

Measuring Time-inconsistency: Evidence from Work-Welfare Decisions in NLSY*

Hanming Fang[†]

Dan Silverman[‡]

March 27, 2001

Abstract

This paper applies a model of potentially time-inconsistent preferences to the problem of dynamic labor supply and welfare program participation. From panel data on the choices of single women with children, we provide estimates of the degree of time-inconsistency. With these estimates we can quantify the utility loss stemming from the inability to commit to future decisions, and the potential value of commitment mechanisms such as welfare time limits.

*This paper is preliminary and incomplete. All comments are welcome. We thank Steve Berry, Stephano Della Vigna, Donghoon Lee, Daniele Passerman, Andrew Postlewaite, Matthew Rabin, John Rust, Kent Smetters, Ken Wolpin and seminar participants at Yale for many helpful discussions. We also thank Ken Wolpin for supplying data on U.S. state welfare benefits. Silverman gratefully acknowledges support from the SSRC Fellowship in Applied Economics.

[†]Department of Economics, Yale University, 37 Hillhouse Avenue, New Haven, CT 06511. Email: hanming.fang@yale.edu.

[‡]Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104. Email: silvermd@econ.sas.upenn.edu.

1 Motivation

Economists studying choice over time typically assume decision makers are impatient. Traditionally, this impatience is modelled in a very particular way. It is assumed that agents discount future streams of utility (or profits) *exponentially* over time. Strotz (1955) showed that exponential discounting is not just an analytically convenient assumption; rather, it represents the only rational form of impatience. Without this assumption, Strotz demonstrated, intertemporal marginal rates of substitution will change as time passes, and preferences will be time-inconsistent.

Recently, a growing theoretical literature has built on the work of Strotz, and others, to explore the consequences of relaxing the standard assumption of time-consistent discounting.¹ Drawing on experimental research in psychology, economists have built models of (quasi-) hyperbolic discounting to reflect the apparently fundamental tendency of decision makers to seize short term rewards at the expense of long-term preferences.² With (quasi-) hyperbolic discounting, the relative value of utility received in period t versus period $t + 1$ increases as period t draws nearer. Agents, that is, have a bias towards near-term utility and immediate gratification. As a result intrapersonal conflicts arise when earlier ‘selves’ prefer a future sequence of trade-offs that their later selves will not find optimal and therefore will not make. Thus, to the extent that they recognize this interpersonal conflict, earlier selves have an incentive to commit their future selves to a preferred sequence. In other words, time-inconsistent preferences may generate an incentive to exercise self-control that is absent from the standard framework.³

This interpersonal conflict, and the incentives to set limits on future behavior have implications for a variety of economic problems. Models of time-inconsistent preferences have been applied to topics ranging from consumption and savings (Laibson, 1997; O’Donoghue and Rabin, 1999a,b), to growth (Barro, 1999), to technology adoption (Jovanovic and Stolyarov, 1999), to belief formation (Carillo and Mariotti, 2000), to job search (Della Vigna and Paserman, 2000). In some of these applications, researchers have used time-inconsistent preferences to parsimoniously explain a set of common but seemingly irrational behaviors. Most often, models of time-inconsistent preferences have been used to show why agents might benefit from, and perhaps pay for, ‘commitment devices’: instruments for restricting their own future choices.

¹See, for example, Barro (1999), Corriolo and Mariotti (2000), Jovanovic and Stolyarov (2000), Laibson (1997), O’Donoghue and Rabin (1999a,b).

²A large body of experimental research in psychology suggests that hyperbolic time discounting is fundamental to the decision making of agents both animal and human. See Ainslie (1992) for a review.

³For thorough presentations of economic problems of self control see Ainslie (1992) O’Donoghue and Rabin (1999b), and Thaler (1991).

While (quasi) hyperbolic discounting has important theoretical implications and accurately describes experimental data on choice over time, models of time-inconsistent preferences have been criticized for adding analytic complexity without explaining much more about actual economic decisions as reflected in survey data.⁴ For example, lifecycle consumption-saving decisions have been studied closely with models of time-inconsistency. Researchers such as Laibson, et al. (1998) find support for time-inconsistency in preferences for saving from evidence that: individuals undersave with respect to their stated preferences; consumers simultaneously maintain positive savings and expensive debt; and individuals commonly use commitment devices such as early withdrawal penalties and holiday clubs. Others have criticized this evidence as having several explanations other than time-inconsistency, or as demonstrating the limited importance of an inability to commit to future choices. For instance, survey data show there is relatively little reliance on commitment mechanisms for saving.⁵ If self-control problems were critical for understanding actual choices over time, it is argued, commitment devices should be in much greater demand.

This paper contributes to the literature on time-inconsistent preferences in two ways. First, by applying a model of time-inconsistent preferences to the problem of labor supply and welfare program participation, we provide an economically significant setting for an empirical evaluation of the importance of time-inconsistency.⁶ As a source of information about the economic importance of time-inconsistency, this context has the advantage that labor supply decisions are among the most consequential economic choices that individuals make. These decisions largely define time use for working-age adults. One might expect, therefore, that choices about whether, and how much, to work would be less subject to prolonged lapses in strict intertemporal rationality. Nevertheless, recent welfare reforms and anecdotal evidence indicate a commonly held view: the trade-off between the short-term costs of entering the labor force at a low wage relative to the government benefit, and long term reward of higher wages from the accumulation of work experience, generates important problems of self-control. Our previous research has shown that this trade-off can in theory produce important observable differences in the behavior of time-consistent and time-inconsistent agents (Fang and Silverman, 2000).

Our second contribution to the literature on time-inconsistency is methodological. Systematic empirical study of time-inconsistency began with the work of experimental psychologists. See

⁴See, for example, William Gale's comments on Laibson, et al. (1998) published with the article, and Mulligan (1997).

⁵Gale estimates that the value of assets in holiday clubs represents just 2 percent of annual savings and just 0.02 percent of all financial assets.

⁶In a related labor economics paper, DellaVigna and Paserman (2000) consider the influence of self control problems on job search.

Ainslie (1992) for a review. More recently, economists have calibrated models of time-inconsistent preferences to match important moments of aggregate data. See, e.g., Laibson, et al. (1998). In this paper, we attempt to directly estimate the structural parameters of the model from a single data set. With simulations based on these estimated parameters, our work can provide quantitative estimates of both the degree of time-inconsistency and the extent of its consequences.

Last, this paper contributes to the literature on welfare program participation and to the welfare reform debate with the first empirical examination of the importance of time-inconsistency for the welfare take-up decision. Our methodology can generate quantitative estimates of the influence of time-inconsistency through, for example, calculations of the value of commitment to potential welfare recipients, or estimations of the predicted response to welfare reforms such as time-limits and work requirements.

This paper joins a small empirical literature on work and welfare decisions that structurally implemented models that incorporate dynamic decision-making. Sanders and Miller (1997) estimate a dynamic model of discrete choice in which women sequentially make monthly decisions about whether or not to work and whether or not to receive welfare. It incorporates the features of wage growth through experience on the job, and preferences which adapt to labor supply and welfare experience, to explain non-participation in AFDC among eligible families and the observed negative duration dependence of families who enroll. As in this paper fertility and marriage are taken as exogenous. Swann (1996) adds marriage to the choice set, and looks at the women's decisions annually. In a recent paper, Keane and Wolpin (2000) endogenize education, employment, fertility and marriage decisions of women. All of these papers take as given a standard exponential discounting preferences, and in fact, Sanders and Miller (1997) had difficulty in pinning down the discount factor. Our paper complements these three studies in focusing on examining the possibility of present biased preference among potential welfare participants.

The remainder of the paper is structured as follows: Section 2 presents our model, describes the intrapersonal game played by the decision maker, and defines the game's solution concept. Section 3 describes the numerical method for obtaining the game's solution. Section 4 presents the estimation strategy. Section 5 describes the data and variable definitions. Section 6 provides preliminary estimation results and associated simulations, and discusses their limitations. Section 7 offers preliminary conclusions and describes our plans for further work.

