Welfare Dynamics under Time Limits

Jeff Grogger
UCLA and NBER

Charles Michalopoulos
MDRC

First draft: June 16, 1998
This draft: August 24, 1999

The authors thank Dan Bloom, Janet Currie, V. Joseph Hotz, David Levine, Robert Moffitt, Seth Sanders, and seminar participants at MDRC and UC Davis for helpful comments and discussions. The views expressed herein are those of the authors, and do not represent the positions of MDRC or the State of Florida. Any errors are the responsibility of the authors.

Abstract

Among the most important changes brought about by the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA) is the imposition of time limits. In this paper, we analyze a simple model in which a potential welfare recipient chooses how to allocate her time-limited endowment of benefits so as to maximize her expected lifetime utility. Not surprisingly, the model reveals that time limits provide an incentive for the consumer to conserve, or bank, her benefits. More interesting is the prediction that these incentives to conserve one's benefits vary inversely with the age of the youngest child in one's family. This implies that the reduction in welfare payments that results from PRWORA will fall disproportionately on families with young children.

We estimate age group-specific effects of time limits and test the prediction of the model using data from a welfare reform demonstration in Florida. Subject to some assumptions that are necessary to distinguish the effects of time limits from the effects of other provisions of the demonstration, we find that time limits indeed reduce welfare use by the greatest amount among the families with the youngest children. Moreover, time limits have substantial effects on welfare utilization, reducing monthly utilization probabilities by 19 percent. Time limits lead families to exit the welfare rolls well before they exhaust their benefits, suggesting that welfare mothers are rational in the sense of being forward-looking.

Introduction

The U.S. welfare system has changed dramatically as a result of the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA). Among the most important of the changes brought about by the Act is the imposition of time limits. Under the old AFDC program, welfare benefits were an entitlement: all poor, single-parent families with at least one child under 18 years of age were eligible to receive aid.[1] The duration of welfare receipt was limited only by the age of the youngest child in the family. Under PRWORA's new TANF program, benefits remain payable largely to poor, single-parent families with children under 18, but the duration of welfare receipt is sharply limited. Federal law now allows families to receive benefits for no more than five years. Many states set stricter limits.[2]

One of the implications of time limits is that the familiar static model long used to study welfare incentives no longer suffices to analyze consumer behavior. Whereas the consumer's problem was essentially static under the old entitlement regime, it is inherently dynamic under time limits. One of the contributions of this paper is to provide a theoretical model that embeds the leisure-consumption choice from the static model within a dynamic framework under which current choices about welfare utilization may affect future constraints.

The model yields both a testable implication and a restriction that is helpful for identifying the effects of time limits in our empirical work. The testable implication is that time limits should affect families differently depending on the age of their youngest child. Families with the youngest children face the longest horizon over which welfare could be used to smooth consumption in the event of an adverse wage realization. Therefore they have the greatest incentive to conserve, or bank, their welfare benefits. This implies that families with the youngest children should be the first to leave the welfare rolls when time limits are imposed. The model also yields an identifying restriction, since it implies that families whose oldest children are above a threshold age--age 13 in the case of a five-year time limit--are unaffected by the imposition of the time limit. For these families, the time limit amounts to a non-binding constraint.

For our empirical analysis we use data from Florida's Family Transition Program (FTP). Subject to a number of identifying assumptions required to isolate the effect of time limits from the effects of other provisions of the program, we find that families with younger children indeed exit the welfare roles more rapidly once time limits are imposed. Moreover, we find strong anticipatory responses to time limits. Our estimates indicate that, in the absence of other provisions of FTP that tended to increase aid utilization, time limits would have caused welfare use to fall by 19 percent before any of the families had

---

[1] Some poor, married-couple families were also eligible under the Unemployed Parent component of AFDC, but this program never accounted for more than a small fraction of the AFDC caseload (Moffitt 1992).

[2] States can exempt up to 20 percent of their caseload from the five-year time limit. In addition, states can allow more individuals to continue receiving cash assistance for more than five years as long as they do not use federal funds for payments past the five year limit. Several states have adopted plans that use this option. Several other states, including California, eliminate only the adults' portion of welfare grants when the five-year limit is imposed, allowing the families to continue receiving the children's portion of the grants.

actually exhausted their benefits. Contrary to expectation, our estimates suggest that even relatively disadvantaged aid recipients respond to the time limit.

By focusing primarily on time limits and using data from a randomized trial, our analysis differs from a number of recent studies that have estimated the effects of welfare reform more generally (Blank 1997; Council of Economic Advisors 1997; Wallace and Blank 1999; Ziliak, et al., 1997). These studies focus on welfare caseloads, utilizing aggregate state-level data measured either monthly or annually. With the exception of Ziliak, et al. (1997), they generally find welfare reform to have at least a marginally significant effect on caseloads, as does Moffitt (1999), who analyzes data from the Current Population Survey (CPS). Their results regarding the effects of time limits are mixed, and at any rate, they do not (and with aggregate data, can not) disaggregate the effects of time limits by age. Moreover, one set of authors constrains time limits to have the same effect as work requirements (a distinct type of welfare reform policy), thus failing to isolate the effects of time limits even on aggregate caseloads (Ziliak, et al. 1997). Others have been criticized as incorrectly characterizing some states' welfare reform plans, causing their models to be misspecified (Council of Economic Advisors 1997; Martini and Wiseman 1997).

To our knowledge, the only other paper to provide both a theoretical and empirical analysis of the effects of time limits is Swann (1998).[3] Swann's model, like ours, captures the important dynamic elements of the consumer's problem under time limits. In most other respects, however, our papers differ substantially. Swann's model is more general than ours in some dimensions, but ours is more general in others. Moreover, our focus differs from Swann's. Whereas he uses numerical simulations to study welfare utilization (and other outcomes) over time, we focus on deriving analytical results. Perhaps the greatest difference between our studies concerns our empirical work. Swann uses individual data collected under the old AFDC system to estimate a dynamic model of welfare utilization, then uses the parameters from that model to simulate how utilization changes in response to time limits. In our empirical work, in contrast, we compare the behavior of consumers who were actually subject to a time limit to that of consumers for whom welfare remained an entitlement.

The Model

We begin with a simple model of welfare receipt under time limits, abstracting from a number of important considerations such as job search, welfare stigma, and human capital formation. Although incorporating these aspects of the consumer's problem might add valuable realism to the model, it would also add substantial complexity. Our goal in focusing on a parsimonious model is not to deny the importance of other elements of the consumer's problem, but rather to focus on the effects of time limits *per se* and derive predictions that can be studied using the experimental data at our disposal.

[3] Bloom et al. (1997, 1998) have studied the FTP program using the data we analyze here, but they did not attempt to isolate the effects of the time limit from the effects of the other program treatments.

The consumer's problem under time limits

As was noted in the introduction, aid is available under TANF (and was available under AFDC) only to families with minor children in the home. This means that the consumer's *period of eligibility* for aid, or equivalently, her *eligibility horizon*, is *T* periods long, where *T* is the number of years until the consumer's youngest child turns 18. We take *T* to be exogenous, but note the implications of endogenous fertility below.

At the beginning of her period of eligibility, the consumer is endowed with an *initial stock of benefits.* This stock is denominated in periods of benefit receipt, consistent with the language of PRWORA. The initial stock, which we denote as *N*, is equal to the *time limit*.

If the consumer works in period *t*, she receives a gross wage of w_t. Wages at time *t* are stochastic, however, and unknown until period *t* begins. We assume that the consumer gets a new i.i.d. wage draw each period from the known distribution function $F(w)$. The new wage draw each period captures the observation that employment in the low-wage sector tends to be fairly unstable, particularly for welfare recipients (Edin and Lein 1997; Loeb and Corcoran 1999). We assume that w_t becomes known at the beginning of period *t*, after which the consumer chooses how much to work and, if she has not already exhausted her benefits, whether to utilize welfare. Of course, the amount of time she spends working may depend on whether she utilizes welfare. We denote her hours of work and hours of leisure at time *t* as $h_t(I_t)$ and $L_t(I_t)$, respectively, where I_t is the welfare utilization indicator, equal to one if the consumer utilizes benefits at time *t* and zero otherwise. Each period the consumer faces a time constraint given by $h_t(I_t) + L_t(I_t) = \overline{L}$, where $\overline{L}$ denotes total time available.

We denote consumption at time *t* by $C_t(I_t)$, the price of which is normalized to one. We assume that there is no borrowing or saving, which is realistic for a low-income population (Edin and Lein 1997). Letting t denote the benefit reduction rate, that is, the rate at which benefits are taxed as the consumer's earnings increase, the consumer's current-period budget constraint is given by $C_t(1) = G + w_t(1-t)h_t(1)$ if she elects to receive welfare. We refer to this as the "welfare budget constraint." Her budget constraint is given by $C_t(0) = w_t h_t(0)$ if she foregoes welfare, which we refer to as the "no-welfare budget constraint."

The consumer's current-period utility function is given by $U(L_t, C_t)$ and is increasing and concave in both arguments. We assume that both consumption and leisure are normal goods.[4] The consumer's problem is to choose hours of work and welfare utilization so as to maximize the expected present value of lifetime utility, subject to the time limit and her time and budget constraints. A formal statement of the consumer's problem is:

[4] We further assume that $U_C \to -\infty$ as $C \to 0$, where U_C denotes the derivative of the utility function with respect to *C*. This final "Inada condition" is not strictly necessary for our results, but it does help simplify some of the diagrams below.

$$\max_{\substack{I_1,\ldots,I_T \\ h_1,\ldots,h_T}} \sum_{t=1}^{T} r^{t-1} EU(L_t, C_t)$$

s.t. $L_t(I_t) + h_t(I_t) = \bar{L}$

$C_t(I_t) = (1 - I_t) w_t h_t(0) + I_t[G + w_t(1-t)h_t(1)]$

$S_1 = N$

$S_t = S_{t-1} - I_{t-1} \qquad t = 2,\ldots,T$

$S_{T+1} \geq 0$

$w_t \sim F(w_t)$ i.i.d.

where $r \in (0,1)$ is the discount factor. The variable S_t is the state variable, giving the stock of benefits remaining at the beginning of period t. At time t =1 it is equal to the time limit, and it must be non-negative at the beginning of period T+1, which is the end of the consumer's eligibility horizon. In between, utilizing benefits at time t depletes the remaining stock of benefits by one unit. Thus $S_t = S_{t-1} - 1$ if benefits are utilized in period t-1 and $S_t = S_t$-1 if not. This simple stock depletion rule has important implications for the solution of the model.

