

# ARE PROGRAM PARTICIPANTS GOOD EVALUATORS?

Jeffrey Smith

Alexander Whalley

Nathaniel T. Wilcox

Version of April 29, 2013

## Abstract

How well do program participants assess program performance ex-post? In this paper we compare participant evaluations based on survey responses to econometric impact estimates obtained using data from the experimental evaluation of the U.S. Job Training Partnership Act. We have two main findings: First, the participant evaluations are unrelated to the econometric impact estimates. Second, the participant evaluations do covary with impact proxies such as service intensity, outcome levels, and before-after outcome differences. Our results suggest that program participants behave as ‘lay scientists’ who seek to estimate the impact of the program but face cognitive challenges in doing so.

JEL Codes: I28, J24, C83

Word count (main text): 10,927 (file) 14,216

We thank the W.E. Upjohn Institute for Employment Research for funding this research. Corresponding author: Jeffrey Smith, Department of Economics, University of Michigan, 238 Lorch Hall, 611 Tappan Street, Ann Arbor MI 48109 USA, 734-764-5359, [econjeff@umich.edu](mailto:econjeff@umich.edu). We are very grateful to Iwan Barankay, Dan Black, Tanya Byker, Sebastian Calónico, Pam Giustinelli, Hilary Hoynes, Guy Michaels, Bob Schoeni, and audience members and discussants at Chicago (CHAS), Concordia, Indiana, LSE, Mannheim, Queen’s, UCL (especially Richard Blundell, Hidehiko Ichimura and Costas Meghir), Warwick, Notre Dame, the Quantitative Methodology Program (QMP) seminar at Michigan, the 2004 APPAM meetings in Atlanta, the 2004 SEA meetings in New Orleans, the 2006 SOLE meeting in Boston, the 2006 IZA/CEPR ESSLE in Amersee, and the 2006 IZA/IFAU Conference on Labor Market Policy Evaluation in Uppsala for helpful comments. We thank participants in a 2002 OECD conference on evaluating local economic development programs for inspiration. Any and all remaining “errors” are actually part of the subgame-perfect equilibrium strategy of a larger game.

## 1.0 Introduction

“Did that program help you?” Many surveys ask people to evaluate the impact some program had on them: We call this participant evaluation. We consider how people might do this, and ask whether they seem to do it well or poorly. If participant evaluations compare favorably with formal impact estimates from controlled experiments, they could be a valuable and cheap substitute for those experiments, as well as a useful check on less compelling non-experimental econometric evaluation strategies. Government bureaucracies have used participant evaluations as part—in some cases, most—of their evidence when evaluating education and labor market programs, both in the United States and Canada (USDOE 2005; HRSD-Canada 2009).

We compare survey-based participant evaluations of a job training program with compelling econometric estimates of the program’s impacts. What should we expect from such comparisons? On the one hand, participants may evaluate programs in a similar fashion to econometricians. They may use their available evidence to construct consistent impact estimates. We call this the “decision theory” view of participant evaluations, and we argue that this view predicts positive covariance of participant evaluations and econometric impact estimates.

On the other hand, participant evaluations may be largely unrelated to econometric impacts, and we describe two theoretical views consistent with no relationship. First, as Nisbett and Ross (1980) discuss, agents may act as “lay scientists” when asked to produce verbal judgments about the causal structure of their social environment, their own behavior or that of other agents. Acting as “lay empiricists,” participants may depend on inherently unrepresentative data (Nisbett and Ross 1980) or fail to correct for potential confounds (Nisbett and Wilson 1977) in forming their evaluations.<sup>1</sup> Moreover, acting as “lay theorists,” participants may simply consult one of their implicit theories and provide a program evaluation based upon it (Ross 1989). In either case, lay scientists may not appropriately construct the counterfactual outcomes required to estimate program impact.

Simon’s (1957) notion of “subjective rationality” offers a second reason why we might observe little relation between econometric impacts and participant evaluations. Participants and econometricians may both make sound inferences about program success, given their own evaluational premises and definitions of success; but their definitions and premises may simply

---

<sup>1</sup> Similarly, Wolfers (2007) argued that voters appear to rely on crude proxies to evaluate the performance of an incumbent’s economic policies.

differ. Consistent econometric impact estimates measure program effects on specific outcomes of policy interest (e.g. earnings or employment, over some specific time period). Participant evaluations instead measure program effects on outcomes, and over time periods, that depend in unobserved ways on the wording of the underlying survey questions and each participant's (possibly idiosyncratic) interpretation of that wording. Under this interpretation, weak relationships between participant evaluations and econometric impact estimates can occur even if participants care a lot about the specific outcomes and time periods analyzed by econometricians, because these may still be but a subset of the outcomes and/or time periods reflected in the participant evaluations.

We present evidence from the National JTPA Study (NJS), an evaluation of the U.S. Job Training Partnership Act (JTPA), which was the largest federal employment and training program for the disadvantaged in the 1980s and 1990s. We use the NJS data for two reasons. First, like many other evaluations of active labor market programs, the NJS includes survey questions that ask participants whether they believe the program helped them in some way. Second, the NJS provides high quality experimental data for a relatively large sample. As a result, we can easily obtain compelling econometric estimates of program impact.