2 The Model

Borrowing from Keane and Wolpin (1994, 1997), we consider a discrete time model of a parent’s (agent’s) work-welfare decisions. Each agent has a finite decision horizon starting from her age at the birth of her first child, a_0 , and ending at age A .⁷ At each age $a \in \{a_0, \dots, A\}$, the agent must choose from among three mutually exclusive and exhaustive alternatives: receive welfare, work in the labor market, or stay at home without work or welfare.⁸ The agent’s decision in period a is denoted by $d_a \in D = \{0, 1, 2\}$, corresponding to welfare, work, or home.

The return from each of the alternatives, which represents all of the current-period benefits and costs associated with the choice, is denoted by $R_a(d)$ and defined as follows:

1. *Welfare.* The current-period return to welfare in period a , $R_a(0)$, depends on the number of the agent’s children n_a ; her municipality j and its cash and food welfare benefits as a function of the number of her children $G_{ja}(n_a)$; her home production skills (alternatively her value of leisure), e_a ; the cumulative number of periods she has received welfare κ ; the net stigma associated with welfare $\phi(d_{a-1})$ denominated in dollars; and an idiosyncratic, choice-specific shock ε_{0a} .⁹ Thus, in the absence of a time limit the return to welfare is given by:

$$R_a(0) = e_a + G_{ja}(n_a) - \phi(d_{a-1}) + \varepsilon_{0a}$$

where

$$\phi(d_{a-1}) = \begin{cases} 0, & \text{if } d_{a-1} = 0 \\ \phi, & \text{otherwise.} \end{cases}$$

In words, we assume stigma lasts for only one period after switching into welfare from some other choice.¹⁰ More generally, given a lifetime time limit of L periods, the return to welfare

⁷In the estimation and simulations we set A to the highest age observed in our data set, 34 years old.

⁸Edin and Lein (1997) report that these three categories are not, in fact, mutually exclusive. According to their study of 379 low-income single mothers, many welfare recipients both work in the (unofficial) labor market and rely on family and neighborhood resources. Our analysis abstracts from the overlap of these categories, and from the distinctions between part- and full-time work.

⁹In the U.S., the welfare benefit may also depend on the assets and income of the agent. We abstract from these additional eligibility restrictions. The net stigma parameter $\phi(d_{a-1})$ has, since Moffitt (1983), become standard in empirical studies of welfare participation. Its primary function is to help explain the significant number of welfare-eligible adults who remain at home without work or welfare. It is *net* stigma because it captures both the psychic and administrative costs associated with welfare takeup and the in-kind consumption benefits of welfare such housing subsidies and health insurance. It is made a function of the previous period decision to allow for the decay of stigma with continued participation.

¹⁰The one shot formulation of welfare stigma has two important consequences: first, it makes staying on welfare more attractive if one is already on welfare last period, that is, it makes the welfare participation more persistent;

is given by:

$$R_a(0) = \begin{cases} e_a + G_{ja}(n_a) - \phi(d_{a-1}) + \varepsilon_{0a}, & \text{if } \kappa < L \\ e_a + \varepsilon_{0a}, & \text{otherwise.} \end{cases}$$

The welfare benefits schedule is assumed to be an affine function of the number of the agent's children and is summarized by the parameter $\theta_{ja} = [\theta_{ja0}, \theta_{ja1}]$ where

$$G_{ja}(n_a) = \theta_{ja0} + \theta_{ja1}n_a .$$

We assume, moreover, that this benefits schedule remains fixed, in real terms, across time and so:

$$G_{ja}(n_a) = G_j(n_a) = \theta_{j0} + \theta_{j1}n_a .$$

We assume residential location at age a_0 is given exogenously and remains unchanged through age A .¹¹ We treat the arrival of additional children as exogenously determined and model births as a stochastic process whereby:

$$n_{a+1} = \begin{cases} n_a + 1, & \text{with probability } \rho(a, n_a, d_a) \\ n_a, & \text{with probability } (1 - \rho(a, n_a, d_a)) \end{cases} .$$

Home production skills/value of leisure is allowed to vary by age, and the number of children such that:

$$e_a = e_0 + e_1\mathbf{I}(14 \leq a \leq 18) + e_2a + e_3n_a + e_4n_a^2$$

where $\mathbf{I}(\cdot)$ is an indicator function equal to one if the expression in parentheses is true, and zero otherwise.^{12,13}

2. *Work.* The current-period return from work at any age, $R_a(1)$, is the wage. Following a standard theory of human capital, we model this wage as the product of a (constant) rental price of human capital r times the number of skill units held by the individual h_a .

$$R_a(1) = rh_a .$$

The number of skill units h_a is a function of the agent's completed years of schooling g_{a0} , and her work experience x_a . We assume the agent's level of schooling g_{a0} remains unchanged after

second, when people are present-biased, then it may keep more people off welfare because they want to avoid/postpone the immediate cost of stigma. These two effects are clearly illustrated in the Figures in Section 6.2.

¹¹CAN WE DETERMINE FOR WHAT FRACTION OF OUR SUBSAMPLE THIS IS TRUE?

¹²The return to leisure is allowed to depend on the age and the number of children to proxy for the effect on the value of leisure of living at home with parents, and of having (more) younger children. In an effort to ease the computational burden, these effects are not modelled explicitly.

¹³We include the term $e_1\mathbf{I}(14 \leq a \leq 18)$ to "explain" the sharp decline in home participation in that age interval. This age effect can arise, for example, if it is less costly to live with parents when young.

the birth of her first child, and so do not model the schooling choice.¹⁴ If her skill endowment at the birth of her first child is denoted by h_{a_0} , then at age $a \geq a_0$ the agent's skills are given by:

$$h_a = \exp [h_{a_0} + \alpha_1 g_{a_0} + \alpha_2 x_a - \alpha_3 x_a^2 - \alpha_4 (1 - \mathbf{I}(d_{a-1} = 1)) + \varepsilon_{1a}]$$

where ε_{1a} is a skill technology shock. Note that the human capital accumulated through work experience is allowed to decay. In particular, the parameter α_4 represents the one-time depreciation of human capital that occurs whenever the agent leaves work for welfare or home. All together, this functional form implies a standard $\ln(\text{wage})$ equation where the constant term consists of $\ln(r) + h_a$.

3. *Home.* The third alternative is to stay home without work or welfare. We assume that by staying home an agent can generate the same return from home production/leisure, e_a , as she could if she were to choose welfare. The total return from staying home is therefore $R_a(2) = e_a + \varepsilon_{2a}$ where ε_{2a} is again a choice specific shock.

The choice-specific shocks $\boldsymbol{\epsilon}_a = (\varepsilon_{0a}, \varepsilon_{1a}, \varepsilon_{2a})$ are assumed to be distributed joint normal, $N(0, \boldsymbol{\Omega})$, and serially uncorrelated. Initial conditions are given by the agent's municipality j , her level of schooling and experience at the birth of her first child, g_{a_0} and x_{a_0} , her skill endowments h_{a_0} , and her decision just prior to the birth of her first child d_{a_0-1} . We denote the state at age a by $\mathbf{s}_a = [j, g_{a_0}, h_{a_0}, d_{a-1}, x_a, n_a, \kappa_a, \boldsymbol{\epsilon}_a]$, and the set of all possible period- a states by \mathbf{S}_a . It will also be useful to identify the predetermined components of the state $\bar{\mathbf{s}}_a = [j, g_{a_0}, h_{a_0}, d_{a-1}, x_a, n_{a-1}, \kappa_a]$. The current-period return to choice $d \in D$ in state \mathbf{s}_a is denoted by $R(d|\mathbf{s}_a)$.

Preferences. We assume the agent consumes all of the returns associated with her choice in each period, and obtains an instantaneous utility $u_a = R_a$.¹⁵ An agent in period a is concerned about both her present and her future instantaneous utilities. Let $U^a(u_a, u_{a+1}, \dots, u_A)$ represent an agent's intertemporal preferences from the perspective of period a . We adopt a simple, and now commonly used, formulation of agents' potentially time-inconsistent preferences, (β, δ) -preferences (Phelps and Pollak 1968):

Definition 1 (β, δ) -preferences are preferences that can be represented by:

$$\text{for all } a, U^a(u_a, \dots, u_A) \equiv \delta^a u_a + \beta \sum_{t=a+1}^A \delta^t u_t, \text{ where } \beta \in (0, 1], \delta \in (0, 1]. \quad (1)$$

¹⁴In fact, 34% of our sample goes on to acquire additional schooling after the birth of their first child. Of this fraction, approximately half goes on to acquire more than one additional year of schooling.

