. 2- and 4-Year Ahead Working Expectations in High and in Low Health

	Percent chance of working if in high health in 2 years	Percent chance of working if in low health in 2 years	Percent chance of working if in high health in 4 years	Percent chance of working if in low health in 4 years
Mean	70.5	41.9	58.7	33
Std. Dev.	36	36.1	39	34.4
Q01	0	0	0	0
Q05	0	0	0	0
Q10	5	0	0	0
Q25	50	5	20	0
Median	90	40	68	20
Q75	100	75	100	50
Q90	100	100	100	99
Q95	100	100	100	100
Q99	100	100	100	100
N of obs.	970	970	839	839
% 0	9.38	22.27	12.04	27.77
% in (0, 100)	50.10	64.33	56.73	62.34
% 100	40.52	13.40	31.23	9.89
% 0	9.38	22.27	12.04	27.77
% in (0, 50)	13.81	28.97	25.03	35.76
% 50	9.48	17.73	9.89	13.11
% in (50, 100)	26.80	17.63	21.81	13.47
% 100	40.52	13.40	31.23	9.89

Table 7. 2- and 4-Year Ahead Subjective *ex ante* Treatment Effects (*SeaTE*) of Health on Work

	2-Year <i>SeaTE</i>	4 Year <i>SeaTE</i>	2 Year <i>SeaTE</i> (if <i>SeaTE</i> < 0)	4 Year <i>SeaTE</i> (if <i>SeaTE</i> < 0)
Mean	-28.5	-25.7	-40.9	-36.8
Std. Dev.	27.9	27.6	24.1	25.1
Q01	-100	-100	-100	-100
Q05	-80	-80	-90	-85
Q10	-70	-70	-75	-75
Q25	-50	-50	-50	-50
Median	-25	-20	-40	-30
Q75	0	0	-20	-15
Q90	0	0	-10	-10
Q95	0	0	-9.4	-5
Q99	5	0	-4.3	-2
N of obs.	970	839	682	594
<hr/>				
% STE = -100	1.55	1.43		
% STE in (-100, 0)	68.76	69.37		
% STE = 0	28.45	28.25		
% STE > 0	1.24	0.95		
<hr/>				
% STE = -100	1.55	1.43		
% STE in (-100, -50)	14.74	12.75		
% STE = 50	14.43	11.32		
% STE in (-50, 0)	39.59	45.29		
% STE = 0	28.45	28.25		
% STE > 0	1.24	0.95		

Note: *SeaTE* = probability of working if in low health – probability of working if in high health.

Table 8. 2- and 4-Year Ahead Average Subjective *ex ante* Treatment Effects

Horizon	2 Years	4 Years
Aggregate TE parameter		
ASeaTE	-28.5	-25.7
ASeaTT	-27.4	-25.2
ASeaTU	-28.8	-25.8
Sample Size	970	839

Note: ASeaTE is the average Subjective *ex ante* Treatment Effect; ASeaTT is the average Subjective *ex ante* Treatment Effect on the Treated; ASeaTU is the average Subjective *ex ante* Treatment Effect on the Untreated. All three are defined in equation (1.8).

Table 9. Predictors of 2- and 4-Year Ahead *SeaTE*

Predictors	2-Year Ahead	4-Year Ahead
	<i>SeaTE</i>	<i>SeaTE</i>
	Coeff (SE)	Coeff (SE)
Constant	-0.150** (0.061)	-0.116* (0.065)
R's age		
age in 60-61	-0.046 (0.031)	-0.038 (0.032)
age = 62	-0.117*** (0.040)	-0.055 (0.041)
age in 63-64	-0.034 (0.031)	-0.034 (0.033)
age = 65	-0.031 (0.045)	-0.111** (0.051)
age in 66-67	-0.021 (0.037)	0.028 (0.039)
age in 68-69	-0.120*** (0.037)	-0.081** (0.040)
age in 70-71	-0.116** (0.046)	-0.088* (0.049)
age ≥ 72	-0.088** (0.034)	-0.086** (0.037)
R's gender		
female	0.001 (0.021)	-0.012 (0.023)
R's education		
some college	-0.002 (0.044)	-0.025 (0.047)
college grad	0.006 (0.042)	-0.010 (0.044)
other adv. degree	-0.042 (0.045)	-0.019 (0.047)
MBA	-0.014 (0.051)	0.003 (0.054)
JD, PhD, MD	-0.031 (0.050)	-0.076 (0.053)
R's occupation		
operative	0.008 (0.025)	-0.008 (0.027)
other services	-0.020 (0.032)	-0.020 (0.034)
R's job type		
bridge	0.008 (0.022)	-0.015 (0.023)
R's marital status		
partnered	-0.012 (0.024)	-0.010 (0.026)
Spouse's work status		
working	-0.014 (0.023)	-0.003 (0.025)