Our first empirical analysis compares econometric impact estimates to participant evaluations using the NJS data. We use heterogeneous program impacts obtained under two different identification strategies to examine whether participants who receive larger program impacts are more likely to report that the program benefited them. Our analysis reveals no relationship between the econometric impact estimates and the participant evaluations.

Our second empirical analysis examines whether participant evaluations correlate with crude impact proxies such as service intensity, outcome levels, and before-after outcome differences. There is little reason to expect a correlation under the subjective rationality view. The lay scientist view in contrast suggests that impact proxies would be correlated with program participant evaluations. Our analysis reveals that crude impact proxies, such as service intensity, outcome levels, and before-after outcome differences, do indeed correlate with participant evaluations, thereby providing evidence in favor of the lay scientist interpretation.<sup>2</sup>

---

<sup>2</sup> Importantly most of the impact proxies are uncorrelated with program impacts. We discuss this issue in Section 6. and report results comparing impact proxies to predicted program impacts in the online appendix Table A6.

We note, but do not consider in detail, other explanations for our negative findings, including low effort by respondents, a desire to help program staff by reporting a positive evaluation regardless of the respondent's actual views, or a similar desire to please the in-person interviewer. These explanations can account for the lack of a relationship between econometric impacts and participant evaluations, but not for our findings regarding impact proxies.

Our paper contributes to several strands of related research. A small number of studies compare survey-based ex post participant evaluations with econometric estimates of program impacts in contexts similar to our own: Kristensen (2006) looks at worker training in a Danish firm, Eyal (2010) studies vocational training in Israel, Calónico and Smith (2012) use data from the U.S. National Supported Work Demonstration and Byker and Smith (2012) use data from experimental evaluation of the Connecticut Jobs First program, respectively. All four papers obtain findings broadly similar to those we report here.

In other contexts, Carrell and West (2010) find that Air Force Academy student evaluations are positively correlated with contemporaneous professor value-added but negatively correlated with follow-on student achievement, while Kelly (2003) finds no link between citizen satisfaction with police and fire services and administrative performance measures (which are themselves problematic as measures of program impact or quality). Another set of studies considers the effect of experiencing a policy on broader political views related to the policy. For example, Di Tella, Galiani and Schargrodsky (2007) show that squatters who receive land titles have more free market beliefs than those who do not. Similarly, Di Tella, Galiani, and Schargrodsky (2012) show that first-hand experience with services provided by a privatized company affects opinions about privatization. A third set of studies examine subjective evaluations by non-participants. For example, Jacob and Lefgren (2008) show the subjective evaluations of school principals positively correlate with econometric estimates of teacher value-added while Bell and Orr (2002) find that caseworkers ex ante evaluations of which welfare recipients participating in the AFDC Homemaker-Home Health Aide Demonstration will benefit most from the program have no predictive content.

Finally, a related literature uses participant decisions to complete or drop out of a program as implicit participant evaluations. For example, Heckman and Smith (1998) consider drop out behavior in the context of the same NJS data we use, Philipson and Hedges (1998)

study dropout from clinical trials in medicine, and Oreopoulos and Hoffman (2009) show that low value-added instructors increase course dropout in a large Canadian university.

The remainder of the paper unfolds as follows. Section 2 introduces a conceptual framework for interpreting the relationships between econometric and participant program evaluations. Section 3 describes the basic structure of the JTPA program, the NJS experiment and the resulting data. Section 4 discusses the construction and interpretation of our econometric estimates of program impact. Section 5 presents results on the relationship between participants' self-reported impacts and impacts estimated using the experimental data. Section 6 examines the relationship between participant evaluations and proxies such as inputs, outcome levels and before-after employment and earnings changes. Section 7 lays out our conclusions.

## 2.0 Conceptual Framework

Three viewpoints shape our empirical analysis and our interpretations of empirical results: A “decision theory” view, a “lay scientist” view and a “subjective rationality” view. These three views draw on literatures that straddle psychology, economics, statistics and survey design. We emphasize that the three views are neither mutually exclusive nor exhaustive; indeed, we expect that all three capture important aspects of the underlying reality, while leaving room for other explanatory factors as well.

To fix ideas and introduce notation, let  $Y_i(D_i)$  be an outcome of interest to policy makers and econometricians, where  $i$  indexes participants and  $D_i = 1$  or  $0$  indicates whether  $i$  was randomly assigned to the treatment or control condition, respectively, of a program evaluation experiment. In the case of an active labor market program,  $Y_a(1)$  might be the earnings of participant Anne ( $i = a$ ) if she was randomly assigned to receive program services and did receive them ( $D_a = 1$ ), while  $Y_a(0)$  would be Anne's earnings if she was randomly assigned to the control group and so did not receive program services ( $D_a = 0$ ).<sup>3</sup> Only one of these outcomes can occur for Anne: Therefore, any estimate of the expected value or sign of  $\Delta_a = Y_a(1) - Y_a(0)$ , the “program impact” on Anne, requires counterfactual reasoning, whether by Anne or an observing econometrician. Usually, treated respondents are the ones surveyed

---

<sup>3</sup> We set aside complicating issues of compliance and substitution. Sometimes, those randomly assigned to treatment drop out and so do not receive it, or those randomly assigned to the control group “cross over” and receive treatment. In other cases, those assigned to the control group obtain other services similar to the treatment in the world outside of the experiment, which is known as “control group substitution”; see e.g. , Heckman et al. (2000).