¹⁵The incentives to save may be important but we leave this analysis for future research.

The parameter δ represents long-run, time-consistent discounting and is called the *discount factor*. The parameter β represents short-term impatience, and is called the *present bias factor*. Note that when $\beta = 1$, (β, δ) -preferences are time-consistent preferences,

$$U^a(u_a, \dots, u_A) \equiv \sum_{t=a}^A \delta^t u_t .$$

When $\beta \in (0, 1)$, (β, δ) -preferences are ‘quasi-hyperbolic’ discounting preferences (Laibson, 1997). We will call an agent with $\beta = 1$, time-consistent (TC), and an agent with $\beta \in (0, 1)$ present-biased.

Following previous studies of time-inconsistent preferences, we will analyze the behavior of an agent in this model by thinking of the individual as consisting of a separate *self* in each period. Each period- a self chooses her current behavior to maximize her current preferences $U^a(u_a, \dots, u_A)$, while her future selves control her subsequent decisions. The literature on time-inconsistent preferences distinguishes between *sophisticated* and (*partially*) *naive* agents (Strotz 1955, Pollak 1968, O’Donoghue and Rabin 1999a,b). An agent is *sophisticated* if in every a , the period- a self knows her future selves’ preferences and correctly anticipates their behavior when making her period- a decision. Thus her anticipated present bias factor, call it $\tilde{\beta}$, equals the true β . She is *partially naive* if at every a , the period- a self underestimates the present-bias of her future selves, i.e. $\tilde{\beta} \in (\beta, 1)$. She is perfectly *naive* if she believes that her future selves are time consistent, i.e. $\tilde{\beta} = 1$. In what follows, we will consider the behavior of both sophisticated and (partially) naive agents.

2.1 Strategies and Payoffs

When a period- a self makes a decision her *choice space* is given by $D \equiv \{0, 1, 2\}$ where, again, 0 represents welfare, 1 work, and 2 home with no welfare. We restrict our attention to Markov strategies and define a *feasible strategy* for a period- a self to be a mapping $\sigma_a : \mathbf{S}_a \rightarrow D$. Because the returns from choosing an alternative in period a represent all the costs and benefits of that alternative, the instantaneous utility $u_a(\mathbf{s}_a, \sigma_a)$ the agent obtains from strategy σ_a in state \mathbf{s}_a can be written as:

$$u_a(\mathbf{s}_a, \sigma_a) = \frac{(1 - \sigma_a(\mathbf{s}_a))(2 - \sigma_a(\mathbf{s}_a))}{2} R(0|\mathbf{s}_a) + \sigma_a(\mathbf{s}_a)(2 - \sigma_a(\mathbf{s}_a)) R(1|\mathbf{s}_a) + \frac{\sigma_a(\mathbf{s}_a)(\sigma_a(\mathbf{s}_a) - 1)}{2} R(2|\mathbf{s}_a). \quad (2)$$

Given \mathbf{s}_a and σ_a , the predetermined elements of state inherited by the period- $(a + 1)$ self are denoted by $\bar{\mathbf{s}}_{a+1}(\mathbf{s}_a, \sigma_a)$ and given by:

$$\bar{\mathbf{s}}_{a+1}(\mathbf{s}_a, \sigma_a) = \left(j, g_{a0}, h_{a0}, e_a, \sigma_a(\mathbf{s}_a), (x_a + \sigma_a(\mathbf{s}_a)(2 - \sigma_a(\mathbf{s}_a))), \kappa + \frac{(1 - \sigma_a(\mathbf{s}_a))(2 - \sigma_a(\mathbf{s}_a))}{2} \right).$$

A *strategy profile* for all selves is $\sigma \equiv (\sigma_a)_{a=a_0}^A$, and specifies for each self her action in all possible states. Given a strategy profile σ , we will write $EV_a(\bar{\mathbf{s}}_a, \sigma)$ as the agent's period- a *long-run* expected utility as a function of the predetermined elements of her current state $\bar{\mathbf{s}}_a$ and the strategy profile σ . The expectation is with respect to the sequence of remaining shocks $\{\epsilon_t\}_{t=a}^A$, and the fertility function $\rho(a, n, d_a)$. An agent's long run expected utility represents her intertemporal preferences from some prior perspective, where both the realization of her current and future shocks are unknown and her own present bias is irrelevant. $EV_a(\bar{\mathbf{s}}_a, \sigma)$ can be written recursively as:

$$EV_A(\bar{\mathbf{s}}_A, \sigma) = Eu_A(\bar{\mathbf{s}}_A, \sigma_A), \text{ and} \quad (3)$$

$$\text{for } a = A - 1, \dots, a_0, \quad EV_a(\bar{\mathbf{s}}_a, \sigma) = E \{u_a(\bar{\mathbf{s}}_a, \sigma_a) + \delta V_{a+1}(\bar{\mathbf{s}}_{a+1}(\bar{\mathbf{s}}_a, \sigma_a), \sigma)\}. \quad (4)$$

To introduce our solution concept due to O'Donoghue and Rabin (1999b), we define a *perception-perfect strategy profile* for TCs as:

Definition 2 *A perception-perfect strategy profile for TCs is a strategy profile σ^{TC} that satisfies for all \mathbf{s}_a , and for all a ,*

$$\sigma_a^{TC}(\mathbf{s}_a) = \arg \max_{\sigma \in A} \{u_a(\mathbf{s}_a, \sigma) + \delta EV_{a+1}(\bar{\mathbf{s}}_{a+1}(\mathbf{s}_a, \sigma), \sigma^{TC})\}.$$

The perception-perfect strategy profile for a TC is simply the optimal plan from her period- a_0 self's point of view. Since her preferences are time-consistent, a TC's subsequent selves will follow this period- a optimal plan.

If the agent is partially naive, then any period- a self believes that, beginning next period, her future selves will behave optimally, given a present bias factor of $\tilde{\beta} \in (\beta, 1)$. I.E., future selves are anticipated to play the continuation strategy profile $\tilde{\sigma}$ given by:

Definition 3 *The perceived continuation strategy profile for (partial) naifs is the strategy profile $\tilde{\sigma}$ that satisfies for all $a = a + 1, \dots, A$, all \mathbf{s}_a ,*

$$\tilde{\sigma}_a(\mathbf{s}_a) = \arg \max_{\sigma \in D} \{u_a(\mathbf{s}_a, \sigma) + \tilde{\beta} \delta EV_{a+1}(\mathbf{s}_{a+1}(\mathbf{s}_a, \sigma), \tilde{\sigma})\}.$$

Anticipating the continuation strategy profile $\tilde{\sigma}$, the perception-perfect strategy profile for partial naifs is given by:

Definition 4 *A perception-perfect strategy profile for partial naifs is a strategy profile σ^{PN} that satisfies for all a , all \mathbf{s}_a ,*

$$\sigma_a^{PN}(\mathbf{s}_a) = \arg \max_{\sigma \in D} \{u_a(\mathbf{s}_a, \sigma) + \beta \delta EV_{a+1}(\mathbf{s}_{a+1}(\mathbf{s}_a, \sigma), \tilde{\sigma})\}.$$

Note that when $\tilde{\beta} = \beta$, the agent is ‘sophisticated.’ The period- a self of a sophisticated agent realizes that her future selves will have different preferences over the same actions and, like a TC, correctly predicts how her future selves will behave given her own action. Unlike TCs, however, the period- a self of a sophisticated, present-biased agent will discount the expected continuation values of her current action by $\beta\delta$. When $\tilde{\beta} \in (\beta, 1)$ the agent is partially naive and underestimates the present-bias that determines the behavior of her future selves. When, $\tilde{\beta} = 1$, the agent is perfectly naive, and her period- a self incorrectly believes that her future selves share her preferences over future trade-offs.

3 Solving σ^s Recursively

When $A < \infty$, the perception-perfect strategy profile for partially naive agents can be solved recursively. The solutions for sophisticated, perfectly naive, and TC agents are merely special cases of the partially naive solution, so we consider only the partial naif’s problem.