The solution of the model

To characterize the solution, let $V_t(S_t)$ denote the optimal value of the consumer's problem starting from time t with remaining benefits S_t. V_t satisfies the usual recursion relationship

$$V_t(S_t) = \max_{h_t, I_t}[U(L_t, C_t) + rEV_{t+1}(S_{t+1})].$$

Because S_t depends on I_t but not h_t, the consumer's problem simplifies considerably. In each period, the consumer solves a two-part problem. First, she chooses optimal current-period labor supply along each of the welfare and no-welfare budget constraints. Second, she chooses whether to utilize welfare so as to maximize expected utility over her remaining eligibility horizon. In other words, she first solves her current-period maximization problem subject to each of her possible current-period budget constraints, then chooses the budget constraint (i.e., decides whether to utilize welfare) so as to achieve the global maximum.[5]

The solutions to the first-part problem satisfy the usual tangency conditions. Define the consumer's current-period marginal rate of substitution as a function of her labor supply and consumption:

$m(h_t(I_t), C_t(I_t)) \equiv U_L(\bar{L} - h_t(I_t), C_t(I_t)) / U_C(\bar{L} - h_t(I_t), C_t(I_t))$. Then the consumer's optimum along the welfare budget constraint satisfies $m(h_t^*(0), C_t^*(0)) \geq (1-t)w_t$, where the inequality holds strictly if $(1-t)w_t < m(0, G)$, that is, if her net wage draw is less than her shadow price of leisure. The consumer's optimum along the no-welfare budget constraint satisfies $m(h_t^*(1), C_t^*(1)) = w_t$.

The second part of the consumer's two-part problem is to choose whether to utilize welfare so as to solve:

---

[5] The mechanics of the problem thus are similar to those in Moffitt's (1983) static model.

$$\max_{I_t} \{U(\bar{L} - h_t^*(0), C_t^*(0)) + rEV_{t+1}(S_t),\ U(\bar{L} - h_t^*(1), C_t^*(1)) + rEV_{t+1}(S_t - 1)\}.$$

It proves useful to write the consumer's second-part problem in terms of her current-period wage realization. To do this, define the consumer's maximized, or indirect, utility function along a budget constraint defined by non-labor income A_t and a net wage of W_t as

$$\begin{aligned} v(W_t, A_t) &= U(\bar{L}, A_t) && \text{if } m(0, A_t) > W_t \\ &= \max_h [U(L_t, C_t) \mid C_t = W_t h_t + A_t;\ \bar{L} = L_t + h_t] && \text{if } m(0, A_t) \le W_t. \end{aligned}$$

Then the second part of the consumer's problem can be written as:

$$\max_{I_t} \{v(w_t, 0) + r\ EV_{t+1}(S_t),\ v(w_t(1-t), G) + r\ EV_{t+1}(S_t - 1)\}.$$

For low wage realizations, utilizing benefits raises current-period utility. Utilizing benefits today, however, reduces the stock of benefits remaining for the future. In the case where $S_t < T - t + 1$, so the remaining stock of benefits is less than the remaining number of periods in the eligibility horizon, utilizing benefits today reduces the expectation of maximized utility beginning in period t +1. In the dynamic model, the consumer will utilize benefits today only if the current-period utility gain is great enough to offset the discounted loss in expected future utility, that, only if

$$v(w_t(1-t), G) - v(w_t, 0) > r[EV_{t+1}(S_t) - EV_{t+1}(S_t - 1)]. \qquad (1)$$

Panel A of Figure 1 depicts the consumer's welfare utilization decision. It plots two functions against the current-period wage realization. The first is the current-period utility gain that results from current-period welfare utilization, denoted by $d(w_t) \equiv v(w_t(1-t), G) - v(w_t, 0)$. The second is the discounted gain in expected future utility that results from foregoing current-period welfare utilization, denoted by $\Delta_{t+1}(S_t) \equiv r[EV_{t+1}(S_t) - EV_{t+1}(S_t - 1)]$. The current-period utility gain $d(w_t)$ is large for low values of w_t, since for a low wage draw, utilizing welfare can raise consumption substantially. The current-period utility gain from utilizing aid falls as the wage realization improves.[6] In contrast, the expected future utility gain from foregoing current-period welfare, $\Delta_{t+1}(S_t)$, is not a function of the current-period wage realization, and hence graphs as a horizontal line.

At the wage corresponding to the intersection of these two curves, the consumer is just indifferent about utilizing welfare in the current period. We term this wage the consumer's *reservation wage for welfare utilization*, and denote it by $\tilde{w}_t(S_t)$. Algebraically, the consumer's reservation wages satisfies:

$$v(\tilde{w}_t(S_t)(1-t), G) + rEV_{t+1}(S_t - 1) = v(\tilde{w}_t(S_t), 0) + rEV_{t+1}(S_t). \qquad (2)$$

Figure 1 shows that the consumer's welfare utilization problem satisfies the reservation wage property: for wage realizations less than or equal to $\tilde{w}_t(S_t)$, the consumer utilizes

---

[6] There is a kink in $d(w_t)$ at $w_t = m(0, G)$ because the consumer will choose not to work for values of $w_t \le m(0, G)$. Thus $v(w(1-t), G)$ is constant (and equal to $U(\bar{L}, G)$) for $w_t \le m(0, G)$, but increasing and concave for $w_t > m(0, G)$, giving rise to a kink.

welfare in the current period; for realizations greater than $\tilde{w}_t(S_t)$, she foregoes current-period welfare utilization.

Implications

A number of implications follow from the reservation wage property of the model. First, the reservation wage is highest in the final period of the eligibility horizon. At period *T*, the value of preserving benefits for the next period is zero. Put differently, at period *T* the consumer's problem is effectively static, since today's choice does not affect future utility. Denote the final-period reservation wage by $\tilde{w} \equiv \tilde{w}_T(S_T)$ for $S_T > 0$.

This notion generalizes. If at any time period *t*' the remaining benefit stock $S_{t'}$ (weakly) exceeds the remaining number of periods in the eligibility horizon, then the consumer's problem is again essentially static, because the period of eligibility will end (weakly) before the remaining benefit stock can be exhausted. Because this result provides an important restriction for our empirical identification analysis in Section IV, we state it as a lemma:

Lemma: Let $S_{t'} \geq T - t' + 1$. Then $\tilde{w}_t(S_t) = \tilde{w}$ for $t' \leq t \leq T$.

It follows that an entitlement is formally a special case of time limits, since for $N = T$, we have $S_t \geq T - t + 1$ for $t = 1,..., T$.

To establish our principal result, we note that the reservation wage in period *t* satisfies equation (2),whereas if $I_t^* = 0$, then the reservation wage in period *t*+1 will satisfy

$$v(\tilde{w}_{t+1}(S_t)(1-t), G) + r\,EV_{t+2}(S_t - 1) = v(\tilde{w}_{t+1}(S_t), 0) + rEV_{t+2}(S_t). \qquad (3)$$

Comparing equations (2) and (3), we see that if benefits are not utilized in period *t*, then the reservation wage at period *t*+1 will exceed the reservation wage at period *t*, provided that $EV_{t+1}(S_t) - EV_{t+1}(S_t - 1) > EV_{t+2}(S_t) - EV_{t+2}(S_t - 1)$, that is, provided that the value (in expected utility terms) of the marginal unit of benefits falls, the closer is the end of the eligibility period. But the value of the marginal unit of benefits falls as *t* rises for a simple reason: the closer one is to the end of one's eligibility horizon, the greater the likelihood that the remaining benefit stock eventually will exceed the remaining eligibility horizon, in which case the consumer's problem becomes static and the marginal unit of benefits has a zero shadow price. We state this result in the form of a proposition, the proof of which is provided in the Appendix:

Proposition (rising reservation wages)—Let $S_t < T - t + 1$. Then provided that $I_t^* = 0$, $\tilde{w}_{t+1}(S_t) > \tilde{w}_t(S_t)$. In other words, all else equal, the consumer becomes less reluctant to utilize welfare as time passes (i.e., as *t* rises).

This result is depicted graphically in Panel B of Figure 1. Since neither $\Delta_{t+1}(S_t)$ nor $\Delta_{t+2}(S_t)$ is a function of the current-period wage realization, both graph as horizontal lines. Since the value of the marginal unit of benefits is lower, the closer the consumer is to the end of her period of eligibility, $\Delta_{t+2}(S_t) < \Delta_{t+1}(S_t)$. Thus $\Delta_{t+2}(S_t)$ intersects $d(w)$ farther to the right than does $\Delta_{t+1}(S_t)$.

This result is fairly intuitive. Welfare allows the consumer to smooth consumption in the face of adverse wage shocks. Because of time limits, however, there are limits to how much smoothing she can do. Early in the eligibility period, the odds of realizing an adverse wage draw at some point in the future are high. The consumer reacts by waiting to draw down her benefits. As time passes, however, the odds of adverse future wage draws fall, so the consumer becomes more willing to utilize her benefits.

Put somewhat differently, the value of preserving one's option to draw benefits in the future is high when the eligibility horizon comprises many periods. As the end of the eligibility horizon approaches, however, the value of preserving one's benefits falls. The consumer can afford to exhibit less reluctance in utilizing welfare.

This result has two empirical implications that can be tested with data on welfare-eligible families observed at the time that time limits are imposed. We term these implications the Exit Hypothesis and the Entry Hypothesis:

Exit hypothesis—Upon the imposition of time limits, families with younger (youngest) children will leave the welfare rolls more quickly than families with older (youngest) children.

Entry hypothesis—Upon the imposition of time limits, families with younger (youngest) children will be less likely to enter the welfare rolls than families with older (youngest) children.

When time limits are imposed initially, families utilizing welfare receive the same initial endowment of benefits, but face substantially different eligibility horizons. Families whose youngest children are very young have long eligibility horizons and thus have low reservation wages. Families whose youngest children are close to age 18 have short eligibility horizons and high reservation wages. For families receiving welfare when time limits are imposed, this means that families with younger (youngest) children are less likely to continue utilizing benefits than families with older (youngest) children. For families not receiving benefits, this means that families with younger (youngest) children are less likely to sign up for welfare than families with older (youngest) children.

Although the entry and exit hypotheses, as stated, pertain to families involved in the transition from AFDC to TANF, the implications of the general result appear to be broader and yield predictions regarding the steady-state age distribution of children receiving welfare under TANF as compared to AFDC. Under AFDC, entering cohorts of welfare-eligible families (predominantly families becoming eligible due to an unwed birth or divorce) had no incentive to hoard their benefits. Under TANF, entering cohorts will have different incentives to hoard their benefits depending on the age of the youngest child in the family. On average, families with older (youngest) children will be more likely to use their benefits, whereas families with younger (youngest) children will be less likely to use theirs. As a result, in an otherwise stationary environment, the move from AFDC to TANF should raise the average age of (youngest) children receiving aid.

The rising reservation wage result may have important substantive implications as well. Taken together, the entry and exit hypotheses indicate that the reductions in welfare expenditures that result from the imposition of time limits will fall disproportionately on families with young children. In other words, any reductions in welfare receipt that a

child experiences due to time limits are more likely to occur earlier, rather than later, during childhood.

If decreasing welfare payments cause poverty rates to rise, then this may affect the child's ability and ultimate educational attainment. Guo (1998) studies ability and achievement tests in the Children of the National Longitudinal Survey of Youth data. She finds that poverty in childhood has a strong negative influence on the child's measured ability, but that poverty in adolescence has little such effect. This is broadly consistent with findings of Duncan et al. (1998), who study data on educational attainment from the PSID. They divide childhood into three segments--birth to age 5, age 5 to age 10, and age 11 to age 15--and find that poverty during the earliest phase of childhood has the greatest negative consequences for the child's ultimate educational attainment. These results suggest that the incentives arising from time limits may have adverse consequences for children, particularly in light of findings that welfare families "neutralize" government policy to only a limited extent (Schoeni 1997; Rosenzweig and Wolpin 1994; Lundberg, Pollak, and Wales 1997).