( $D_i = 1$  for all surveyed  $i$ ), and are asked whether or not the program services were beneficial or helpful (in some sense specified by some question). Formally this seems a request that a treated Anne, having directly experienced only  $Y_a(1)$ , estimate the sign of  $\Delta_a = Y_a(1) - Y_a(0)$ .

### 2.1 A Benchmark Case

We can imagine conditions under which some form of participant evaluation could reveal more information than econometric impact estimates. This optimistic benchmark case is characterized by five assumptions, each of which may or may not be true.

A1: Complete outcome resolution. At the time of questioning respondent  $i$ , all events on which the outcomes  $Y_i(D_i)$  are conditioned have occurred. The outcomes are determined (or in the case of counterfactual outcomes, would now be fully determined): They are wholly things of the past.

A2: Event omniscience. Respondent  $i$ 's information at the time of questioning is complete enough to compute both  $Y_i(0)$  and  $Y_i(1)$  with certainty. If  $Y_i(0)$  and  $Y_i(1)$  are conditioned on non-identical sets of events, the respondent knows all events in the union of those sets.

A3: Mutual outcome correspondence. Let  $Y_i(D_i)$  be the outcomes of interest to researchers and policy makers, and let  $\mathcal{Y}_i(D_i)$  be the outcomes respondent  $i$  answers questions about. Mutual outcome correspondence holds if  $\mathcal{Y}_i(D_i) \equiv Y_i(D_i)$ : When answering questions, respondents consider exactly the same outcomes as the researcher.

A4: Conscious cognitive competence. Respondent  $i$  is cognitively able to compute the counterfactual outcome in consciousness and verbally report it.

A5: Motivational dominance. Respondent  $i$  is sufficiently motivated (by intrinsic and/or extrinsic rewards) to consciously compute the counterfactual outcome and report it. Any subjective costs of doing so are dominated by sufficient positive motivation to do so (cf. Smith 1982).

Suppose for instance that Anne was randomly assigned program services, began them on January 1<sup>st</sup> of 2008 and completed them six months later. She also may accumulate earnings at any time during the 2008 calendar year, including during the six months of the program. The earnings Anne accumulated from January 1<sup>st</sup> of 2008 to December 31<sup>st</sup> of 2008 is the outcome  $Y_a(1)$  available to the econometrician.

*Complete outcome resolution* (assumption A1) says that we interview Anne on or after January 1<sup>st</sup> of 2009, so that the outcome (2008 earnings) is wholly a thing of the past: There is no remaining uncertainty about the outcome, and Anne understands that she should not include

forecasts of future earnings when answering the question. *Event omniscience* (assumption A2) says that even though Anne received program services, she knows any and all past events needed to compute  $Y_a(0)$ , what she would have earned in 2008 if she had not received those services. *Mutual outcome correspondence* (assumption A3) says that if we ask Anne whether the services made earnings better, Anne fully understands that we are asking about earnings accumulated between January 1<sup>st</sup> 2008 and December 31<sup>st</sup> 2008, and exactly what that means: If the econometrician’s data does not include unreported tips, Anne understands that and does not include those either. *Conscious cognitive competence* (assumption A4) says that Anne is cognitively equipped to consciously compute and verbally declare whether  $\Delta_a = Y_a(1) - Y_a(0) > 0$  or not. Economic theory usually assumes that agents automatically know everything logically entailed by other things they know. However, since the cognitive sciences do not assume this, we explicitly state A4 as an additional assumption. To wrap things up, *motivational dominance* (assumption A5) says that Anne is sufficiently well-motivated to cooperate perfectly with the questioner.

Under these assumptions, Anne knows  $Y_i(1)$  with certainty through experience, and is willing, able and sufficiently informed to compute  $Y_i(0)$  without error or any remaining uncertainty. She compares the two outcomes and her yes/no participant evaluation  $R_a$  is

$$(2.1) \quad R_a = 1[\Delta_a > 0],$$

where  $1[\textit{expression}]$  denotes the indicator function, equal to “1” when *expression* is true and “0” when it is not. In this benchmark case, these reports by respondents will reveal the direction of each respondent’s impact perfectly. We could also imagine asking (treated) respondents to report the counterfactual  $Y_i(0)$  (or  $\Delta_i$ ) directly, though this is rarely done. Note that in this benchmark case, the information potential of participant evaluations actually exceeds that of an observing researcher. Even with a fully randomized experiment, a researcher only observes any participant in one “state”—the treated state or the control state. Event omniscience and conscious cognitive competence (assumptions A2 and A4) essentially imply that the participant can compute her own counterfactual outcome and report it (or whether  $\Delta_i > 0$ ) to the researcher.

Now we consider three ways in which the benchmark case can break down.