Suppose an agent enters the terminal period A with a deterministic state $\bar{\mathbf{s}}_A$. In period A the individual draws random shocks from the joint distribution over ϵ and realizes n_A , to complete the state \mathbf{s}_A . Then for all \mathbf{s}_A , $\sigma_A^{PN}(\mathbf{s}_A)$ is simply given by:

$$\sigma_A^{PN}(\mathbf{s}_A) = \arg \max_{\sigma \in D} u_A(s_A, \sigma) = \tilde{\sigma}_A(\mathbf{s}_A);$$

Now define

$$EV_A(\bar{\mathbf{s}}_A) = Eu_A(\bar{\mathbf{s}}_A, \tilde{\sigma}_A(\mathbf{s}_A)).$$

Thus, by way of recursion, for every $a = A - 1, \dots, a_0$, for every \mathbf{s}_a :

$$\begin{aligned} \tilde{\sigma}_a(\mathbf{s}_a) &= \arg \max_{\sigma \in D} \{u_a(\mathbf{s}_a, \sigma) \\ &\quad + \tilde{\beta}\delta EV_{a+1}(\bar{\mathbf{s}}_{a+1}(\mathbf{s}_a, \sigma))\}, \\ \sigma_a^{PN}(\mathbf{s}_a) &= \arg \max_{\sigma \in D} u_a(\mathbf{s}_a, \sigma) + \beta\delta EV_{a+1}(\bar{\mathbf{s}}_{a+1}(\mathbf{s}_a, \sigma)) \end{aligned}$$

and

$$EV_a(\bar{\mathbf{s}}_a) = Eu_a(\bar{\mathbf{s}}_a, \tilde{\sigma}_a(\mathbf{s}_a)) + \delta EV_{a+1}(\bar{\mathbf{s}}_{a+1}(\mathbf{s}_a, \tilde{\sigma}_a(\mathbf{s}_a))).$$

This completes the recursion.

Less formally, in equilibrium the individual’s decision making proceeds as follows: beginning at age a_0 , and given the deterministic elements of her state $\bar{\mathbf{s}}_{a_0}$, the individual draws a fertility shock and three choice specific shocks from the joint ϵ_{a_0} distribution. She then calculates both the realized current rewards and the expected future rewards from each of her the three alternatives,

given the anticipated behavior of her future selves. This calculation yields $\sigma_a^{PN}(\mathbf{s}_a)$, representing the alternative that yields the highest discounted present value. The state space is then updated according to the alternative chosen, and the process is repeated. The perception-perfect strategy at each age a is identified by the set of regions in the three dimensional (ϵ_a) space over which each of the alternatives would be optimal. There is no closed-form representation of this solution. Therefore, in the estimations and simulations described below the game is solved numerically by backward recursion using crude Monte Carlo integration to approximate each of the expected continuation utilities EV_{a+1} .¹⁶

4 Estimation Method

The solution to the intrapersonal game described above provides the inputs for estimating the parameters of the model by the following method: Consider having data on a sample of individuals each of whom is assumed to be solving the intrapersonal game. The decision at every age is deterministic for the agent, but probabilistic from our perspective since we do not observe the shocks ϵ_a . For each individual $i = 1, \dots, N$, the data consist of two sequences: 1) the predetermined elements of the state space, $[\bar{s}(a)]_{a=a_{i0}}^{\bar{a}_i}$, where a_{i0} is the age at which individual i gave birth to her first child and $\bar{a}_i \leq A$ is the age after which we no longer observe the agent i ; and 2) the welfare benefits available to the agent, and her wages if she worked $[G_{ja}, R_a(1|d_{ia} = 1)]_{a=a_{i0}}^{\bar{a}_i}$.

The solution to the game provides the probability of choosing alternative d and receiving $R_a(d)$ at age a , given deterministic state $\bar{s}(a)$, number of children n_a and the model's parameters Θ :

$$\Pr [d, R_a(d) \mid \bar{s}(a), n_a, \Theta].$$

We can therefore consistently estimate Θ by maximizing the sample likelihood

$$\mathcal{L} = \prod_{i=1}^N \prod_{a=a_{i0}}^{\bar{a}_i} \Pr [d_i, R_a(d_i) \mid \bar{s}_i(a), n_{ai}, \Theta]$$

with respect to Θ .¹⁷

To ease the computational burden, we separate this estimation into three parts. First, the parameters of the welfare benefits function $G_j(n_a)$, θ_j , are taken as the mean estimated real

¹⁶The numerical solution method we employ follows closely Keane and Wolpin (1994). However, because the state space of our model is relatively small (roughly 150,000 elements at age $A = 34$), we do not use Keane and Wolpin's method for approximating the expected continuation utilities using only a subset of the state space. Instead we approximate the expected continuation utility for *every* element of the state space by Monte Carlo integration. Based on sensitivity analysis we chose to rely on 225 draws from the ϵ distribution to perform this integration.

¹⁷In our estimation, the terminal value at age 34 is modelled as a linear function of the state variables.

benefit in the agent’s municipality j over the period observed.¹⁸ Table A1 of the appendix presents these estimated parameters for the 20 U.S. states represented in the sample.

Second, the parameter of the fertility function $\rho(a, n_a, d_a) = \Pr(n_{a+1} = n_a + 1 | a, n_a, d_a)$ is estimated by a logit of the form:

$$\ln \left(\frac{\rho(a, n_a, d_a)}{1 - \rho(a, n_a, d_a)} \right) = \gamma_0 + \gamma_1 a + \gamma_2 n_a + \gamma_3 I(d_a = 0) + \gamma_4 I(d_a = 1) + \nu. \quad (5)$$

Finally, given this set of estimated parameters denoted by $\hat{\Theta}'$ the remaining parameters of the utility function, the returns functions, and the variance-covariance matrix of the shocks Ω , denoted by Θ'' are estimated by choosing Θ'' to maximize:

$$\mathcal{L}'' = \prod_{i=1}^N \prod_{a=a_{i0}}^{\bar{a}_i} \Pr \left[d_i, R_a(d_i) \mid \bar{s}_i(a), n_{ai}, \hat{\Theta}', \Theta'' \right].$$

For each observation i , $\Pr \left[d_i, R_a(d_i) \mid \bar{s}_i(a), n_{ai}, \hat{\Theta}', \Theta'' \right]$ is a three-dimensional integral which we approximate using 500 Monte Carlo draws to form kernel-smoothed simulators of the probabilities.¹⁹

5 Data

5.1 Sample Definition

The sample data are taken from the 1979 youth cohort of the National Longitudinal Surveys of Labor Market Experience (NLSY). The NLSY began in 1979 with 6,283 women ages 14-21, and has interviewed this cohort annually up to the present. We restrict our subsample to the 675 women who, as of 1992, had both remained unmarried and given birth to at least one child during the years they were surveyed. We then consider only the decisions each individual made after the birth of her first child and during the calendar years 1978-1991, assuming she continued to reside in the state in which they lived at age 14.

Our purpose in selecting this subsample of individuals and years is threefold. First, to be consistent with our model, we want to restrict attention to those who, if they do not work, are

¹⁸The estimates of real welfare benefits as a function of the number of children are taken from “SOURCE” provided by Ken Wolpin. The municipality is defined as the U.S. state in which the respondent resided at age 14.

¹⁹We chose 500 draws after tests for sensitivity of the simulated probabilities to changes in the number of repetitions. The kernel of the simulated integral is given by:

$$\exp \left[\frac{U_d^a - \max_{d \in D} (U_d^a)}{\tau} \right] / \sum_{d=0}^2 \exp \left[\frac{U_d^a - \max_{d \in D} (U_d^a)}{\tau} \right]$$

In the estimation results that follow, the smoothing parameter τ is set to 150 – again based on sensitivity analysis.

eligible for welfare by virtue of having a child and being unmarried.²⁰ Second, to better justify our assumption that marriage decisions are not germane, we restrict attention to women who, because they never marry, are relatively far from margin making them indifferent between marriage and a single life. Third, we want to restrict our attention to decisions made before the major changes in welfare eligibility rules beginning in 1993, and easily anticipated by 1992. Finally, again to ease the computational burden, we further limit our sample to residents of the 20 states best represented in the data. This final restriction leaves us with 483 individuals taken from the NLSY’s core random sample and its oversamples of blacks and Hispanics.²¹ The women in our subsample were observed with at least one child for an average of 9.6 of the 14 years from 1978-1991, providing us with 4,652 state-choice observations for the estimation.