Of course, other analysts have argued that welfare receipt is damaging for children because it encourages them to bear their own children at early ages and rely on welfare themselves when they grow up (Murray 1984). Indeed, since welfare spells that begin when children are young last longer, on average (O'Neill et al., 1987), preventing such spells might be the most effective way to reduce welfare utilization overall. This line of reasoning would lead one to conclude that time limits may have favorable consequences for children.

Our objective in pursuing this discussion is not to take a stand on how time limits will affect child welfare. Rather, it is to point out that, although the conclusions of these arguments differ widely, they both suggest that the rising reservation wage result may have important substantive implications for the well-being of children. Therefore it is important to test the theory and determine whether it adequately describes the data.

Before turning to a discussion of our data, however, it is useful to mention how potentially endogenous fertility might affect our results. Additional births increase the consumer's period of eligibility by increasing the number of years before the youngest child in the family turns 18. They have no effect on the family's time limit, however. Thus additional births increase the family's reluctance to use welfare, all else equal.

Florida's Family Transition Program: Background, Data, and Program Treatments

Our empirical work focuses on the exit hypothesis, which predicts that families with younger (youngest) children will spend less time on aid than families with older (youngest) children when time limits are initially imposed. Our data come from Florida's Family Transition Program (FTP). Although these data have some limitations, they have two advantages which may help to offset their shortcomings. First, they provide some of the first individual-level data that permit one to study the effects of time limits. Although other analysts have studied time limits that were imposed under the states' pre-PRWORA waivers, most of them use state-level aggregate data.[7] Moffitt (1999) uses data from the CPS, but he does not analyze the effects of time limits separately from the effects of other pre-PRWORA welfare reform provisions. The second advantageous feature of FTP is

[7] This includes CEA (1997), Blank (1997), Ziliak, et al. (1997), and Wallace and Blank (1999).

that it involved random assignment to treatment, which aids in identifying the effect of time limits.

Background

Florida's experimental FTP program was implemented in Escambia County (Pensacola) starting in May 1994. New applicants for cash welfare were randomized into one of two groups: the experimental FTP group, which was subject to time limits, or the control group, which was enrolled in the AFDC program. Persons already receiving welfare were randomized into treatment or control groups at the time of their biannual recertification interviews. Recruitment into the experiment continued until October 1996. Bloom et al. (1997, 1998, 1999) provide evaluations of the early impacts of the program and more detailed information about the data, program operations, and program rules.

Data

In this paper we analyze data on persons drawn into the experiment between May 1994 and February 1995. This yields a sample of 2,623 persons, 1,312 of whom participated in the FTP program and 1,311 of whom received AFDC. Sample members are observed for 24 months beginning with the month after random assignment. In this study we include only single-parent cases, the vast majority of which are headed by women.

The data analyzed here come primarily from two sources, Florida administrative records and a short survey instrument known as the Background Information Form (BIF). Florida administrative records provide data on monthly benefit receipt. The BIF was administered to welfare applicants (and re-certificants) at the time they applied (or were re-certified) for benefits, prior to random assignment. It contains baseline demographic characteristics.

FTP program treatments and conditions of AFDC receipt

Although time limits were arguably the central component of the FTP program, the program included a number of other treatments as well. Since these other treatments may affect recipients' decisions about exiting welfare, they may hinder our ability to isolate the effects of time limits. We describe the full set of FTP treatments in some detail, as well as the corresponding conditions imposed on AFDC recipients under the Florida welfare reform waiver. This discussion motivates our approach to estimating the effect of time limits.

FTP participants were subject to five different categories of treatments: the time limit, financial work incentives, enhanced support services, employment and training mandates, and child care subsidies. Most FTP participants faced a 24-month time limit, after which their benefits could be, and nearly always were, terminated. Particularly disadvantaged participants, who make up roughly 40 percent of the sample, received a 36-month time limit. FTP participants also enjoyed relatively generous work incentives: the first $200 of monthly earnings were disregarded from income in determining their monthly benefits, and earnings in excess of $200 were subject to a benefit reduction rate of 50 percent. Relative to AFDC recipients, FTP participants had enhanced social support services, including enhanced employment and training services. The FTP group was subject to a 30 hour/week employment and training mandate, with exemptions only

for women with children under six months old. Finally, the FTP program provided subsidies for child care.

In some respects, conditions facing the AFDC group were the same as those facing participants in a conventional AFDC program. Most importantly for our purposes, AFDC recipients faced no time limits. The AFDC group also faced conventional AFDC work incentives. During the first four months on aid, the AFDC group faced an income disregard of $120 and a benefit reduction rate of 67 percent After the first four months, the benefit reduction rate rose to 100 percent, and after the 12th month on aid, the disregard fell to $90. Thus AFDC participants had lower income than FTP participants for any given level of earnings.

In other respects, however, the AFDC group faced conditions quite unlike those of a conventional AFDC program. As part of Florida's welfare reform waiver, AFDC participants were subject to the same 30 hours/week employment and training mandate as the FTP group. Unlike the FTP group, however, AFDC mothers were exempted from the mandate if they had children younger than three. Furthermore, AFDC recipients were in principle eligible for the same child care subsidies as FTP participants. Table 1 summarizes the various FTP treatments and the corresponding conditions facing AFDC recipients.

Identification

On the one hand, the presence of multiple program treatments complicates our effort to estimate the effects of time limits. On the other hand, since some of these treatments, such as greater financial work incentives, are components of many states' TANF plans (Moffitt 1998), it may be useful to learn how those treatments affect behavior as well.

Given the number of treatments, however, one might expect that the effects of time limits could be identified only under a number of assumptions. Therefore we approach the identification problem as an effort to establish conditions under which the effects of the time limit are identified. We argue that some of these conditions are likely to be satisfied either by the rules of the FTP program or as a result of the theory developed above. Other conditions, however, are more specialized, and as we stress below, the estimates we provide are valid only under those conditions.

In section V, we present estimates of the effects of time limits and some of the other program treatments under the assumptions we develop here. We then attempt to test some of these assumptions in section VI, acknowledging that the tests are fairly indirect. Our goal with this approach is not to assert that we can identify the effects of time limits under general conditions, but rather to use what are among the first available data on time limits to learn what we can about this new and important feature of the U.S. welfare system. A more robust analysis will have to await the collection of data more specifically designed to isolate the effect of time limits.

The first step of our identification analysis is to determine what is identified by random assignment to the FTP program. Since the program consists of several different treatments, randomization by itself does not identify the effects of any specific treatment. Rather, randomization identifies what we term the general program effect: the effect of participating in FTP rather than AFDC. Because the effects of the specific treatments may vary with the age of the youngest child in the family, the general program effect may vary by age as well. Let g_a denote the general program effect for a mother whose youngest child is *a* years old (or alternatively, is in age group *a*).

Assuming the general program effects to be the sum of the specific effects of time limits (TL), financial work incentives (WI), enhanced support services (SS), employment and training mandates (ET), and child care subsidies (CC), we have

$$g_a = t_a^{TL} + t_a^{WI} + t_a^{SS} + t_a^{ET} + t_a^{CC}$$

where t_a^j denotes the effect of the *j*th specific treatment on the *a*th age group. In general, there are five specific treatment effects to be identified from each general program effect. Obviously, the effects of time limits can be isolated from the general program effects only under a number of identifying restrictions.

One type of restriction that may aid in identification is age-invariance of a particular treatment effect. If the effect of the *j*th treatment does not vary by the age of the youngest child in the family, then $t_a^j = t^j$ for all *a*. A special case of age-invariance arises when a specific treatment does not vary between the experimental and control groups. If the *j*th treatment does not vary between the AFDC and FTP groups, then $t_a^j = 0$ for all *a*. Some restrictions of this type derive from program rules.

For example, all aid recipients from both the FTP and the AFDC groups, with the exceptions of AFDC mothers with children under three and FTP mothers with infants

under six months old, were subject to the same employment and training mandates. Thus $t_a^{ET} = 0$ for $a \geq 3$ (and for $a \leq 6$ months). Similarly, both the AFDC mothers and the FTP mothers had access to the same child care subsidies, at least in principle. Under this condition, $t_a^{CC} = 0$ for all *a*.

Another important restriction derives from the theory developed in section II. Our Lemma implies that, for a consumer whose benefit stock exceeds her eligibility horizon, a time-limited program is equivalent to an entitlement program. Therefore, for mothers whose youngest children are sufficiently close to their 18th birthdays, time limits will have no effect on behavior. Put differently, both FTP mothers and AFDC mothers whose children are above a threshold age will face the same effective time limit, which is the categorical limit that arises from the requirement that there be a minor child in the household. For those with a two-year time limit, the threshold age is 16; for those facing a three-year time limit, the threshold age is 15. Denote this threshold age by $\bar{a}$. Then the Lemma from section II implies that $t_a^{TL} = 0$ for $a \geq \bar{a}$.

Table 2 helps to illustrate the consequences of these various restrictions. Corresponding to our approach to the empirical analysis below, we consider the relationship between the general program effects and the specific treatment effects for families whose youngest children fall into one of four age categories. The youngest category, age group 1, consists of mothers whose children are older than six months (0.5 year) and less than 3 years. This is the age range over which employment and training mandates differ between the FTP and AFDC groups. The oldest category, age group 4, consists of mothers whose youngest children are older than the threshold age $\bar{a}$. Categories 2 and 3 are intermediate age categories, of which at least two are needed to estimate whether time limits lead mothers with younger children to exit the welfare roles sooner than mothers with older children. Age group 2 includes families whose youngest child is 3 to 11 years old; age group 3 includes families whose youngest child is between 12 and 14 or 15, where the upper limit depends on the length of the family's time limit.

Column (3) of Table 2 simply defines the general program effects as the sum of the unrestricted specific treatment effects, reiterating that the specific treatment effects are unidentified without further restrictions. Column (4) imposes the age-invariance restrictions that $t_a^{ET} = 0$ for $a > 1$ and $t_a^{CC} = 0$ for all *a*. Column (5) imposes the restriction from the Lemma that $t_4^{TL} = 0$.

Column (5) reveals that an additional set of conditions sufficient to isolate the effects of time limits, at least on the two intermediate age groups, is that the effects of the financial work incentives and the enhanced social support services, t_a^{WI} and t_a^{SS}, be age-invariant. These assumptions follow neither from the program rules nor, in general, from the theory developed in section II. We do not attempt to argue that these conditions hold generally, but instead provide a set of assumptions under which they hold. In section VI we provide some evidence on the validity of these assumptions.

To analyze the effects of the work incentives we draw on the familiar leisure-consumption diagram long used to study welfare incentives under AFDC. One lesson that follows from the theory developed in section II is that program features that affect the current-period budget constraint may affect current-period utilization decisions by affecting the choice of hours conditional on utilization. This suggests that the current-

period budget diagram may help us to understand how program features such as the benefit reduction rate influence welfare utilization decisions.

Figure 2 shows how the lower benefit reduction rate facing the FTP group affects current-period utilization incentives. Under each program, there is a level of labor supply at which income is the same, whether or not the consumer utilizes welfare. This is known as the "breakeven" level of labor supply. By increasing breakeven hours from h_B^{AFDC} to h_B^{FTP}, and by raising the welfare payment to consumers working less than the breakeven number of hours, the lower benefit reduction rate under FTP should increase welfare utilization by providing a greater incentive to combine work and welfare. The higher earned income disregard under FTP (not illustrated) reinforces this effect.