## 2.2 Decision theory

The decision theory view becomes relevant if A1 and/or A2 fail to hold. Manski (1990, p. 940) popularized this view though he gives large credit to Juster (1964, 1966): “Divergences [between responses and outcomes] may simply reflect...events not yet realized at the time of the survey. Divergences will occur even if responses...to questions are the best predictions possible given the available information.” Manski recognized that much of his analysis was more general than his particular subject matter (p. 935): “Although the substantive concern of this article is the use of intentions data, most of the analysis applies equally to a larger class of prediction questions asked in surveys.” Section 2.3 of Manski (1999) explicitly discusses counterfactual scenarios.

There are two key features of the decision theory view. The first is uncertainty: The outcome in question depends either on future events (genuine uncertainty about the future, violating the assumption A1 of complete outcome resolution), or on past events that are simply unknown to the respondent (most likely, we think, uncertainty about past events relevant to the counterfactual outcome, violating the assumption A2 of event omniscience).

The second key feature is a discrete answer format—for instance, the binary yes/no response to questions such as “Did the program help?” With uncertainty and a binary decision, the respondent’s situation resembles that of “the statistician” in modern statistical decision theory, and this is the perspective Manski (1990) develops. In this view, Anne does not know for sure whether  $\Delta_a = Y_a(1) - Y_a(0) > 0$  or not. However, the decision theory view assumes that Anne knows the distribution  $F(\cdot | Z_a)$  of  $\Delta_a$  conditional on information  $Z_a$  available to her at the time of the survey, and that she can and does compute an objectively rational probabilistic forecast of the event  $\Delta_a = Y_a(1) - Y_a(0) > 0$  given  $Z_a$ : denote this probability as

$$(2.2) \quad P_a = 1 - F(0|Z_a) = \text{Prob}[\Delta_a > 0|Z_a].$$

When answering a binary response participant evaluation question, Anne can make two kinds of errors: affirming that  $\Delta_a > 0$  when it is false, and denying that  $\Delta_a > 0$  when it is true. For many kinds of loss functions (associated with these two possible mistakes), we can characterize Anne’s decision rule and participant evaluation in terms of a critical value  $c_a$  as follows:

$$(2.3) \quad R_a = 1[P_a > c_a].$$



The critical value  $c_a$  simply reflects the relative importance of the two kinds of mistakes (to Anne). We follow Manski (1990) in treating  $c_i = c = 0.5 \forall i$  as a maintained hypothesis.<sup>4</sup>

Manski (1990) shows that binary response forecast data only provide very weak bounds on future binary outcomes under a decision-theoretic formulation like (2.3). Yet even these weak bounds allow for a test of the decision-theoretic view. Suppose we have a vector of observed characteristics  $X_i$  for each respondent  $i$ : Following Manski, we assume these observed characteristics are a subset of the information  $Z_i$  that each respondent knows when making their own evaluation. Let  $S$  be a “subgroup” of respondents  $i$  who share the same values of the vector  $X_i$ : Formally, subgroup  $S$  is  $\{i | X_i = X_S\}$ . Let  $R_S = E(R_i | X_i = X_S)$ : This is just the expected participant evaluation (2.3) in subgroup  $S$ . Similarly, let  $P_S = Prob[\Delta_i > 0 | X_i = X_S]$ : This is the expected proportion of participants in subgroup  $S$  who have a positive impact. Manski (1990, p. 937, eq. 8) shows that under the decision-theoretic view,

$$(2.4) \quad P_S \in [cR_S, c + (1 - c)R_S] \text{ for any subgroup } S.$$

Expression (2.4) may be viewed as a predicted relationship between two different estimators—a consistent estimator of  $R_S$  and a consistent estimator of  $P_S$ —under the null hypothesis of the decision theory view. We can easily generate a consistent estimate of  $R_S$  as the proportion of positive participant evaluations in subgroup  $S$ . In contrast, we lack a straightforward estimator for of  $P_S$ , a difficulty that leads us to introduce a second version of the decision theoretic view producing an alternative test, to which we now turn.

Manski (1990) considered surveyed intentions (yes/no binary responses) concerning a future binary outcome, such as a future purchase or vote decision. In program evaluation many interesting outcomes are not discrete:  $Y_i(D_i)$  may be earnings, or weeks of employment, and so forth. In such instances, it is far less clear that respondents ought to inspect the probability of some discrete event in order to answer participant evaluation survey questions. Above, we implicitly assumed that respondents would in fact do so: Equations (2.2) and (2.3) say that respondent  $i$  inspects the *probability that  $\Delta_i$  is positive*—discretizing an underlying continuous outcome into a threshold crossing event and its complement. This implicit assumption made our setting formally identical to the one considered by Manski (1990) and allowed us to borrow

---

<sup>4</sup> Perhaps the most natural assumption is that Anne and most respondents view the two kinds of mistakes as equally costly. In that case  $c_a = 0.5$ , and Anne’s participant evaluation is “Yes, the program helped” if she believes it more likely that  $\Delta_a > 0$  than not.

Manski's derivations to get expression (2.4). But respondents might inspect other features of the underlying distribution  $F(\Delta_i|Z_i)$ , most notably its expectation.