Total HH wealth		
1 st quintile	-0.039 (0.033)	-0.039 (0.036)
2 nd quintile	-0.045 (0.032)	-0.084** (0.034)
3 rd quintile	-0.020 (0.030)	-0.032 (0.032)
4 th quintile	-0.044 (0.029)	-0.044 (0.031)
R's replacement rate		
1 st quintile	-0.013 (0.031)	-0.022 (0.033)
2 nd quintile	0.002 (0.031)	0.028 (0.033)
3 rd quintile	-0.023 (0.030)	-0.023 (0.032)
4 th quintile	-0.018 (0.030)	-0.018 (0.032)
R's current salary		
1 st quintile	-0.039 (0.037)	-0.032 (0.039)
2 nd quintile	-0.067** (0.034)	-0.040 (0.036)
3 rd quintile	-0.0001 (0.032)	0.008 (0.034)
4 th quintile	-0.005 (0.030)	-0.003 (0.031)
Sample size	970	839
R^2	0.0484	0.0528

Note: *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 10. DP Values in High and Low Health

	2-year Ahead		4-year Ahead	
	\tilde{v}^H	\tilde{v}^L	\tilde{v}^H	\tilde{v}^L
Mean	0.97	-0.35	0.48	-0.72
Standard deviation	1.70	1.65	1.80	1.60
Q01	-2.58	-2.58	-2.58	-2.58
Q10	-1.64	-2.58	-2.58	-2.58
Q25	0	-1.64	-0.84	-2.58
Median	1.28	-0.25	0.47	-0.84
Q75	2.58	0.67	2.58	0
Q90	2.58	2.58	2.58	2.33
Q99	2.58	2.58	2.58	2.58
% > 0	67.32	31.03	53.04	23.36
Sample Size	970	970	839	839

Notes: Table shows distribution of the measured differenced conditional value functions \tilde{v}^H and \tilde{v}^L .

Table 11. Quantifying Cross-Sectional Heterogeneity in DP Values

Horizon	2-year Ahead		4-year Ahead	
	$h=H$	$h=L$	$h=H$	$h=L$
Health state				
α^h	0.97*** (0.05)	-1.04*** (0.04)	0.48*** (0.06)	-1.04*** (0.04)
γ		0.71*** (0.02)		0.66*** (0.02)
$\sigma(v^h)$	1.70	1.12	1.80	1.06
Sample Size	970	970	839	839

Notes: Table shows mean, covariance, and variability of the measured differenced conditional value functions \tilde{v}^H and \tilde{v}^L as specified in equation (1.28) of text. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 12. Quantifying Cross-Sectional Heterogeneity in DP Values, With Covariates

Horizon	2-year Ahead		4-year Ahead	
Health state	<i>h=H</i>	<i>h=L</i>	<i>h=H</i>	<i>h=L</i>
Constant	1.34*** (0.35)	-0.39 (0.24)	0.43 (0.40)	-0.44* (0.24)
R's age				
age in 60-61	-0.41** (0.17)	-0.34*** (0.12)	-0.47** (0.20)	-0.32*** (0.12)
age = 62	-0.18 (0.22)	-0.51*** (0.16)	-0.99*** (0.25)	-0.48*** (0.16)
age in 63-64	-1.08*** (0.18)	-0.46*** (0.13)	-1.23*** (0.20)	-0.61*** (0.13)
age = 65	-1.43*** (0.25)	-0.57*** (0.18)	-1.01*** (0.31)	-0.79*** (0.19)
age in 66-67	-1.07*** (0.21)	-0.34** (0.15)	-1.54*** (0.24)	-0.41*** (0.15)
age in 68-69	-1.22*** (0.21)	-0.86*** (0.15)	-1.36*** (0.24)	-0.77*** (0.15)
age in 70-71	-1.15*** (0.26)	0.72*** (0.18)	-1.16*** (0.30)	-0.74*** (0.18)
age ≥ 72	-1.03*** (0.19)	-0.71*** (0.14)	-1.01*** (0.23)	-0.72*** (0.14)
R's gender				
female	-0.09 (0.12)	-0.04 (0.08)	0.03 (0.14)	-0.08 (0.08)
R's education				
some college	-0.16 (0.25)	-0.16 (0.17)	0.13 (0.29)	-0.08 (0.18)
college grad	-0.08 (0.24)	-0.03 (0.16)	0.07 (0.27)	-0.01 (0.16)
other adv. degree	0.15 (0.25)	-0.13 (0.18)	0.15 (0.29)	0.03 (0.18)
MBA	0.13 (0.29)	-0.06 (0.20)	0.16 (0.33)	0.01 (0.20)
JD, PhD, MD	0.18 (0.28)	-0.09 (0.20)	0.55 (0.32)	-0.19 (0.20)
R's occupation				
operative	0.04 (0.14)	0.06 (0.10)	0.04 (0.17)	-0.01 (0.10)
other services	0.11 (0.18)	0.05 (0.13)	0.12 (0.21)	-0.00 (0.13)
R's job type				
bridge	-0.05 (0.12)	0.01 (0.09)	0.03 (0.14)	-0.08 (0.09)
R's marital status				
partnered	-0.35*** (0.13)	-0.14 (0.09)	-0.16 (0.16)	-0.16* (0.10)