The real question for us, however, is whether these utilization incentives vary by the age of the youngest child in the family. Presumably, utilization incentives would vary by the age of the youngest child due to interactions between the financial work incentives and child care costs. Child care costs vary inversely with the age of the youngest child. Therefore a mother who had the same wage as that depicted in Figure 2, but who had higher child care costs, would face a lower budget constraint. Thus FTP's financial work incentives might affect her differently than they would affect an otherwise identical mother with lower child care costs. If the child care subsides that were available under FTP were sufficient to equalize child care costs, however, then the two mothers would face the same budget constraint, and be affected similarly by the financial work incentives. Thus an assumption sufficient to justify age-invariance of the effects of the financial work incentives is that the child care subsidies equalized child care costs among mothers with youngest children of different ages.

Ideally, we would test this assumption directly using survey data on net child care expenditures. Unfortunately, no such data exist. Instead, we propose in section VI an indirect test that relies on a simple observation: in order for FTP's financial work incentives to affect the consumer's income, the consumer must combine work and welfare. Thus we test for age-invariant financial work incentives by asking whether FTP affects the consumer's probability of combining work and welfare differently according to the age of her youngest child.

Beyond age-invariance of the effects of the financial work incentives, another condition required to identify the effects of time limits is age-invariance in the effects of the enhanced social supports. To the extent that these services include valuable training opportunities, human capital theory suggests that they would be more valuable to younger mothers than to older mothers. If service utilization were tied to welfare utilization, then this would give the youngest mothers the greatest incentive to remain on the welfare rolls. This would tend to mask the effects of time limits, which provide the youngest mothers with the greatest incentive to leave the welfare rolls. In section VI we test for age-invariance of the effects of enhanced social services by asking whether, within the FTP group, young mothers are more likely than older mothers to remain on aid.

Under the assumptions that the effects of the financial work incentives and the enhanced social supports are age-invariant, some of the age-specific time limit effects are identified, as shown in column (6) of Table 2. Since mothers in the oldest age group are not affected by the time limit, they serve to identify the sum of the effects of the financial work incentives and enhanced support services, $t^{WI} + t^{SS}$. This in turn identifies the

effect of time limits for the two intermediate age groups, since for them, the gross program effect is equal to the sum of the effects of the time limit, the work incentives, and the support services. Thus with these assumptions, we can estimate t_2^{TL} and t_3^{TL} and determine whether $t_2^{TL} < t_3^{TL}$ (< 0) as the theory predicts.

Another potential problem with this strategy is that, prior to 1996, child care subsidies for the AFDC group may have been less generous than those for the FTP group owing to general budgetary shortfalls facing the State of Florida (Bloom et al., 1998). If indeed the subsidies available to the AFDC group were less generous than those available to the FTP group, then our strategy will fail to identify the effects of time limits if the effects of child care subsidies vary with the age of the youngest child. Fortunately, the temporary nature of any shortfalls in child care subsidies allows us to present some evidence as to their effects on our estimates. We return to this issue in section VI.

We turn now to our main results. These provide consistent estimates of the effects of time limits under the full set of identifying assumptions that is spelled out in Table 2. This includes the assumptions based on program rules, the assumption provided by the Lemma in section II, and the age-invariance assumptions on the effects of the financial work incentives and enhanced social supports.

Estimation and Results

Difference-in-difference estimates

The central prediction from the model is that the imposition of time limits reduces time on welfare for mothers of young children by a greater amount than for mothers of older children. The outcome that we study in the initial portion of the analysis is thus the number of months on aid during the first two years after entering the program. In the regression analysis that follows, we also analyze sequences of monthly welfare utilization indicators.

As suggested above, the randomization that was used to assign participants to the FTP and AFDC groups simplifies estimation of the general program effects. We initially compute mean months on aid by age group separately for the FTP participants and the AFDC participants, and use the difference in these means as estimates of g_a. Randomization ensures that these estimates are consistent.

Table 3 presents mean months on aid for AFDC and FTP participants and estimates of g_a for the four age groups defined above.[8] The estimates of the general program effects, in the next-to-last column, show substantial age variation, ranging from -0.30 for the youngest age group to 2.79 for the oldest age group. For the sample as a whole, the general program effect is -0.13. This overall estimate is consistent with the finding that, as of mid-1997, the program had had little effect on aggregate months on aid (Bloom et al. 1998). The small negative effects among the larger, younger age groups

[8] As noted above, the age thresholds for the two oldest groups depend on the length of the family's time limit. Our classification scheme is based on an imputed time limit, which can be constructed for all families. The actual time limit, which was determined by county welfare officials, is available only for the FTP group. When we replace the imputation with the actual time limit for the FTP group, we obtain similar results.

cancel out the large positive effects among the smaller, older age groups, leaving a null effect overall.

Under the identifying assumptions spelled out in section IV, the estimate of g_4 provides an estimate of $t^{WI} + t^{SS}$, as shown in column (6) of Table 2. It suggests that the financial work incentives and enhanced social services raised time on aid by an average of 2.79 months. This is consistent with the discussion above, where we showed that a reduction in the benefit reduction rate should increase welfare utilization. This estimate has a t-statistic of 1.59.

Under the full set of identifying assumptions, we can estimate the specific time limit effects as $\hat{t}_2^{TL} = \hat{g}_2 - \hat{g}_4$ and $\hat{t}_3^{TL} = \hat{g}_3 - \hat{g}_4$. These simple difference-in-difference estimates suggest that time limits reduce time on aid by 2.84 months on average among mothers in the 3-to-11 age group. They reduce time on aid by only 0.40 months on average among mothers in the 12-to-14/15 age group. This pattern is consistent with the prediction from the theory: on average, mothers of younger children spend less time on aid following the imposition of time limits than do mothers of older children. The t-statistic for $\hat{t}_2^{TL}$ is 1.75 and the t-statistic for $\hat{t}_3^{TL}$ is less than one.

Under our identifying assumptions, the difference in the age-specific effects of time limits, $\hat{t}_2^{TL} - \hat{t}_3^{TL}$, is equal to the difference in the corresponding general program effects, $\hat{g}_2 - \hat{g}_3$. The difference between $\hat{g}_2$ and $\hat{g}_3$ is -2.44, with a standard error of 1.36. The t-statistic is thus -1.79. The estimates from this initial estimation exercise follow a pattern that is consistent with the prediction from the theory, and the difference in the effects of the time limit by age group is marginally significant.

Regression estimates

In this section we use regression methods to estimate the effect of time limits on welfare receipt under the full set of identifying assumptions. The regression approach has two advantages over the simple difference-in-difference strategy. First, with a regression we can control for personal characteristics of the experiment participants, reducing unobservable error that contributes to the imprecision of the estimates. Second, many of the specific treatments under FTP are actually time-varying. For example, AFDC mothers who begin the sample period with children under three, who therefore are exempt initially from the employment and training mandates of the program, become subject to those mandates once their children turn three. Likewise, children of FTP mothers who are less than 16 at the beginning of the experiment may eventually reach the age at which the FTP time limit is no longer binding on their mothers. Mothers whose youngest children turn 18 during the sample period become ineligible for aid; as a result, we drop them from the sample starting with the month of their youngest child's 18th birthday.

We analyze monthly welfare utilization sequences, so the unit of observation is the person-month. The model is given by:

$$y_{it} = a_0 + a_1 A_{1it} + a_2 A_{2it} + a_3 A_{3it} + t_1 A_{1it} E_i + t_2^{TL} A_{2it} E_i + t_3^{TL} A_{3it} E_i + t_4 E_i + X_{it} b + m_i + e_{it}$$

for $i = 1,..., n$ and $t = 1,..., 24$, where n is the number of persons in the sample. The dependent variable y_{it} equals one if family i utilized welfare in period t, and equals zero

otherwise. The age group dummies are defined as $A_{jit} = 1$ if the youngest child in the ith family falls into age group j at time t and $A_{jit} = 0$ otherwise, for $j = 1,..., 4$. The variable E_i is the FTP, or experimental, dummy; $E_i = 1$ if family i is in the FTP group and $E_i = 0$ if family i is in the AFDC group. The vector X_{it} includes a number of exogenous regressors including the mother's age at time t; a dummy equal to one if she is black and zero otherwise; a dummy equal to one if she had a three-year time limit and equal to zero if she had a two-year limit; the number of children in the family; the mother's years of schooling; the number of months in which the family had utilized welfare during the 24 months prior to random assignment; the number of quarters of employment during the year preceding random assignment; and a vector of year dummies. Time is measured in months from the date of random assignment, so $t = 0$ is the month of entry into the program. The a, b, and t terms are parameters to be estimated from the data. Under the full set of identifying assumptions, this parameterization of the regression model estimates the effects of time limits for age groups two and three, t_2^{TL} and t_3^{TL}, as the coefficients on the interactions between the corresponding age group dummies and the FTP dummy.

The error term consists of two components, m_i and e_{it}. The former is a family effect, which gives rise to dependence among the monthly observations from a particular family. Since this groupwise dependence may cause conventional OLS standard errors to be biased downwards, we employ a Huber-White covariance matrix estimator that corrects for such dependence. It also accounts for heteroskedasticity, which is an inherent property of the linear regression estimator when applied to a binary dependent variable.

Results of this specification are reported in Table 4. Before proceeding to the estimates of the effects of time limits, it is useful to consider the coefficients of the other regressors in the model. The second row of the table shows that recipients who received 36-month time limits have utilization rates that are 4.7 percentage points higher on average than recipients with the shorter time limit. Black women have higher utilization rates as well, which is consistent with previous work showing that African-American women tend to have longer spells of welfare use than others (O'Neill, Bassi, and Wolf 1987). Neither the mother's age nor the number of children she has have much effect on welfare utilization, but education, past welfare use, and past employment are strongly related to current welfare utilization. The next three rows show a strong trend toward declining welfare use over time. Also worth noting is the high level of significance of most of the coefficients discussed so far, which would lead one to expect that the effects of time limits would be more precisely estimated by this regression model than by the simple difference-in-difference approach.

The coefficients on the age-group dummies show patterns of age dependence within the AFDC sample. The base group is mothers whose youngest children are in the oldest age group, so the coefficients on the age group dummies measure average utilization rates relative to that group. Women with the youngest children are 9.4 percentage points more likely to receive aid on average. Women with youngest children in the 3-to-11 age group are slightly less likely to utilize welfare. Mothers whose youngest children are in the early teen age group are somewhat less likely to receive aid.

Our primary interest centers on the coefficients of the interactions between the age-group dummies and the FTP dummy. Since we include no interaction between the

oldest age-group dummy and the FTP dummy, age patterns within the FTP group are estimated relative to the oldest group. This means that the coefficient on the FTP dummy provides an estimate of $t^{WI} + t^{SS}$; the coefficient on the interaction between the youngest age-group dummy and the FTP dummy provides an estimate of $t_1^{TL} + t_1^{ET}$; and the interactions between the FTP dummy and the second and third age-group dummies provide estimates of t_2^{TL} and t_3^{TL}, respectively.