A second decision theory view is based on the respondent's expectation of  $\Delta_i$  at the time of the survey. We maintain the two "decision theory assumptions," that (1) there is remaining uncertainty about  $\Delta_i$  at survey time, and that (2) the survey response format is binary. As before, respondents have information  $Z_i$  and a conditional distribution  $F(\Delta_i|Z_i)$  of  $\Delta_i$ . Given conscious cognitive competence (assumption A4), respondents can then produce an expected impact  $E(\Delta_i|Z_i)$  in consciousness, and we can imagine some critical value  $k_i$  such that the respondent is willing to declare that the expected program impact is positive, given the relative costs of judgment errors and the quality of her information  $Z_i$ . Under this "second decision theory view," the yes/no binary response is

$$(2.5) \quad R_i = 1[E(\Delta_i|Z_i) > k_i].$$

The researcher can also estimate program impacts for respondents  $i$ . Again following Manski (1990), we assume that a researcher's information about respondent  $i$  is a vector of observed characteristics  $X_i$  that is a subset of the respondent  $i$ 's own information  $Z_i$ . Because  $X_i \subseteq Z_i$ , we can think of  $E(\Delta_i|X_i)$  as the conditional expectation of  $E(\Delta_i|Z_i)$ , so that  $E(\Delta_i|Z_i) = E(\Delta_i|X_i) + u_i$  where  $u_i$  is a mean zero error that is orthogonal to  $X_i$ . Therefore, across the sampled population of respondents  $i$ ,

$$(2.6) \quad Cov[E(\Delta_i|Z_i), E(\Delta_i|X_i)] = Cov[E(\Delta_i|X_i) + u_i, E(\Delta_i|X_i)] = Var[E(\Delta_i|X_i)].$$

Equation (2.6) says that (a) *if* there is some variation in conditional expected impacts, and (b) *if* respondents directly reported consistent expected impact estimates  $E(\Delta_i|Z_i)$ , these would be positively correlated with the researcher's own consistent impact estimates  $E(\Delta_i|X_i)$ . In other words, because respondents use all of the information used by researchers to estimate program impacts (and possibly more), the two impact estimates must be positively related as long as expected impacts vary with this information (i.e.  $Var[E(\Delta_i|X_i)] > 0$ ).

Suppose now that we are willing to assume that  $k_i = k \forall i$ , just as we assumed that  $c_i = c \forall i$  in the previous section. If  $R_i$  is given by (2.5) and  $Var[E(\Delta_i|X_i)] > 0$ , then:<sup>5</sup>

$$(2.7) \quad Cov[R_i, E(\Delta_i|X_i)] > 0.$$

---

<sup>5</sup> Actually, expression (2.7) requires one more assumption that rules out a pathological case: We discuss this later in Section 4.1.

In words, expression (2.7) says that a respondent's binary report of program impact will be positively correlated with expected conditional program impacts as long as conditional impacts vary with the information both the researcher and respondent use to construct their consistent impact estimates. We take expression (2.7) to be the central prediction of the second decision theory view, and the one we test in our analysis.

Rejection of expression (2.4) and/or expression (2.7) does not mean that participant evaluations are simply unmotivated mental coin flips.<sup>6</sup> As we show in what follows participant evaluations are not random. Rather, they are correlated with several variables, including variables that are wholly unrelated to econometric impact estimates. The “lay scientist” view, discussed next, gives some shape to these findings.

### *2.3 Lay scientists*

In psychology, the term “lay scientist” goes back to Kelly (1955), but today it is most widely associated with Nisbett and Ross (1980).<sup>7</sup> Nisbett and Ross discuss the idea that agents act as “lay scientists” when asked to produce verbal judgments about the causal structure of their social environment, their own behavior or that of other agents. Like real scientists, lay scientists make these judgments using either empirical or theoretical reasoning, or some mixture of these, depending on how they interpret questions and, perhaps, which approach appears reasonable or appropriate to them. Yet lay scientists are not idealized professional scientists, in two critical senses. First, when acting as “lay empiricists” at the behest of a survey questioner, they are not compelled to follow canons of formal inference on pain of professional embarrassment if they do not. The collection of sometimes biased “judgment heuristics” popularized by Kahneman, Slovic and Tversky (1982) is, in the view of Nisbett and Ross, a large part of lay empiricists’ arsenal. Second, when acting as “lay theorists,” lay scientists may apply some theories that are generally less well supported by formal canons of evidence than the theories of idealized professional

---

<sup>6</sup> We say little about motivational dominance (assumption A5)—that respondents are sufficiently motivated to do what the survey interviewer wants them to do. If conscious counterfactual reasoning was simple this might not be a problem, but experimental work suggests that motivation (specifically, the presence and/or magnitude of extrinsic incentives for good choice or judgment) starts to matter when choice and judgment involve relatively large stakes (Holt and Laury 2002) or relatively complex alternatives or stimuli (Wilcox 1993).

<sup>7</sup> The term “lay scientist” was most current during the 1980s. Although the currency of the term itself has diminished, the influence of the ideas is ongoing and widespread. Just since 2000, Google Scholar reports nearly 3000 citations of Nisbett and Wilson (1977), about 2000 citations of Nisbett and Ross (1980) and about 650 citations of Ross (1989). We like the term because it creates a natural contrast to the researcher, the professional scientist, who (hopefully!) estimates program impacts using formal canons of statistical inference.

scientists. These two possibilities, of course, can interact: following Ross (1989), if lay scientists use a poorly supported theory as an identifying restriction for empirical inference, their inferences may be relatively flawed.