5.2 Period and Variable Definition

At each interview, the NLSY collects welfare participation data as a monthly event history recorded back to the preceding interview. The survey’s employment data are collected as a weekly event history. We assume the decision period of the model corresponds to a calendar year, and identify an agent as age a in a year if she was a years old for at least half of that year. We then define the decisions at each age a as follows: An individual is identified as having chosen welfare at age a if she received Aid to Families with Dependent Children (AFDC) for at least six months of the calendar year during which she was a years old. We will say an individual chose work at age a if she was employed for at least 1,500 hours of the calendar year during which she was a years old.²² We will say the agent chose to stay home if she chose neither of the above.²³

5.3 Descriptive Statistics

Descriptive statistics of the subsample are presented, by age, in Table A2 of the Appendix. Since none of the women in the subsample marries during the period she is observed, we know the

²⁰In fact, before PRWORA there were national asset limits for welfare eligibility. Prior to 1981 the national maximal liquid asset limit was \$2,000 and after 1981 it was \$1,000. States could and did set lower asset tests. We abstract away from the asset limits. It is not clear to us how our estimates would be systematically biased by this omission.

²¹Our subsample does not include members of the NLSY’s oversamples of poor whites and the military.

²²We do not distinguish between full- and part-time employment. The earnings from work are defined as the full time equivalent, real annual income from salary and wages. Like welfare benefits, earnings are denominated in 1987 New York dollars, as de(in)flated by the consumer price index.

²³On 19 occasions a respondent reported that she both received AFDC for at least 6 months of the previous calendar year *and* was employed for at least 26 weeks of that year. In these cases the agent was defined as having chosen welfare.

subsample is not typical of the population as a whole. To better understand the ways in which members of the subsample differ from average women, their statistics are compared with those of the entire sample of women in the NLSY from 1978-1991 in Table A3 of the Appendix.

Recognizing that the subsample is in some ways atypical of the population as a whole, we move on to consider its choices over time. The distribution of choices over welfare, work and home, is presented by age in Table 1 below. We concentrate on those ages 16-32 who represent 98% of the data. The fraction of the subsample choosing welfare increases sharply between ages 16 and 22. Of the 16 year-olds with at least one child, just 25% chose welfare while 51% of 22 year-olds with children chose welfare. The proportion choosing work exhibits a comparable increase over the same period; rising from 2.5% of 16 year-olds with children to 25.5% of 22 year-olds. Over this same age transition we observe a still steeper decline in the fraction of women with children choosing to remain at home; with 72.5% choosing to stay home at age 16 and just 23.3% choosing to stay at home at age 22.

Age	Choice			Total
	Welfare	Work	Home	
16	31.9	0.0	68.1	100.0
	15	0	32	47
17	38.5	0.0	61.5	100.0
	35	0	56	91
18	39.0	1.9	59.1	100.0
	62	3	94	159
19	46.2	8.8	45.0	100.0
	110	21	107	238
20	49.5	11.8	38.8	100.0
	143	34	112	289
21	49.9	14.9	35.2	100.0
	167	50	118	335
22	52.5	17.1	30.4	100.0
	190	62	110	362
23	50.4	20.7	28.9	100.0
	192	79	110	381
24	46.4	26.6	26.9	100.0
	183	105	106	394
25	45.4	27.9	26.7	100.0
	187	115	110	412
26	47.3	30.9	21.7	100.0
	196	128	90	414
27	43.1	32.6	24.4	100.0
	168	127	95	390
28	42.6	34.8	22.5	100.0
	142	116	75	333
29	41.0	38.1	20.9	100.0
	110	102	56	268
30	45.1	36.8	18.1	100.0
	92	75	37	204
31	40.8	41.4	17.8	100.0
	62	63	27	152
32	34.0	45.4	20.6	100.0
	33	44	20	97
Total	45.7	24.6	29.7	100.0
	2087	1124	1355	4566

It is important to note that not all of the movements in these age-decision profiles reflect the changing choices of the same individuals. The transitions are due in part to the fact that the composition of the sample is changing as the women age and, by virtue of having a child, join the subsample. To begin to assess the degree to which the choices of same individuals change over time, Table 2 presents the one-period transition rates between decisions by the same agent. Here we see evidence of considerable state-dependence of individuals' choices over time. The rows of Table 2 represent the choices made in period $t - 1$; the columns describe the choices made in the subsequent period. The top figure in each cell represents the fraction of the subsample that made the row choice in period $t - 1$ who went on to make the column choice in period t . The bottom figure in each cell shows the fraction of the subsample that made the column choice in period t who made the row choice in the previous period. We see 83.3% of those who chose welfare in period $t - 1$ go on to choose it again in period t . Similarly, 79.9% of those who chose work in period $t - 1$ go on to choose it again in period t . Decisions to remain at home are considerably less persistent. Of those who chose to stay home in period $t - 1$, 60.3% chose it again in period t .

Choice (t-1)	Choice (t)		
	Welfare	Work	Home
Welfare			
Row %	83.3	4.0	12.6
Column %	76.8	6.9	17.9
Work			
Row %	5.1	79.9	15.0
Column %	2.6	75.4	11.8
Home			
Row %	27.2	12.6	60.3
Column %	20.6	17.7	70.3

Table 3 explores in greater detail the state-dependence of decisions to work or receive welfare. We see, for example, that while 35.7% of the subsample with exactly one child chose welfare, more than 80% of those with exactly one child continued to choose welfare if they chose welfare in the previous period. Similarly, while 53.8% of the subsample with exactly two children chose welfare, 85.5% of those with exactly two children continued to choose welfare if they chose welfare in the previous period. We note that the level of persistence in the choice of AFDC declines somewhat with years of schooling at the birth of the first child. Among those with 11 years of schooling at the

birth of their first child, 85.9% continued to choose welfare if they chose it in the previous period. This fraction declines to 82.5% among those with 12 years of schooling, and to 76.1% of those with 13 years of schooling. There is a commensurate increase in the persistence of decisions to work with years of schooling.

Number of children	1	2	3	4	5	6	7	
Percent choosing AFDC	35.7	53.8	63.4	63.5	76.1	89.5	100	
If chose AFDC previous period	80.5	85.5	85.0	87.6	88.6	87.5	100	
Years of Schooling at Birth of First Child	9	10	11	12	13	14	15	16
Percent choosing AFDC	52.7	54.3	54.9	39.9	26.6	15.1	14.5	0.0
If chose AFDC previous period	82.4	84.8	85.9	82.5	76.1	78.6	62.5	0.0
Percent choosing Work	12.1	12.2	19.7	33.7	44.4	58.5	55.1	81.8
If chose Work previous period	75.4	81.4	78.5	79.1	82.7	88.3	80.6	86.2

6 Preliminary Results

6.1 Estimates of Θ'

The parameters of the government benefits and fertility functions (Θ'), estimated in the first stage, are presented in Tables A1 and A4 of the Appendix, respectively. Note that the estimates of these parameters, which are obtained separately from the estimates of the parameters of the utility function, do not vary with assumptions on the time-consistency of preferences or the degree of naiveté. The estimate of the fertility function's parameters suggests that the probability of an additional birth is decreasing with age and with the number of children. The estimate also indicates that the probability of an additional birth is lower for workers and higher for those on welfare, relative to staying home. We note, however, that our simple exogenous model of fertility explains very little of the variation in the timing of births in this subsample. The pseudo- R^2 is less than two percent.

6.2 Identification of β and δ

An important question concerning the estimation method is the separate identification of both discount parameters β and δ . In a companion paper, Fang and Silverman (2000), we show the potential for a variety of observable differences in the behavior of time-consistent and time-inconsistent agents in a simplified and deterministic version of the present model. With regard to identification of the two time discount parameters, the results of that paper suggest a pair of basic relationships.