The coefficient on the FTP dummy is positive, indicating that, under our identifying assumptions, the combined effects of the financial work incentives and enhanced social services act to raise welfare utilization. As noted above, this accords with the prediction that higher disregards and lower tax rates raise welfare use by providing greater incentives to combine welfare with work.

The coefficient on the interaction between the FTP dummy and the youngest age-group dummy is negative and marginally significant. It shows that the combined effect of time limits and employment and training mandates reduces aid utilization by 9.5 percentage points relative to families in the same age group who do not face these constraints. Relative to the mean utilization rate of 51.5 percent, this amounts to a reduction of 18 percent.[9]

The next coefficient, on the interaction between the FTP dummy and the second age-group dummy, provides an estimate of the effect of time limits on families whose youngest children are 3 to 11 years old. This effect is negative, as predicted by the theory, and significant. Subject to the validity of our identifying assumptions, it shows that time limits reduce monthly aid utilization by 12.1 percentage points on average within this age group. This is a sizeable effect. Relative to the mean utilization rate, it represents a 23 percent reduction.[10]

Under our identifying assumptions, the effect of time limits on families with children between 12 and 14 or 15 is estimated by the interaction between the FTP dummy and the third age-group dummy. This coefficient is negative but smaller in absolute value than the coefficient for the 3-to-11 age group, as predicted by the theory. It indicates that time limits reduce aid utilization by 3.4 percentage points among families with youngest children in this age group.

The exit hypothesis predicts that $t_2^{TL} < t_3^{TL} < 0$. The difference $\hat{t}_2^{TL} - \hat{t}_3^{TL}$ equals -0.087 with a standard error of 0.040, yielding a t-statistic of -2.18. The regression framework has improved the power of this test; it is now significant at the 5 percent level. The data thus provide evidence that is supportive of the theory.

This result is fairly robust to alternative specifications of the age relationship. As one alternative, we split the 3-to-11 age group in two. The coefficient on the interaction between the FTP dummy and the dummy for the resulting 3-to-5 age group was nearly

---

9 Summary statistics for all variables included in the regression model appear in Appendix Table 1.

10 This coefficient is larger in absolute value than the preceding coefficient, whereas the theory predicts that the time limit by itself should have a greater negative effect, the younger the youngest child in the family. Of course, the coefficient on the interaction between the FTP dummy and the youngest age-group dummy estimates not just the effect of the time limit, but the sum of the effect of the time limit and the employment and training mandate imposed on this group. The observed pattern in the results would be consistent with the theory if the employment and training mandates were satisfied largely by participation in unpaid training activities, which presumably would lead to increased welfare utilization.

identical to that for the interaction between the FTP dummy and the dummy for the 6-to-11 age group. Although the evidence in favor of the model would have been stronger if the former coefficient had been more negative than the latter, both coefficients were significantly different from the coefficient on the interaction between the FTP dummy and the dummy for the 12 to 14/15 age group. Moreover, when we allowed the effect of the youngest child's age to be linear between the ages of 3 and 14/15, the coefficient on the interaction between the FTP dummy and age was positive (as predicted by the theory) and significant.

Subject to the validity of our identifying assumptions, our estimates suggest that time limits substantially reduce welfare utilization among families with young children. If we weight the estimated age-specific proportionate reductions in welfare utilization by the fraction of cases in each age group, taking the coefficient on the interaction between the FTP dummy and the youngest age group dummy as an estimate of the effect of the time limit on the youngest age group, we calculate that time limits by themselves would have led to a 19 percent reduction in welfare utilization among the FTP group. This is particularly striking because our sample period extends only over the first 24 months after the time limit was imposed. This suggests that time limits had a negative effect on welfare utilization, at least among families with younger children, well before any of the families in our sample could have actually exhausted their benefits. In other words, time limits appear to induce families to leave the welfare rolls well before they actually exhaust their benefits.

An issue of particular interest is whether these apparent responses to time limits are concentrated among those recipients who are most able to find and keep a job. A frequently voiced concern is that welfare reform will lead only the most able recipients to leave the welfare rolls, with the result that only the most disadvantaged will remain on aid long enough to actually exhaust their benefits (Duncan, Harris, and Boisjoly 1997; Moffitt 1998; Sawhill and Zedlewski 1995). To address this question, we divide the sample according to the length of the time limit received by the mother. Whereas two-year time limits were the default, women with particularly low levels of education, high levels of past welfare utilization, or low levels of past employment were assigned a three-year limit. Since these factors are known to predict lengthy spells on aid (Pavetti 1995), the length of the time limit provides a useful summary measure of relative disadvantage within the welfare population.

Regression results by the length of the family's time limit are presented in Table 5. These regressions include all of the variables reported in Table 4, but we report only the coefficients on the FTP-age group interactions in order to save space. The coefficients are fairly similar between the two groups. In both cases, $\hat{t}_2^{TL} < \hat{t}_3^{TL} < 0$. Although dividing the sample has rendered most of the coefficients insignificant, the estimates suggest that responses to time limits are not concentrated among the most work-ready recipients. The relatively disadvantaged group appears to be leaving the welfare rolls at about the same rate as the relatively more advantaged group.

Evidence on the Identifying Assumptions

As we have stressed above, our estimates rely on a number of identifying assumptions. In this section we provide evidence on the validity of some of those

restrictions. The first two involve the age-invariance of the effects of the financial work incentives and the age-invariance of the effects of the social support services. There is also a question of whether access to child care subsidies varied between the control and treatment groups, contrary to the program design. We present some admittedly indirect evidence on each of these assumptions in turn.

Age-invariance of the effects of the financial work incentives

We first consider whether the effects of the financial work incentives are invariant with respect to the age of the youngest child in the family. Our evidence stems from a simple observation: in order for FTP's financial work incentives to affect the consumer's income, the consumer must combine work and welfare. If, alternatively, she receives welfare but does not work, then she simply receives the maximum welfare benefit, which is the same under both AFDC and FTP and is labeled as *G* in Figure 2. If she works but does not receive welfare (i.e., works more than h_B^{FTP} hours in Figure 2), then again she receives the same income under both AFDC and FTP. Either way, her behavior is not affected by the FTP work incentives. Only if she combines welfare and work does she benefit from FTP's more generous income disregards. Put differently, if the financial work incentives affect welfare use, then they must do so by affecting probability of combining welfare and work. Thus we test whether FTP affected the probability of working and receiving welfare in a manner that varies by the age of the youngest child.

Table 6 presents estimates from a regression in which the dependent variable is a dummy equal to one if the sample member utilized welfare during a quarter in which she was employed. Ideally, we would like to know whether she truly combined work and welfare, that is, whether she worked and received welfare in the same month. Because Florida's Unemployment Insurance (UI) system does not reveal when during the quarter the person worked, however, we cannot determine precisely whether parents were receiving welfare at the same time they were working. Because we have quarterly employment data, the unit of observation in Table 6 is the person-quarter, rather than the person-month as in all of the other tables. The results reported in Table 6 are from a regression that included all the variables shown in Table 4. Only the coefficients of the FTP dummy and the FTP-age group interactions are shown in order to save space.

Although the coefficient on the FTP dummy is only marginally significant, it is positive as predicted by the theory, suggesting that the financial work incentives did increase welfare utilization by increasing the probability that the typical FTP mother combined work and welfare. More important for testing the age-invariance hypothesis, however, are the interactions between the FTP dummy and the age-group dummies. The coefficients on these terms vary as to their sign and are exceeded by their standard errors. They are jointly insignificant as well, with an F-statistic of only 1.03. These estimates suggest that the financial work incentives increase the probability of combining work and welfare in a manner that is essentially uniform across the child-age distribution. Thus these results lend some support to the assumption that the effects of FTP's financial incentives are age-invariant.

Age-invariance of the effects of the enhanced social services

Regarding the effects of the enhanced social services, we noted that the most likely alternative to age-invariance is that younger mothers would find the services more

valuable than others, and thus be more likely to remain on the welfare rolls in order to take advantage of them. If so, then an interaction between the mother's age and the FTP dummy should enter the model significantly and with a negative sign.

Results are shown in column (1) of table 7, where the dependent variable is once again the monthly welfare utilization dummy. With the exception of the added variables, the specification used to estimate these models is exactly the same as that reported in Table 4. In table 7 we save space by reporting only the newly added variables and the interactions between the age group dummies and the FTP dummy.

The coefficient on the interaction is indeed negative but it is no larger than its standard error. In the second column, we include an interaction between the FTP dummy and the mother's education on the grounds that the attractiveness of the enhanced training services might vary by the mother's education level. In the third column, we include interactions between the FTP dummy and variables measuring the extent of the mother's employment over the year prior to random assignment and welfare receipt during the 24 months preceding random assignment. We include these variables as proxies for work-readiness, on the grounds that the attractiveness of the enhanced training services might vary according to the mother's general employability. None of these interaction terms is significant. More importantly, although the coefficient on the FTP dummy is sensitive to the inclusion of these interactions, the key coefficients, $\hat{t}_2^{TL}$ and $\hat{t}_3^{TL}$, are essentially the same across the different specifications. Thus none of these tests yields evidence contrary to our assumption that the effects of the enhanced social services are age-invariant.

There are two possible explanations for this result. The first has to do with the manner in which aid recipients were assigned to training activities. These assignments were made by caseworkers rather than being chosen by the recipient. Even with the possibility of negotiation between the two parties, aid recipients may not always have received their desired assignment. The second is that the enhanced training services may not have been perceived by the FTP mothers as offering valuable opportunities to gain skills. In this case, there would be no reason for younger, better educated, or more employable mothers to take greater advantage of these services.

More direct evidence on use of social services comes from MDRC's Two-Year Client Survey (TYCS), which was administered to a sample of about 600 mothers drawn equally from the AFDC and FTP groups in early- to mid-1997. This survey included questions about the use of education, training, and employment services by respondents in the two years following random assignment. If enhanced social services provide a greater opportunity to mothers with younger children, then FTP should have induced the largest increase in use of these services for mothers with younger children. Results of the survey indicate that it did not. Families did increase their use of services, but this increase does not appear to be related to age of youngest child. For families with children between 3 and 11 years of age, about 51 percent of the control group used such services compared to about 73 percent of the program group, an increase of 22 percentage points. However, the increase for other families was quite similar – about 30 percentage points

for families with children 12 to 15 years old and about 28 percentage points for families with children 16 to 18 years old.[11]

This is consistent with evidence from California's GAIN program, which also provided welfare mothers with enhanced employment and training services. As in FTP, mothers believed to have few job skills were encouraged to increase their human capital through adult education, while more job-ready mothers were required to enroll in programs designed to help them apply for and obtain jobs. Also like FTP, GAIN was evaluated by MDRC using a random assignment evaluation. According to data from that evaluation, GAIN reduced welfare use by about the same amount for mothers with younger children as for mothers with older children. In the first 3 years after random assignment, mothers whose youngest child was 6 to 11 years old reduced their welfare use by 0.80 quarters (out of a potential 12 calendar quarters over 3 years) and mothers whose youngest child was older than age 11 reduced their welfare use by 0.68 quarters. Although families with children younger than 6 were not required to participate in GAIN, some welfare recipients with young children did volunteer for GAIN. Among mothers with children between 4 and 5 years old, GAIN reduced welfare use by a similar amount, 0.79 calendar quarters.[12] While these reductions in welfare use are highly statistically significant, differences across age groups are not.