For a simple and pertinent example of lay theory in action, suppose participants have a lay theory that program impacts on any participant generally increase with input expense or resource intensity. They may then be more likely to say that a program service had a positive impact on them if it seemed relatively expensive or resource-intensive, *ceteris paribus*, even without thinking about any actual evidence concerning themselves or others.

Many potential mistakes of lay empiricism are quite humdrum—the kind of things one teaches new graduate students in a research methods class to avoid. For instance, Nisbett and Wilson (1997) point out that any potential cause that is not salient to people at the time of judgment, and hence ignored, can be a source of bias; this is simply omitted variable bias. Participants might wholly depend on relatively crude proxies, such as simple before-after comparisons, in order to make judgments, without accounting for the fact that other causes might have produced the change (or lack of it). The example of “Ashenfelter’s Dip” considered in Heckman and Smith (1999) provides an example of such a confounding omitted variable; in that case, the omitted confounding variables would have caused outcomes for program participants to improve even had they not participated in the program. Interestingly, before-after comparisons and other crude impact proxies are commonly collected and used in administrative performance standards systems for employment and training programs, perhaps because they provide quick and inexpensive bureaucratic alternatives to the more difficult construction of consistent impact estimates.<sup>8</sup> Participants may rely on the very same proxies to construct their responses to participant evaluation survey questions.<sup>9</sup>

---

<sup>8</sup> See e.g. Heckman, Heinrich, Courty, Marschke and Smith (2011) for a critical discussion of the literature on administrative performance measures based on simple impact proxies.

<sup>9</sup> We should emphasize that the lay science view is no necessary challenge to the notion that people’s minds accurately represent alternatives and their future consequences when skillfully making decisions. Recall our assumption of “conscious cognitive competence” (assumption A4)—that a respondent is cognitively able to compute the counterfactual outcome in consciousness and verbally report it. The heart of Nisbett and Wilson’s (1977) and Ross’s (1989) surveys is a dissociation between experimentally measured causes of subject behavior and subjects’ own verbal reports on those causes. Nisbett and Wilson remind us that though subjects sometimes tell more than they can know, they also clearly know more than they can tell (Polanyi 1964). There is no necessary paradox here. Skilled performance frequently depends on information and processes hidden from consciousness. Therefore, skillful decision making does not imply any capacity for accurate verbal reports on all of the information and processes that undergird that skill. A neoclassical economist gives no important ground by embracing a lay science

## 2.4 Subjective Rationality

Subjective rationality will matter whenever A2 and/or A3 fail(s) to hold. Put simply, a person's choice, estimate or report depends on their choice set, what outcomes they value and/or what they believe. The term "subjective rationality" comes from Simon (1957, p. 278). Simon's reference to "consequences" is like the outcomes  $Y_i(D_i)$  of interest to the formal analyst and the outcomes  $\mathcal{Y}_i(D_i)$  considered by respondent  $i$  when answering questions.

Subjective rationality becomes a potential problem when  $\mathcal{Y}_i(D_i) \neq Y_i(D_i)$ . To illustrate this, consider the example of "rationality as viewed by the researcher." Schochet, Burghardt and McConnell (2008) estimate the impact of Job Corps participation on very specific outcomes  $Y_i(D_i)$  such as earnings, GED receipt and arrests. Presumably they chose these outcomes for several reasons (including data availability), but also because the program was designed to increase earnings and contained a significant GED preparation component.

"Rationality as viewed by the respondent" could be very different. One Job Corps evaluation survey question asks respondents (like our Anne) whether or not they would recommend the program to a friend. Would Anne naturally think of earnings, GED receipt and/or arrests when answering this question? Perhaps she would, but nothing about this question demands or clearly suggests that she should. Suppose instead that Anne has a positive overall emotional remembrance of her Job Corps experience—what psychologists might call "positive affect" for the Job Corps experience. Outcomes such as meeting new friends may loom large among the outcomes  $\mathcal{Y}_i(D_i)$  that gave Anne her overall positive affect for Job Corps and so may be the primary reasons she reports "Yes (I would recommend Job Corps to a friend)." This report is neither frivolous nor a mistake. After all, Anne and her friends may value similar outcomes.

In this example, the outcomes  $\mathcal{Y}_i(D_i)$  implicitly considered by Anne are clearly not the outcomes  $Y_i(D_i)$  analyzed by Schochet, Burghardt and McConnell (2008), nor are they the outcomes valued by policymakers. We have a failure of A1 in this instance, the most pure form of a subjective rationality problem where either (1) the respondent's interpretation of the question causes her to consider outcomes far different from the ones that interest researchers and policymakers, or (2) the respondent bases her answer on some overall evaluation of the program

---

interpretation of participants' inability to accurately *report* program impacts. The capacity for verbal report and the capacity for decision making may simply be two different things.

which mostly reflects outcomes that are not of interest to researchers or policy makers. In this case, Anne’s participant evaluation is

$$(2.8) \quad R_a = 1[y_a(1) - y_a(0) > 0],$$

and whenever  $y_i(D_i) \neq Y_i(D_i)$ , this is not identical to (2.1).