The first is that when stigma $\phi(d_a)$ lasts for only one period after switching into welfare (as it does here), there is a large set of circumstances under which changes in the standard discount factor δ – holding β constant – will almost exclusively influence the margin between welfare and work, and not the margins between either welfare or work and home. With naive agents, the theory suggests the asymmetry of this influence should be particularly sharp. In particular, if the insights from the simple, deterministic setting of Fang and Silverman (2000) carry through to this more realistic and uncertain environment, we would expect that, holding β constant, variations in δ would affect the conditional probability of choosing work and welfare, but leave the probability of choosing home unchanged.

Figures 1-3 below summarize the results of four simulation experiments suggesting that at least some of these insights carry through, and that this first relationship holds in this more realistic model. Each of the experiments simulates 1,000 sequences of decisions of a naive agent from ages 18 to 34. In each experiment, the present bias factor β is set to 0.80. In experiment 1, the agent has a high degree of standard patience ($\delta = 0.97$). In the subsequent three experiments the standard discount factor is set to 0.95, 0.90, and 0.85, respectively. The remaining parameters of the model are held constant across experiments. Figures 1-3 show that, as predicted by theory, the probability of choosing home (conditional on age) appears largely unaffected by changes in δ , while the probability of choosing welfare (work) increases (decreases) as δ decreases. The intuition is that the standard discounting factor δ drives the *long-run* decisions, which here is “work”. When δ increases, the workers substitutes their choices from AFDC to work, while leaving the choices of home almost the same.

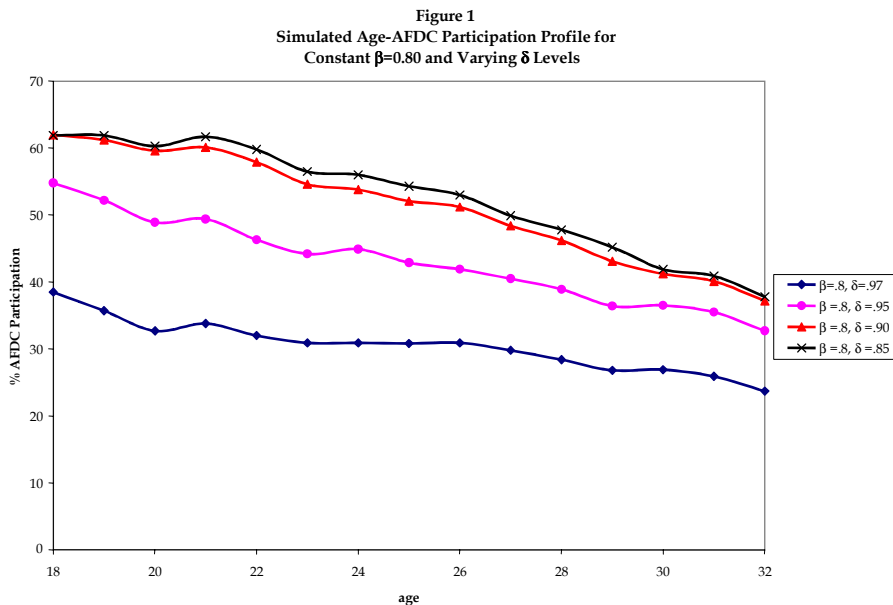


Figure 2
 Simulated Age-Work Participation Profile for
 Constant $\beta=0.80$ and Varying δ Levels





Figure 3
 Simulated Age-Home Profile for
 Constant $\beta=0.80$ and Varying δ Levels



The second relevant insight from Fang and Silverman (2000) is, again when stigma $\phi(d_a)$ lasts for only one period after a switch to welfare, that there is a large set of circumstances under which changes in the standard discount factor β – holding δ constant – will almost exclusively affect the margin between welfare and home, and not the margins between either welfare or home and work. Again, with naive agents the asymmetry of this influence should be most pronounced. In particular, we would expect that, holding δ constant, variations in β would affect the conditional probability of choosing welfare and home, but leave the probability of choosing home unchanged.

Figures 4-6 below summarize the results of another four simulation experiments suggesting that this second relationship also carries through. Again, each of the experiments simulates 1,000 sequences of decisions of a naive agent from ages 18 to 34. In this set of experiments, δ is always set to 0.95. In experiment 1, the agent is time consistent ($\beta = 1$). In the subsequent three experiments β is set to 0.90, 0.80, and 0.70, respectively. The remaining parameters of the model are held constant across experiments. Figures 4-6 show that, as predicted by theory, the probability of choosing work conditional on age appears largely unaffected by changes in β , while the probability of choosing home (welfare) increases (decreases) as β decreases. The intuition is that the present bias factor β drives the *short-run* decisions. Here because of our assumption of one-time stigma cost, varying β affect the margin between home and AFDC choices. As β increases, the agents will substitute their choices from home to AFDC while leaving the choice of work almost the same.

Figure 4
 Simulated Age-AFDC Participation Profile for
 Constant $\delta=0.95$ and Varying β Levels




Figure 5
 Simulated Age-Work Participation Profile for
 Constant $\delta=0.95$ and Varying β Levels





Figure 6
 Simulated Age-Home Profile for
 Constant $\delta=0.95$ and Varying β Levels



Of course, none of these exercises provides a definitive mathematical statement about the identification of the model. Taken together, however, the experiments suggest that the theoretical basis for identification described in our companion paper may be valid in this more realistic model.

6.3 Parameter Estimates and Simulations

Table 4 presents preliminary estimates of the parameters of the model under the assumption that agents are naive. We have postponed the calculation of standard errors until we are more

confident in the point estimates themselves. The estimated present bias factor $\beta = 0.69$ and the estimated $\delta = 0.87$, together imply a one-year ahead discount rate of 66.5%. Experimental work in psychology estimates this figure to be considerably lower – approximately 40% in an average population.²⁴ The estimated standard deviations of the choice specific shocks ε_0 (AFDC) and ε_2 (home) are large, but not unusual in comparable estimates such as those in Keane and Wolpin (1997, Pages 501 and 509).

Table 4
Estimated Parameters for Naifs


		parameter	point estimate	standard error
utility parameters	time	β	0.69	not calculated
	discounts	δ	0.87	not calculated
	net stigma	ϕ	9534.71	not calculated
	utility of	e_0	10190.02	not calculated
	leisure	e_1	-2795.20	not calculated
		e_2	315.75	not calculated
		e_3	474.53	not calculated
		e_4	22.88	not calculated
wage and skill parameters	constant	$\ln(r)+h_{a0}$	8.2100	not calculated
	education	α_1	0.0372	not calculated
	experience	α_2	0.0777	not calculated
	experience2	α_3	0.0010	not calculated
	experience decay	α_4	0.3264	not calculated
variance/covariance	std. dev ε_0		3193.12	not calculated
	std. dev ε_1		0.53	not calculated
	std. dev ε_2		10600.00	not calculated
	corr $\varepsilon_{0,2}$		0.26	not calculated

Using the estimated values of the parameters, Figure 7 presents wage offers as a function of experience for a high school graduate who begins a career of uninterrupted work at age 19. These predicted wage offers are compared with the estimated current-period return to uninterrupted welfare receipt (including the value of leisure, and the disutility of stigma) for this same agent with an average fertility experience, who has her first child at age 19 in a high, a medium and a low benefit state.²⁵

²⁴See sources cited in Laibson, et. al (1998).

²⁵The high benefit is the mean benefit among the seven most generous states in the sample. The medium benefit is the mean benefit among the next seven most generous states, the low is the mean benefit among the six least generous of the 20 states represented in the subsample.

Figure 7
 Estimated Undiscounted Returns to Uninterrupted Work and Welfare, Ages 19-34
 for HS Graduate in High, Medium and Low Welfare Benefits State



We note several basic features of the estimated streams of returns to work and welfare. First, for this agent who has neither worked nor received welfare in the past, work is estimated to have a current-period advantage over welfare net of stigma. This advantage is substantial (more than \$3,000 1987 dollars) in low benefits states, but negligible (\$289) in high benefits states. After the stigma has decayed, however, the current period return to welfare exceeds that of work for at least the next six years. Only after seven total years of work experience is the expected current-period return to work again greater than that of welfare in low benefits states. Ten years of work experience is required for the wage to catch up with the return to welfare in high benefits states.