Equal access to child care subsidies between FTP and AFDC groups

The potential problem regarding access to child care subsidies arises due to funding shortages which may have caused cutbacks in the subsidies available to AFDC mothers. Our strategy to investigate this problem makes use of the fact that any such cutbacks had been restored by the beginning of 1996. Those cutbacks that took place arose due to general budgetary limitations, and those limitations had ended by the end of 1995. Thus any differences in child care subsidies between the AFDC and FTP groups existed only in 1994 and 1995. After 1995, this particular treatment did not vary between the experimental and control groups.

Thus we ask whether AFDC mothers had higher utilization rates relative to FTP mothers in 1994 and 1995 than in 1996 and 1997. It would be damaging to our identification strategy if increased welfare usage within the AFDC group during this time showed substantial negative age dependence. To address this issue, we first add to the regression an interaction between a pre-1996 dummy, equal to one in 1994 and 1995 and equal to zero in 1996 and 1997, and a control group dummy, equal to one if $E_i = 0$ and equal to zero if $E_i = 1$. We then interact this term with the age group dummies to allow for age-dependence.

The first column of table 8 reports the coefficient on the interaction between the pre-1996 dummy and the control group dummy. It indicates that, contrary to expectation, AFDC mothers actually were less likely to utilize welfare, relative to FTP mothers, in the early part of the sample period. Results in the second column show no significant age pattern. Moreover, the coefficients of the age group/experimental dummy interactions

---

[11] In addition to use of any activity, the survey contains information on individual activities such as enrollment in adult basic education and attendance at job club. For no activity were the impacts of FTP statistically significant different by age of youngest child.

[12] For more information on the GAIN program and its evaluation, see Riccio et al, 1994. Tabulations in this paragraph were made by the authors using data made available by MDRC.

are largely unchanged by the addition of these variables. To the extent that there were differential care subsidies between the AFDC and FTP groups during this time, they did not lead to greater welfare use on the part of AFDC mothers, nor did the effects of any shortfalls exhibit an age-dependent pattern. Thus the evidence against our assumption that child care subsidies were effectively equal between groups is fairly weak.

Further evidence on this point comes from the TYCS, which was administered to a sample of mothers drawn equally from the AFDC and FTP groups in early- to mid-1997.[13] Although this survey asked no specific questions about child care utilization, it did ask a number of questions pertaining to links between child care and work. The strongest evidence from the TYCS suggesting that there were important differences in child care subsidies comes from questions posed to women who were neither working nor looking for work at the time of the TYCS. Among FTP mothers, 7 percent indicated that an inability to find or afford child care was the reason for their (non-)employment status. Among AFDC mothers, the corresponding proportion was 17 percent. Since only 20 percent of the FTP mothers, and 27 percent of the AFDC mothers, were neither working nor seeking employment, however, this means that only 1.5 percent of the FTP group, and 4.5 percent of the AFDC group, were neither working nor looking for work primarily due to child care problems.

Other data from the survey, however, are even less suggestive of important differences in the level of child care support between groups. One question asked people who worked less than 30 hours per week why they were not working full time. Five percent of the AFDC group cited child care as the reason. In the FTP group, the corresponding fraction was four percent. Another question focused on women who had worked at some point since program intake, but were not working at the time of the TYCS because they had quit their job. AFDC and FTP mothers were equally likely to say that they had quit for reasons related to child care.

In view of the prominence that was given to discussions of child care shortfalls in the original FTP evaluation (Bloom et al, 1998), it is worthwhile to discuss how these shortfalls could have had such seemingly little effect on our results. We offer two possible explanations. First, as noted in the original evaluation report, it took some time for the FTP program to start running at full efficiency. Thus start-up issues may have affected the delivery of child care services to FTP participants at the same time that budgetary shortfalls affected service delivery to AFDC recipients.

Second, many of the recipients may have had access to free child care, in which case they would have been unaffected by the availability of subsidies. Although there are no data on child care utilization for the FTP program, the evaluators of an earlier Florida welfare reform experiment found that 55 percent of aid recipients had access to free care from friends or relatives, and that only 15 percent utilized any child care subsidies at all (Kemple and Haimson 1994). If participants in the FTP program were similarly situated, then one might expect differences in child care subsidies between the AFDC and FTP groups to have little effect on welfare receipt.

Conclusions

[13] This paragraph and the next borrow heavily from Bloom, et al. (1997, pp. 57-59)

One of the important contributions of this paper is the theoretical model. Essentially, we have taken the static model of welfare utilization, which long served as the basis for analyzing welfare incentives under AFDC, and generalized it to incorporate the effects of time limits. The solution to this model reveals that the consumer is most reluctant to utilize welfare at the beginning of her period of eligibility, but that she becomes less reluctant as time passes. This has a number of implications for the age distribution of children leaving and entering the welfare rolls. First, upon the imposition of time limits, families with younger (youngest) children should be more likely to exit welfare than families with older (youngest) children. Second, families with younger (youngest) children should be less likely to enter the welfare rolls than families with older (youngest) children. Together, this implies that the reduction in welfare expenditures that results from the imposition of time limits will fall disproportionately on families while their children are young. For a number of reasons, this result may have substantive implications for the well-being of poor children.

Using data from Florida's Family Transition Program, which imposed time limits in 1994 under welfare reform waivers, we attempt to test the first of these empirical implications. The data do accord with the theory, in that families with younger (youngest) children do spend less time on welfare after the imposition of time limits than families with older (youngest) children. Indeed, our estimates suggest that time limits by themselves resulted in a substantial decrease in welfare utilization that was concentrated among families with young children.

A number of assumptions were necessary to isolate the effects of time limits from other features of FTP, however. Tests of those assumptions generally failed to reject, but the tests were indirect. As a result, our empirical results should be taken as suggestive rather than definitive. An important agenda item for future research is to better isolate the effect of this new and important feature of the U.S. welfare system.

References

Blank, Rebecca M (1997). "What Causes Public Assistance Caseloads to Grow?" National Bureau of Economic Research working paper.

Bloom, Dan, Mary Farrell, and James J. Kemple. The Family Transition Program: Implementation and Interim Impacts of Florida's Initial Time-Limited Welfare Program. New York: Manpower Demonstration Research Corporation, 1997.

Bloom, Dan, Mary Farrell, and James J. Kemple. The Family Transition Program: Implementation and Interim Impacts of Florida's Initial Time-Limited Welfare Program. New York: Manpower Demonstration Research Corporation, 1998.

Council of Economic Advisers (1997). Technical Report: Explaining the Decline in Welfare Receipt, 1993-1996. Washington, DC: Council of Economic Advisers.

Duncan, Greg J., Kathleen Mullan Harris, and Johanne Boisjoly (1997). " Time Limits and Welfare Reform New Estimates of the Number and Characteristics of Affected Families." Mimeo, Northwestern University.

Duncan, Greg J., W. Jean Yeung, Jeanne Brooks-Gunn, and Judith R. Smith. "The Effects of Childhood Poverty on the Life Chance of Children." American Sociological Review 63, June 1998, 406-423.

Edin, Kathryn and Laura Lein (1997). Making Ends Meet. New York: Russell Sage Foundation.

Guo, Guang (1998). "The Timing of the Influences of Cumulative Poverty on Children's Cognitive Ability and Achievement." Social Forces 77: 257-288.

Kornfeld, Robert and Howard S. Bloom. 1999. "Measuring Program Impacts on Earnings and Employment: Do Unemployment Insurance Wage Reports from Employers Agree with Surveys of Individuals?" Journal of Labor Economics. 17(1): 168-197.

Kemple, James J. and Joshua Haimson. Florida's Project Independence: Program Implementation, Participation Patterns, and First-Year Impacts. New York: Manpower Demonstration Research Corporation, January 1994.

Loeb, Susanna, and Mary Corcoran. "Welfare, Work Experience, and Economic Self-Sufficiency." Mimeo, February 1999.

Lundberg, Shelly J., Robert A. Pollak, and Terence J. Wales (1997). "Do Husbands and Wives Pool their Resources?" Journal of Human Resources 32, 463-480.

Martini, Alberto and Michael Wiseman (1997). "Explaining the Recent Decline in Welfare Caseloads: Is the Council of Economic Advisers Right?" http://www.urban.org/welfare/cea.htm.

Moffitt, Robert (1983). "An Economic Model of Welfare Stigma." American Economic Review 73, 1023-1035.

Moffitt, Robert. "Incentive Effects of the U.S. Welfare System: A Review." Journal of Economic Literature 30, March 1992, 1-61.

_____. "Time Limits." Mimeo, Johns Hopkins University, November 1998.

_____. "The Effect of Pre-PRWORA Waivers on AFDC Caseloads and Female Earnings, Income, and Labor Force Behavior." Working paper, Joint Center for Poverty Research, May 1999.

Murray, Charles. Losing Ground. New York: Basic Books, 1984.

O'Neill, June, Lauri Bassi, and Douglas A Wolf. "The Duration of Welfare Spells." Review of Economics and Statistics 69, May 1987, 241-248.

Pavetti, LaDonna. "Who is Affected by Time Limits?" In Sawhill, Isabel, ed., Welfare Reform: An Analysis of the Issues. Washington, DC: Urban Institute Press, 1995.

Riccio, James, Daniel Friedlander, and Stephen Freedman. 1994.
GAIN: Benefits, Costs, and Three-Year Impacts of a Welfare-to-Work Program. Manpower Demonstration Research Corporation. New York.

Rosenzweig, Mark R. and Kenneth I Wolpin. "Parental and Public Transfers to Young Women and Their Children." American Economic Review 84, December 1994, 1195-1212.

Sawhill, Isabel and Sheila Zedlewski. "Assessing the Personal Responsibility Act." In Sawhill, Isabel, ed., Welfare Reform: An Analysis of the Issues. Washington, DC: Urban Institute Press, 1995.

Schoeni Robert F.(1997) "Does Aid to Families With Dependent Children Displace Familial Assistance?" Mimeograph, RAND Corporation.

Swann, Christopher A. "Welfare Reform When Agents are Forward-Looking." Mimeo, University of Virginia, December 1998.

Wallace, Geoffrey and Rebecca M. Blank (1999). "What Goes Up Must Come Down? Explaining Recent Changes in Public Assistance Caseloads." Mimeo, Northwestern University.

Ziliak, J.P., Figlio, D.N., Davis, E.E., and Connolly, L.S. "Accounting for the Decline in AFDC Caseloads: Welfare Reform or Economic Growth?" Institute for Research on Poverty Working Paper 1151-97, 1997.