Anne might also focus on different time periods than the econometrician, who is constrained by the amount of follow-up data available on the outcomes. For instance, Anne could look ahead and consider the effects of having participated in Job Corp on future educational and employment outcomes. The “subjective rationality” view assumes that both the participants who respond to the evaluation questions and the econometricians seeking to estimate program impacts make rational judgments about program success. If their conclusions differ, this happens merely because they consider different sets of valued consequences or outcomes.

### **3.0 Data and institutions**

#### *3.1 The JTPA program*

The U.S. Job Training Partnership Act program was the primary federal program providing employment and training services to the disadvantaged from 1982, when it replaced the Comprehensive Employment and Training Act (CETA) program, to 1998, when it was replaced by the Workforce Investment Act (WIA) program. All of these programs share more or less the same set of services and serve the same basic groups. They differ primarily in their organizational details (e.g. do cities or counties play the primary role) and in the emphasis on, and temporal ordering of, the various services provided. Nonetheless, the commonalities dominate with the implication that our results for JTPA likely generalize to WIA (and CETA).

The JTPA eligibility rules included categorical eligibility for individuals receiving means tested transfers such as Aid to Families with Dependent Children (AFDC) or its successor Temporary Aid to Needy Families (TANF) or food stamps. Individuals in families with incomes in the preceding six months below certain cutoffs were also eligible. Finally, an “audit window” allowed up to 10 percent of participants at each site not to satisfy these rules.<sup>10</sup>

The JTPA program provided five major services: classroom training in occupational skills (CT-OS), subsidized on-the-job training (OJT), job search assistance (JSA), adult basic

---

<sup>10</sup> Devine and Heckman (1996) provide more detail on the JTPA eligibility rules while Heckman and Smith (1999, 2004) study the JTPA participation process.

education (ABE) and subsidized work experience (WE). Local sites had the flexibility to emphasize or de-emphasize particular services in response to the perceived needs of the local population and the availability of local service providers. In general, CT-OS was the most expensive service, followed by OJT, ABE and WS. JSA cost much less.<sup>11</sup>

Services were assigned to individuals by caseworkers, typically as the result of a decision process that incorporated the participant's abilities and desires. This process led to clear patterns in terms of the characteristics of participants assigned to each service (Kemple, Doolittle and Wallace 1993). For example, the most job ready individuals typically were assigned to JSA or OJT, while less job ready individuals typically were assigned to CT-OS, BE or WE, where CT-OS was often followed by JSA. This strongly non-random assignment process has implications for our analyses below in which we examine the relationship between the participant evaluations and types of services received.

### *3.2 The National JTPA Study data*

The NJS evaluated the JTPA program using a random assignment design. Random assignment in the NJS took place at a non-random sample of 16 of the more than 600 JTPA Service Delivery Areas (SDAs). The exact period of random assignment varied among the sites, but in most cases random assignment ran from late 1987 or early 1988 until sometime in the spring or summer of 1989. A total of 20,601 individuals were random assigned, usually but not always with the probability of assignment to the treatment group set at 0.67.

The NJS data come from multiple sources. First, nearly all those randomly assigned completed a Background Information Form (BIF) at the time of random assignment. The BIF collected basic demographic information along with information on past schooling and training and on labor market outcomes at the time of random assignment and earlier. Second, all experimental sample members were asked to complete the first follow-up survey around 18 months after random assignment. This survey collected information on employment and training services (and formal schooling), as well as information on employment, hours and wages, from which a monthly earnings measure was constructed. Third, a random subset (for budgetary reasons) of the experimental sample was asked to complete a second follow-up survey around 32

---

<sup>11</sup> See Heinrich, Marschke and Zhang (1998) for a detailed study of costs in JTPA and Wood (1995) for information on costs at the NJS study sites.

months after random assignment. Response rates to both follow-up surveys were around 80 percent. Finally, administrative data on quarterly earnings and unemployment benefit receipt from state UI records in the states containing the 16 NJS sites were collected.<sup>12</sup>

### *3.3 The participant evaluation questions*

Two questions from the first follow-up survey, taken together, define the participant evaluation measure we use in this paper. The skip pattern in the survey excludes control group members from both questions. The first question asks treatment group members whether they participated in JTPA:

(D7) According to (LOCAL JTPA PROGRAM NAME) records, you applied to enter (LOCAL JTPA PROGRAM NAME) in (MONTH/YEAR OF RANDOM ASSIGNMENT). Did you participate in the program after you applied?

The question assumes application because it is implied by the respondent having been randomly assigned. The JTPA program had different names in the various sites participating in the evaluation; the interviewer included the appropriate local name in each site as indicated in the question. The second question was asked only of those with a positive response to the first:

(D9) Do you think that the training or other assistance that you got from the program helped you get a job or perform better on the job?

This question has a number of problems. It does not explicitly prompt the respondent to think in counterfactual terms. The outcome is vague and composite, though at least it is clear that the respondent is to think about labor market outcomes. No explicit time period is specified, so that a

---

<sup>12</sup> See Doolittle and Traeger (1990) on the design of the NJS, Orr et al. (1996) and Bloom et al. (1993) for the official impact reports and Heckman and Smith (2000) and Heckman, et al. (2000) for further interpretation.



respondent who had not yet found a job might answer in the affirmative if she thought the program would help her find a job in the future.