Each of these features of the wage-experience profile relative to the welfare benefits schedule has important implications for optimal decisions, especially for time-inconsistent agents. The fact that, net of welfare stigma, work has a slight current period advantage for an agent who has neither worked nor received welfare before may make present-biased agents more likely to choose work or home rather than endure the immediate cost of stigma. The effects of this feature are counterbalanced, however, by the fact that (once stigma has decayed) it takes considerable time for the return to work to exceed that of welfare – even in low-benefits states. As a result, especially if preferences are time-inconsistent, these decision-makers may simply have difficulty beginning a career because to do so requires overcoming a bias towards the immediate reward of welfare over work.

Summarizing the interaction of these potentially complex and countervailing effects, Figures 8-10 compare the estimated model’s predicted distributions over the three alternatives welfare, work and home, to the actual distributions in the data, by age. The model’s predicted distributions

represent the simulated decisions of 1,000 agents in each of 16 cells defined to reflect the sample variation in initial conditions $j, a_0, g_0, x_0,$ and d_{-1} . There are four different j categories defined as high, medium-high, medium-low, and low benefits municipality. Similarly there are four g_0 categories defined as less than 11 years of education, 11 years of education, high school graduate, or some college at the birth of the first child. Within each of these 16 cells, the initial conditions are given by the sample average benefits (age, schooling, experience) level in the cell. The distribution of the 1,000 simulated decisions in each of these cells is then weighted by the proportion of the data falling into that initial condition cell to generate the predicted distributions appearing in Figures 8-10.

Figure 8
Age-AFDC Participation Profile:
Comparison of Data and Simulation with Estimated Parameters (Naïve Agents)




Figure 9
Age-Work Participation Profile:
Comparison of Data and Simulation with Estimated Parameters (Naïve Agents)





Figure 10
 Age-Home Profile:
 Comparison of Data and Simulation with Estimated Parameters (Naïve Agents)



The simulated age profiles match the data quite well. Each of the profiles implied by the estimated model assumes roughly the correct shape, and often matches the levels of the data closely. At young ages (17-18) that are poorly represented in the data, the model predicts too much welfare participation and too few remaining at home. Similarly, at older ages the model overestimates both the rate of decline in welfare participation and the rate of increase in the size of the population at work.

The estimated model’s predicted transitions are presented in Table 5. The model matches the persistence of both welfare and work choices quite well but greatly underestimates the persistence in home choices. To illustrate, the estimated model predicts that 80.6% of those who chose welfare in period $t - 1$ will go on to choose it again the following period. This figure should be compared with the 83.3% we observe in the data. Similarly, the model predicts that of those who chose to work in period t , 68.3% had chosen to work in the previous period. The comparable fraction in the data is 75.4%. However, the model predicts that of those who chose to stay home in period $t - 1$, just 37.8% will make the same choice the subsequent period. In the data, this proportion is 60.3%.²⁶

²⁶The inability of the model to match the persistence in choices to stay home may be due, in part, to specification error. Many of those defined as choosing home because they did not receive welfare or work 1,500 hours in the previous year may, in fact, have been working part time.

TABLE 5
Estimated Transition Matrix:
Naïve Agents

Choice (t-1)	Choice (t)		
	Welfare	Work	Home
Welfare			
Row %	80.6	2.1	17.3
Column %	81.3	2.5	28.4
Work			
Row %	4.5	71.5	24.1
Column %	3.6	68.3	32.1
Home			
Row %	23.5	38.8	37.8
Column %	15.1	29.2	39.5

Last, Figures 11 and 12 compare, respectively, the model’s mean wage-age and wage-experience profiles, with the analogous moments in the data. The model somewhat overestimates the average wages for those who choose to work at young ages and somewhat underestimates average wages for workers at older ages. The model similarly underestimates the mean wages of workers with three to five years of work experience while matching the average wages of workers with both less and more experience.

Figure 11
Mean Accepted Wages (FTE) for Workers by Age
Comparison of Data and Simulation with Estimated Parameters (Naïve Agents)





Figure 12
 Mean Wage Experience Profile
 Comparison of Data and Simulation with Estimated Parameters (Naïve Agents)



7 Conclusions and remaining work

Preliminary estimates of the structural parameters of a dynamic model of the work-welfare decision provide evidence of time-inconsistent preferences. Crude comparisons of the model’s predictions with corresponding moments of the data suggest the model matches the data reasonably well. As work on this project continues, we will give priority to testing the precision and robustness of our estimates with respect to changes in assumptions regarding sophistication and the decay of stigma and work experience. In particular, we will attempt to further assess whether the two discount parameters are estimated with adequate precision under varying assumptions. If we are confident in the robustness of our estimates we will then explore the predicted effects of policy experiments such as imposing welfare time limits or welfare work requirements. In each case, the emphasis would be placed on assessing the value of commitment to time-inconsistent agents.

8 Appendix

Table A1
Estimated Annual Welfare Benefits Function
and Summary Statistics (1987 Dollars)

U.S. State*	θ_0^{**}	θ_1^{**}	annual benefit for 1 child	annual benefit for 2 children	percent receiving welfare ⁺
1	2380.45	1238.01	3618.46	4856.48	39.6%
2	2467.68	1301.31	3768.99	5070.30	50.0%
3	2962.66	1203.84	4166.50	5370.34	32.9%
4	2979.62	1280.44	4260.06	5540.50	22.5%
5	3128.33	1340.02	4468.35	5808.38	39.2%
6	3493.63	1186.81	4680.45	5867.26	29.6%
7	3541.08	1251.03	4792.11	6043.13	50.3%
8	3985.20	1212.98	5198.18	6411.15	46.6%
9	4348.62	1098.98	5447.60	6546.58	28.2%
10	4358.47	1318.76	5677.23	6995.99	71.0%
11	4279.58	1419.96	5699.54	7119.50	51.2%
12	4509.59	1368.62	5878.21	7246.83	29.4%
13	4183.05	1539.27	5722.32	7261.59	13.6%
14	4592.94	1343.95	5936.89	7280.83	20.2%
15	4511.30	1411.63	5922.93	7334.57	66.8%
16	5005.98	1480.68	6486.65	7967.33	52.2%
17	4988.00	1577.07	6565.07	8142.15	27.0%
18	5634.63	1661.86	7296.49	8958.35	61.7%
19	5317.42	1851.81	7169.23	9021.04	69.7%
20	6264.03	1613.01	7877.04	9490.05	68.5%
mean	4146.61	1385.00	5531.62	6916.62	43.5%
std. dev.	1042.30	187.46	1182.51	1334.38	17.9%

Mean percentage receiving welfare in five most generous states = 56%

Mean percentage receiving welfare in five least generous states = 37%

* States are left unnamed to maintain the anonymity of survey respondents.

** The estimated annual welfare benefits function takes the form

⁺ Percent of sample living in the corresponding state that choose welfare.

Relative to all women in the NLSY during this period, the subsample has on average a similar number of children at every age. However, the subsample has an average of 0.87 fewer years of work experience, and 1.55 more years on AFDC at every age. At the birth of her first child, the typical woman in the subsample has also had fewer years of schooling (11.2 years in the subsample versus 11.9 years in the whole NLSY), and work experience (2.08 years versus 2.7 years). This difference reflects in part the fact that the average member of the subsample had her first child at age 20, while the average age of first birth among all women surveyed is 21.6 years old.