Spouse's work status	0.36***	0.03	0.02	0.04
working	(0.13)	(0.09)	(0.15)	(0.09)
Total HH wealth				
1 st quintile	0.92***	0.19	1.02***	0.16
	(0.19)	(0.13)	(0.22)	(0.14)
2 nd quintile	0.68***	0.01	1.01***	-0.09
	(0.18)	(0.12)	(0.21)	(0.13)
3 rd quintile	0.40**	0.02	0.39**	-0.04
	(0.17)	(0.12)	(0.20)	(0.12)
4 th quintile	0.24	-0.02	0.28	-0.10
	(0.16)	(0.12)	(0.19)	(0.12)
R's replacement rate				
1 st quintile	0.24	0.10	0.47**	0.14
	(0.18)	(0.12)	(0.20)	(0.12)
2 nd quintile	0.29*	0.09	0.09	0.22*
	(0.17)	(0.12)	(0.20)	(0.12)
3 rd quintile	0.19	0.008	0.09	-0.00
	(0.17)	(0.12)	(0.20)	(0.12)
4 th quintile	-0.05	-0.14	-0.11	-0.01
	(0.17)	(0.12)	(0.20)	(0.12)
R's current salary				
1 st quintile	-0.49**	-0.20	-0.07	-0.11
	(0.21)	(0.14)	(0.24)	(0.15)
2 nd quintile	0.22	-0.11	0.33	0.00
	(0.19)	(0.13)	(0.22)	(0.13)
3 rd quintile	-0.06	0.02	0.09	0.09
	(0.18)	(0.12)	(0.21)	(0.12)
4 th quintile	-0.19	0.01	-0.12	-0.04
	(0.17)	(0.12)	(0.19)	(0.12)
γ		0.64***		0.61***
		(0.02)		(0.02)
$\sigma(\nu^h)$	1.56	1.09	1.67	1.02
Sample Size	970	970	839	839
R^2	0.19	0.58	0.16	0.61

Notes: Table shows mean as a linear index of covariates, covariance, and variability of the measured differenced conditional value functions $\tilde{\nu}^H$ and $\tilde{\nu}^L$ as specified in equation (1.28) of text. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 13. Relationship between Health and Work with Simulated Realizations

Horizon	2-year Ahead			4-year Ahead		
Case	Uncorrelated	Empirical	Higher correlation	Uncorrelated	Empirical	Higher correlation
Constant	0.703 (0.011)	0.709 (0.011)	0.730 (0.011)	0.586 (0.019)	0.594 (0.014)	0.621 (0.018)
Health h	-0.282 (0.040)	-0.305 (0.039)	-0.415 (0.039)	-0.254 (0.040)	-0.286 (0.035)	-0.371 (0.039)
SEE	0.463	0.461	0.448	0.488	0.485	0.474
Sample size	970	970	970	839	839	839

Note: Table reports mean values from the 1000 replications. Uncorrelated case has health transition probability fixed at mean; empirical case uses individual-specific elicited health transition probabilities. Highly correlated case has stronger correlation of health transition probability and value of work as described in text. The LHS variable is the simulated realized decision to work (d). The RHS variable is simulated realized health state (h).