Appendix: Proof of the Rising Reservation Wage Proposition

To begin we adopt some simplifying notation. Let

$$F_t(S_t) \equiv F(\tilde{w}_t(S_t)), \tag{A1}$$

$\overline{F}_t(S_t) = 1 - F_t(S_t)$, and

$$Eu_t(S_t) = F_t(S_t)E[v(w_t(1-t),G) \mid w_t \le \tilde{w}_t(S_t)] + \overline{F}_t(S_t)E[v(w_t,0) \mid w_t > \tilde{w}_t(S_t)]. \tag{A2}$$

Equation (A1) is simply notation for the probability of utilizing welfare in period *t*, given a remaining benefit stock of S_t. Equation (A2) defines expected current-period utility at time *t* as a function of the remaining benefit stock at time *t*. Finally, let

$$E[(u_{t'}(S_{t'}) \mid S_t]$$

denote the expectation of current-period utility at time *t*' as a function of the remaining benefit stock at *t*', given that the consumer had S_t units of benefits at time *t* and has followed the optimal reservation wage strategy in periods *t*, *t*+1, ..., *t*'-1.

Our objective is to establish that

$$EV_{t+1}(S_t) - EV_{t+1}(S_t - 1) > EV_{t+2}(S_t) - EV_{t+2}(S_t - 1),$$

or equivalently, that

$$EV_{t+1}(S_t) - EV_{t+2}(S_t) > EV_{t+1}(S_t - 1) - EV_{t+2}(S_t - 1). \tag{A3}$$

With the definitions above, we can write

$$EV_t(S_t) = \sum_{s=t}^{T} r^{s-t} E[u_s(S_s) \mid S_t].$$

Then since $EV_t(S_t)$ gives the maximized value of the consumer's problem starting from period *t* with benefit stock S_t, we have

$$EV_{t+1}(S_t) \ge EV_{t+2}(S_t) + r^{T-t-1} E[u_T(S_T) \mid S_{t+1} = S_t].$$

Applying similar logic to $EV_{t+2}(S_t - 1)$, we have

$$EV_{t+2}(S_t - 1) \geq EV_{t+1}(S_t - 1) - r^{T-t-1} E[u_T(S_T) \mid S_{t+1} = S_t - 1].$$

Therefore a sufficient condition for (A3) is

$$E[u_T(S_T) \mid S_{t+1} = S_t] > E[u_T(S_T) \mid S_{t+1} = S_t - 1]. \quad \text{(A4)}$$

To evaluate $E[u_T(S_T) \mid S_{t+1} = S_t]$, use the recursion formula and the reservation wage property to write $EV_{t+1}(S_t)$ in terms of $E[u_s(S_s) \mid S_{t+1} = S_t]$ for $s = t+1, t+2,...,T$ and collect terms in T. $E[u_T(S_T) \mid S_{t+1} = S_t - 1]$ can be evaluated the same way. Simplifying the resulting expressions, using the fact that $Eu_T(S_T) = Eu_T(1)$ for all $S_T > 1$ since benefits have no value after period T, shows that (A4) holds if and only if

$$Eu_T(0) + [1\text{-}P(\text{utilizing welfare in } S_t \text{ periods})]\ [Eu_T(1) - Eu_T(0)] >$$

$$Eu_T(0) + P(\text{utilizing welfare in fewer than } S_t \text{ -1 periods})\ [Eu_T(1) - Eu_T(0)]. \quad \text{(A5)}$$

Expected utility in the final period is equal to $Eu_T(0)$, expected utility in the absence of any remaining benefits, plus the product of the probability of having at least one unit of benefits at the beginning of period T and the difference in expected utility from having at least one unit of benefits and having none. The last term is given by $[Eu_T(1) - Eu_T(0)]$. Beginning with S_t units of benefits, the probability of having at least one unit remaining at the beginning of period T is one minus of probability of utilizing benefits in S_t periods. Beginning with S_t - 1 units of benefits, the probability of having at least one unit remaining is equal to the probability of utilizing benefits in fewer than S_t -1 periods. Simplifying and rearranging terms, (A5) can be re-written as

P(utilizing welfare in S_t periods) < 1-P(utilizing welfare in fewer than S_t -1 periods)

or

P(utilizing welfare in S_t periods) < P(utilizing welfare in S_t -1 or more periods), establishing the result.

Figure 1
The Consumer's Reservation Wage for Welfare Utilization

Panel A

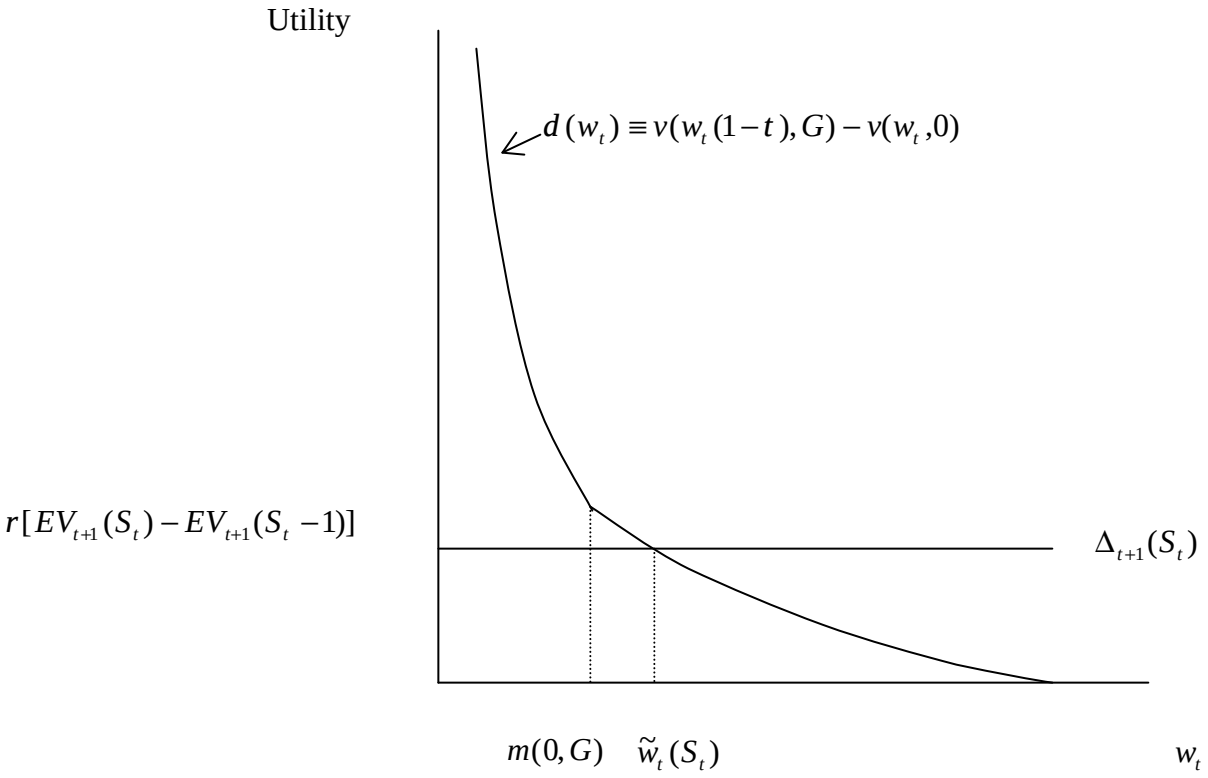

Panel B

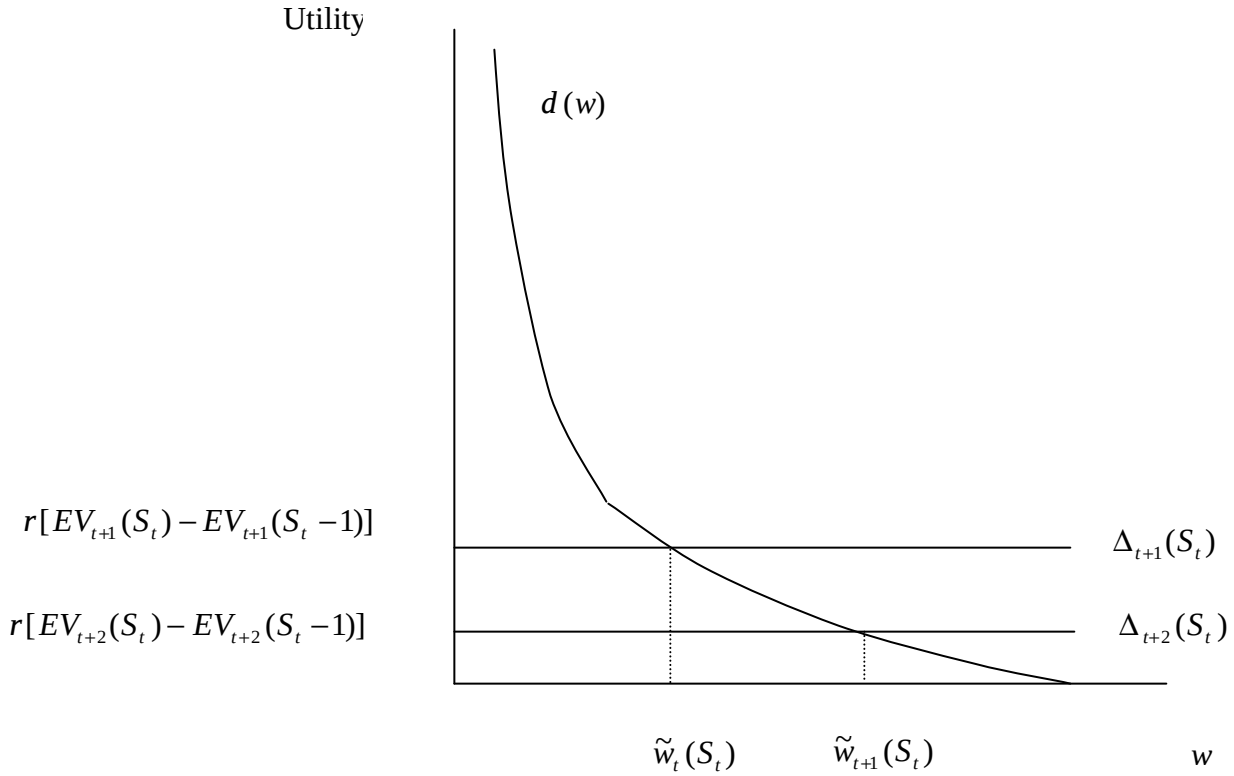

Figure 2
The Effect of Financial Work Incentives on Welfare Utilization

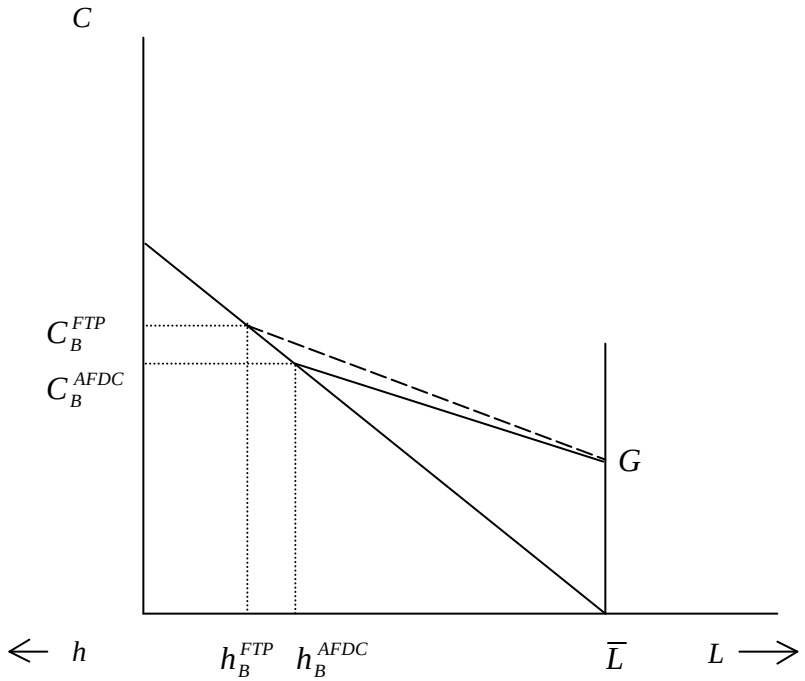

Table 1

Summary of FTP Treatments and Corresponding Conditions of AFDC Receipt

Type of Treatment	FTP Group	AFDC Group
Time limits	24-month time limit for most recipients; 36-month time limit for particularly disadvantaged recipients.	No time limits.
Financial work incentives	$200 earned income disregard and 50 percent benefit reduction rate.	$120 disregard and 67 percent benefit reduction rate for first four months; 100 percent benefit reduction rate after four months; $90 disregard after 12 months.
Support services	Enhanced employment and training services and social support services.	Conventional AFDC services.
Employment and training mandates	30 hours/week either working or in training. Exemptions only for mothers with infants under 6 months of age.	30 hours/week either working or in training. Exemptions for mothers with children under 3 years old.
Child care subsidies	In theory, both groups had access to same child care subsidies. In practice, subsidies for AFDC group may have been less generous prior to 1996.	