TABLE A2
Summary Statistics: Ages 16-32
Never Married Females With At Least One Child

Age	Mean				
	Number of Children	Years Work Experience	Years of Schooling*	Earnings for Workers**	Yrs. Received AFDC
16	1.23 (0.07)	0.00 (0.00)	8.81 (0.17)	n/a	0.13 (0.06)
17	1.21 (0.06)	0.00 (0.00)	9.41 (0.14)	n/a	0.23 (0.06)
18	1.25 (0.05)	0.00 (0.00)	10.04 (0.12)	10822.56 (2254.07)	0.35 (0.06)
19	1.34 (0.04)	0.05 (0.01)	10.41 (0.10)	6904.63 (753.06)	0.49 (0.06)
20	1.41 (0.04)	0.14 (0.02)	10.63 (0.09)	7361.80 (623.75)	0.78 (0.07)
21	1.46 (0.04)	0.25 (0.03)	10.79 (0.09)	7041.01 (581.24)	1.09 (0.08)
22	1.54 (0.05)	0.43 (0.05)	10.87 (0.08)	8053.38 (498.75)	1.46 (0.09)
23	1.66 (0.05)	0.68 (0.07)	10.91 (0.08)	8913.09 (370.28)	1.90 (0.10)
24	1.71 (0.05)	0.95 (0.08)	10.95 (0.08)	9270.13 (380.44)	2.30 (0.11)
25	1.77 (0.05)	1.31 (0.10)	11.03 (0.08)	9741.56 (402.59)	2.63 (0.13)
26	1.86 (0.05)	1.68 (0.12)	11.06 (0.08)	10095.84 (378.03)	3.00 (0.14)
27	1.89 (0.05)	2.02 (0.14)	11.09 (0.09)	10347.46 (393.13)	3.34 (0.16)
28	1.88 (0.06)	2.46 (0.17)	11.21 (0.09)	10551.26 (445.87)	3.53 (0.19)
29	1.92 (0.07)	2.68 (0.20)	11.34 (0.10)	10974.26 (488.75)	3.91 (0.22)
30	1.96 (0.08)	2.96 (0.25)	11.44 (0.12)	11265.91 (614.89)	4.38 (0.27)
31	2.02 (0.10)	3.30 (0.31)	11.54 (0.15)	13366.75 (1037.81)	4.64 (0.32)
32	2.02 (0.13)	4.16 (0.43)	11.68 (0.20)	11924.40 (837.73)	4.43 (0.41)

Standard errors are in parentheses

*Years of schooling at the birth of the first child. **Earnings are full-time equivalent, 1987 dollars.

TABLE A3
Summary Statistics: Ages 16-32
All Females With At Least One Child

Age	Mean				
	Number of Children	Years Work Experience	Years of Schooling*	Earnings for Workers**	Yrs. Received AFDC
16	1.00 (0.07)	0.01 (0.01)	9.49 (0.09)	3070.11 (891.05)	0.07 (0.02)
17	1.00 (0.04)	0.08 (0.02)	9.80 (0.07)	6101.60 (478.93)	0.10 (0.02)
18	1.06 (0.03)	0.24 (0.02)	10.10 (0.05)	6072.58 (350.36)	0.17 (0.02)
19	1.14 (0.02)	0.48 (0.02)	10.44 (0.04)	7260.15 (235.60)	0.21 (0.02)
20	1.24 (0.02)	0.83 (0.03)	10.68 (0.04)	8078.61 (192.50)	0.29 (0.02)
21	1.38 (0.02)	1.20 (0.03)	10.88 (0.04)	8113.51 (185.78)	0.40 (0.02)
22	1.47 (0.02)	1.64 (0.03)	11.03 (0.04)	9033.08 (167.89)	0.53 (0.02)
23	1.58 (0.02)	2.16 (0.04)	11.17 (0.03)	9371.49 (161.68)	0.65 (0.03)
24	1.67 (0.02)	2.71 (0.04)	11.30 (0.03)	10035.28 (152.23)	0.76 (0.03)
25	1.75 (0.02)	3.36 (0.05)	11.42 (0.03)	10429.77 (163.53)	0.87 (0.03)
26	1.83 (0.02)	4.07 (0.05)	11.55 (0.03)	11000.68 (162.03)	0.96 (0.03)
27	1.91 (0.02)	4.82 (0.06)	11.66 (0.03)	11641.47 (170.46)	1.06 (0.04)
28	1.98 (0.02)	5.54 (0.07)	11.74 (0.03)	12024.83 (183.73)	1.15 (0.04)
29	2.03 (0.02)	6.14 (0.08)	11.80 (0.03)	12198.27 (206.30)	1.18 (0.04)
30	2.07 (0.02)	6.73 (0.09)	11.85 (0.03)	12413.90 (224.45)	1.20 (0.05)
31	2.12 (0.03)	7.25 (0.11)	11.87 (0.03)	13085.85 (290.35)	1.23 (0.06)
32	2.13 (0.03)	7.95 (0.14)	11.90 (0.03)	12539.09 (320.50)	1.29 (0.07)

Standard errors are in parentheses

*Years of schooling at the birth of the first child. **Earnings are in 1987 dollars.

Table A4: Logit Estimate of the Fertility Function (5)

parameter	Estimate	Std. Error	P-Value
γ_0	-.8105135	.322526	.012
γ_1	-.0440472	.014757	.003
γ_2	-.0766624	.0586878	.191
γ_3	.0943619	.1154835	.414
γ_4	-.4935787	.161612	.002
No. of Obs.:	3911		
LR $\chi^2(4)$:	38.20		
Log Likelihood:	-1387.1602		
Pseudo R^2 :	0.0136		

References

- [1] Ainslie, George (1992). *Picoeconomics: The Strategic Interaction of Successive Motivational States Within the Person*. Cambridge [England]; New York: Cambridge University Press.
- [2] Barro, Robert J. (1999). “Ramsey Meets Laibson in the Neoclassical Growth Model.” *Quarterly Journal of Economics* **114(4)**: 1125-1152.
- [3] Carillo, J. and T. Mariotti (2000). Strategic Ignorance as a Self-Disciplining Device.” *Review of Economic Studies* **67(3)**:529-544.
- [4] DellaVigna, Stefano, and M. Daniele Paserman (2000). “Job Search and Hyperbolic Discounting.” Mimeo Harvard University.
- [5] Edin, Kathryn and Laura Lein (1997). *Making Ends Meet: How Single Mothers Survive Welfare and Low-wage Work*. New York: Russell Sage.
- [6] Fang, Hanming, and Dan Silverman (2000). “On the Compassion of Time-limited Welfare Programs.” Mimeo University of Pennsylvania.
- [7] Jovanovic, Boyan, and Dmitriy Stolyarov (2000). “Ignorance is Bliss.” Mimeo New York University.
- [8] Keane, Michael P., and Kenneth I. Wolpin, (1994). “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence.” *The Review of Economics and Statistics* **76(4)**:648-672.
- [9] _____ (1997). “The Career Decisions of Young Men.” *Journal of Political Economy* **105(3)**:473-522.
- [10] ——— (2000). “Estimating the Effect of Welfare on the Education, Employment, Fertility and Marriage Decisions of Women”, mimeo, New York University and University of Pennsylvania.
- [11] Laibson, David (1997). “Golden Eggs and Hyperbolic Discounting.” *Quarterly Journal of Economics* **112**: 443-477.
- [12] Laibson, David, Andrea Repetto, Jeremy Tobacman, (1998). “Self-Control and Saving for Retirement.” *Brookings Papers on Economic Activity* **1**:91-196.
- [13] McFadden, Daniel (1989). “A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration.” *Econometrica* **57(5)**:995-1026.

- [14] Moffit, Robert (1983). "An Economic Model of Welfare Stigma." *American Economic Review* **73(5)**:1023-35.
- [15] Mulligan, Casey B. (1997). "A Logical Economist's Argument Against Hyperbolic Discounting." Mimeo University of Chicago.
- [16] O'Donoghue, Ted and Matthew Rabin (1999a). "Doing It Now or Later." *American Economic Review* **89(1)**: 103-124.
- [17] _____ (1999b). "Addiction and Self-Control." In Jon Elster, editor, *Addiction: Entries and Exits*. New York: Russell Sage.
- [18] Phelps, Edmund S., and Robert A. Pollak (1968). "On Second-best National Saving and Game-equilibrium Growth." *Review of Economic Studies* **35(2)**: 185-199.
- [19] Pollak, Robert A. (1968). "Consistent Planning." *Review of Economic Studies* **35(2)**: 201-208.
- [20] Sanders, Seth G. and Robert A. Miller (1997), "Human Capital Development and Welfare Participation", *Carnegie-Rochester Conference Series on Public Policy*, **Vol** 46, 1-43.
- [21] Strotz, Robert H. (1955). "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* **23(3)**: 165-180.
- [22] Swann, Christopher (1996). "Welfare Reforms when Agents are Forward-looking", mimeo, University of Virginia.
- [23] Thaler, Richard H. (1991). *Quasi Rational Economics*. New York: Russell Sage.