Table 2
Identification of Specific Treatment Effects from General Program Effects

Age group (Age of youngest child) (1)	Gross program effect... (2)	Without restrictions, equals... (3)	With restrictions on t_a^{ET} and t_a^{CC} derived from program rules, equals... (4)	Adding the restriction $t_4^{TL} = 0$ from the Lemma in section II, equals... (5)	Adding age-invariance of t_a^{WI} and t_a^{SS}, equals (6)
1 (6 mos. to 2 years)	g_1	$t_1^{TL} + t_1^{WI} + t_1^{SS} + t_1^{ET} + t_1^{CC}$	$t_1^{TL} + t_1^{WI} + t_1^{SS} + t_1^{ET}$	$t_1^{TL} + t_1^{WI} + t_1^{SS} + t_1^{ET}$	$t_1^{TL} + t^{WI} + t^{SS} + t_1^{ET}$
2 (3 to 11)	g_2	$t_2^{TL} + t_2^{WI} + t_2^{SS} + t_2^{ET} + t_2^{CC}$	$t_2^{TL} + t_2^{WI} + t_2^{SS}$	$t_2^{TL} + t_2^{WI} + t_2^{SS}$	$t_2^{TL} + t^{WI} + t^{SS}$
3 (12 to 14/15)	g_3	$t_3^{TL} + t_3^{WI} + t_3^{SS} + t_3^{ET} + t_3^{CC}$	$t_3^{TL} + t_3^{WI} + t_3^{SS}$	$t_3^{TL} + t_3^{WI} + t_3^{SS}$	$t_3^{TL} + t^{WI} + t^{SS}$
4 (15/16 to 17)	g_4	$t_4^{TL} + t_1^{WI} + t_1^{SS} + t_4^{ET} + t_4^{CC}$	$t_4^{TL} + t_4^{WI} + t_4^{SS}$	$t_4^{WI} + t_4^{SS}$	$t^{WI} + t^{SS}$

Table 3
Mean Months on Aid during First 24 Months after Random Assignment, by Age of Youngest Child

	AFDC Group		FTP Group		Gross program effect	Time limit effect
Age group	N	Mean	N	Mean	($\hat{g}_a$)	($\hat{t}_a^{TL}$)
6 months to 2 years	555	12.96	557	12.66	-0.30	
		(0.38)		(0.37)	(0.53)	
3 to 11	612	12.26	626	12.21	-0.05	-2.84
		(0.36)		(0.35)	(0.50)	(1.83)
12 to 14/15	102	9.82	92	12.21	2.39	-0.40
		(0.87)		(0.93)	(1.27)	(2.17)
15/16 to 17	42	7.64	37	10.43	2.79	
		(1.15)		(1.34)	(1.76)	
Total	1,311	12.42	1,312	12.35	-0.13	
		(0.25)		(0.24)	(0.34)	

Table 4
Linear Regression Estimates of the Effects of Time Limits on the Monthly Probability of Receiving Aid

Variable	Coefficient	Standard error
Constant	0.605	0.070
36-month time limit	0.047	0.013
Mother black	0.084	0.014
Mother's age	-0.001	0.001
Number of children	0.001	0.006
Years of schooling	-0.018	0.004
Months of welfare utilization in 24 months prior to random assignment	0.018	0.001
Quarters of employment in year preceding random assignment	-0.024	0.004
1995 dummy	-0.138	0.009
1996 dummy	-0.308	0.011
1997 dummy	-0.407	0.021
Youngest child between 6 months and 2 years	0.094	0.043
Youngest child between 3 and 11	0.082	0.040
Youngest child between 12 and 14/15	0.049	0.043
FTP dummy * Youngest child between 6 months and 2 years	-0.095	0.058
FTP dummy * Youngest child between 3 and 11 ($\hat{t}_2^{TL}$)	-0.121	0.055
FTP dummy * Youngest child between 12 and 14/15 ($\hat{t}_3^{TL}$)	-0.034	0.062
FTP dummy	0.105	0.053
$\hat{t}_2^{TL} - \hat{t}_3^{TL}$	-0.087	0.040
R^2	0.214	
Sample size: 62,229 observations		

Note: Huber-White standard errors in third column account both for groupwise dependence arising from multiple observations per person and for heteroskedasticity.

Table 5
Linear Regression Estimates of the Effects of Time Limits on the Probability of Receiving Aid, by the Length of the Family's Time Limit

	Length of Time Limit	
Variable	24 months	36 months
FTP dummy * Youngest child between 6 months and 2 years	-0.108 (0.076)	-0.097 (0.087)
FTP dummy * Youngest child between 3 and 11 ($\hat{t}_2^{TL}$)	-0.090 (0.072)	-0.162 (0.085)
FTP dummy * Youngest child between 12 and 14/15 ($\hat{t}_3^{TL}$)	-0.023 (0.079)	-0.008 (0.102)
FTP dummy	0.072 (0.069)	0.149 (0.082)
R^2	0.200	0.190
Sample size	35,757	26,472

Note: Huber-White standard errors, in parentheses, account both for groupwise dependence arising from multiple observations per person and for heteroskedasticity. In addition to the variables shown, all regressions include all other variables shown in Table 4.

Table 6
Estimates of the Effects of the FTP Program on the Probability of Combining Work and Welfare, by the Age of the Youngest Child in the Family

Variable	Coefficient (standard error)
FTP dummy * Youngest child between 6 months and 2 years	0.026 (0.042)
FTP dummy * Youngest child between 3 and 11	-0.011 (0.044)
FTP dummy * Youngest child between 12 and 14/15	-0.001 (0.043)
FTP dummy	0.058 (0.038)
R^2	0.072
Sample size	20,665

Note: Huber-White standard errors, in parentheses, account both for groupwise dependence arising from multiple observations per person and for heteroskedasticity. In addition to the variables shown, all regressions include all other variables shown in Table 4.

Table 7
Linear Regression Estimates of the Effects of Time Limits on the Probability of Receiving Aid, with Interactions between the FTP Dummy and Maternal Characteristics

Variable	(1)	(2)	(3)
FTP dummy * Youngest child between 6 months and 2 years	-0.123	-0.098	-0.090
	(0.063)	(0.057)	(0.057)
FTP dummy * Youngest child between 3 and 11 ($\hat{t}_2^{TL}$)	-0.140	-0.123	-0.122
	(0.058)	(0.055)	(0.055)
FTP dummy * Youngest child between 12 and 14/15 ($\hat{t}_3^{TL}$)	-0.040	-0.035	-0.035
	(0.062)	(0.062)	(0.062)
FTP dummy	0.182	0.011	0.071
	(0.094)	(0.094)	(0.056)
Mother's age * FTP dummy	-0.002		
	(0.002)		
Mother's education * FTP dummy		0.009	
		(0.007)	
Quarters of employment in year preceding random assignment * FTP dummy			0.014
			(0.008)
Months of welfare utilization in 24 months preceding random assignment * FTP dummy			0.001
			(0.001)
R^2	0.214	0.214	0.214

Note: Huber-White standard errors, in parentheses, account both for groupwise dependence arising from multiple observations per person and for heteroskedasticity. In addition to the variables shown, all regressions include all other variables shown in Table 4. Sample size is 62,229

Table 8
Further Linear Regression Estimates of the Effects of Time Limits on the Probability of Receiving Aid

Variable	(1)	(2)
FTP dummy * Youngest child between 6 months and 2 years	-0.101	-0.110
	(0.058)	(0.067)
FTP dummy * Youngest child between 3 and 11 ($\hat{t}_2^{TL}$)	-0.120	-0.113
	(0.055)	(0.062)
FTP dummy * Youngest child between 12 and 14/15 ($\hat{t}_3^{TL}$)	-0.034	-0.019
	(0.062)	(0.072)
FTP dummy	0.085	0.080
	(0.053)	(0.060)
Pre-1996 dummy * AFDC dummy * Youngest child between 6 months and 2 years		-0.011
		(0.050)
Pre-1996 dummy * AFDC dummy * Youngest child between 3 and 11		0.013
		(0.045)
Pre-1996 dummy * AFDC dummy * Youngest child between 12 and 14/15		0.026
		(0.059)
Pre-1996 dummy * AFDC dummy	-0.036	-0.045
	(0.014)	(0.045)
R^2	0.214	0.214

Note: Huber-White standard errors, in parentheses, account both for groupwise dependence arising from multiple observations per person and for heteroskedasticity. In addition to the variables shown, all regressions include all other variables shown in Table 4. Sample size is 62,229.

Appendix Table 1
Summary Statistics of Variables Used in Regression Models

Variable	Mean (Standard deviation)
Welfare utilization indicator (dependent variable)	0.515
36-month time limit	0.436
Mother black	0.520
Mother's age	29.5 (7.3)
Number of children	2.1 (1.2)
Years of schooling	11.1 (1.6)
Months of welfare utilization in 24 months preceding random assignment	12.4 (9.6)
Quarters of employment in year preceding random assignment	1.2 (1.5)
1995 dummy	0.487
1996 dummy	0.394
1997 dummy	0.014
Youngest child between 6 months and 2 years	0.262
Youngest child between 3 and 11	0.612
Youngest child between 12 and 14/15	0.085
FTP dummy	0.501
FTP dummy * Youngest child between 6 months and 2 years	0.134
FTP dummy * Youngest child between 3 and 11	0.305
FTP dummy * Youngest child between 12 and 14/15	0.043
Pre-1996 dummy	0.592
AFDC dummy	0.489
Pre-1996 dummy * AFDC dummy	0.295
Pre-1996 dummy * AFDC dummy * Youngest child between 6 months and 2 years	0.092
Pre-1996 dummy * AFDC dummy * Youngest child between 3 and 11	0.165
Pre-1996 dummy * AFDC dummy * Youngest child between 12 and 14/15	0.025
Number of monthly observations	62,229