We code the responses to both questions as indicator variables. The participant evaluation measure employed in our empirical work consists of the product of the two indicator variables. Put differently, our self-reported evaluation measure equals one if the respondent replies “YES” to question (D7), and “YES” to question (D9), so that individuals who do not recall participating get coded as a negative participant evaluation. Otherwise, it equals zero. The unconditional percentages of respondents with positive participant evaluations equals 63% for adult males, 65% for adult females, 67% for male out-of-school youth and 72% for female out-of-school youth.<sup>13</sup> Appendix Table A2 documents that these percentages do not have a strong positive correlation with experimental impact estimates for the four demographic groups.

### *3.4 Outcome variables*

Our outcome variables consist of earnings and employment measured using both the self-report and UI data, given that previous research finds differences (Kornfeld and Bloom 1999). We separately examine outcomes over the full 18 months after random assignment and in just month 18 using the self-reported outcome data, and the analogous outcomes, namely six calendar quarters and just the sixth calendar quarter, using the UI data. We examine earnings as well as employment in order to capture the “perform better on the job” aspect of the participant evaluation question, as better performance should result in increased hours, wages or both.

## **4.0 Econometric framework**

### *4.1 Predicted impacts: subgroups*

The first method we employ for generating impact estimates that vary among participants takes advantage of the experimental data and the fact that random assignment remains valid for subgroups defined on characteristics unaffected by treatment, as discussed in, e.g. Heckman (1997).

To create estimates of  $E(\Delta_i | X_i = X_S)$ , we estimate regressions of the form

$$(4.1) \quad Y_i(D_i) = \beta_0 + \beta_D D_i + \beta_X X_i + \beta_1 D_i X_i + \eta_i,$$

---

<sup>13</sup> Table A1 in the online appendix provides additional detail on the responses to the underlying survey questions.

where  $Y_i(D_i)$  is an outcome measure,  $D_i$  is an indicator equal to “1” for experimental treatment group members and “0” for experimental control group members,  $X_i$  denotes a vector of characteristics and  $D_i X_i$  represents interactions between the characteristics and the treatment indicator. The interaction terms yield variation in predicted impacts among individuals at the subgroup level. For treated participants  $i$  in subgroup  $S$ , we want “predicted subgroup impacts”  $\widehat{\Delta}_S(X_S)$ , that is, estimates of  $E(\Delta_i|X_i = X_S)$ , based on the estimated coefficients in (4.1). These are given by

$$(4.2) \quad \widehat{\Delta}_S(X_S) = \widehat{\beta}_D + \widehat{\beta}_1 X_S = \widehat{E}(\Delta_i|X_i = X_S).$$

Though quite straightforward conceptually, our experimental subgroup impact estimates do raise some important issues. The first issue concerns the choice of variables to interact with the treatment indicator. We address this issue by presenting two sets of estimates based on characteristics selected in different ways. One set borrows the vector of characteristics employed by Heckman, Heinrich and Smith (2002); the notes to Table 1 list these variables. We select the second set using the method of stepwise regression. While economists typically shun stepwise procedures as atheoretic, for our purposes that bug becomes a feature, as it makes the selection procedure mechanical. Thus, we can be assured of not having stacked the deck in one direction or another. In both cases, we restrict our attention to main effects in order to keep the problem manageable.<sup>14</sup>

The second issue concerns the amount of subgroup variation in impacts in the NJS data within the four demographic groups – adult males and females ages 22 and older and male and female out-of-school youth ages 16-21 – for which both we and the official reports conduct separate analyses. Subgroup variation corresponds to the term  $Var[E(\Delta_i|X_i)]$  in (2.6); clearly, each definition of subgroups (choice of a particular vector  $X$ ) yields a distinct division of the overall variation in impacts into systematic and idiosyncratic components. Although the NJS impact estimates differ substantially between youth and adults, the experimental evaluation reports – see Exhibits 4.15, 5.14, 6.6 and 6.5 in Bloom et al. (1993) and Exhibits 5.8, 5.9, 5.19 and 5.20 in Orr, et al. (1994) – do not reveal much statistically significant variation in impacts among subgroups defined by the baseline characteristics reported. If impacts do vary among individuals, but not in ways that are correlated with our choice of baseline characteristics, we may reach the wrong conclusion about the quality of the participant evaluations. This case has

---

<sup>14</sup> The online appendix describes the stepwise procedure in detail.

more than academic interest given that Heckman, Smith and Clements (1997, Table 3) calculate a lower bound on the impact standard deviation of \$675 for adult women in the NJS data (with a standard error of \$138).<sup>15</sup>

Our methods address concerns regarding a lack of meaningful subgroup variation in impacts in two ways. First, the stepwise procedure only retains a subset of  $D_i X_i$  interactions that attain statistical significance (putting aside, in the spirit of the procedure, concerns about pre-test bias). Second, we examine the quantile treatment effect estimates described in the next section; these estimates contain substantial individual variation.

The third issue concerns an additional assumption required to interpret our results in the way that we do. When we find evidence that  $\text{cov}(\hat{\Delta}_S(X_{S(i)}), R_i) \leq 0$  we take this as evidence against the positive covariance property (2.7) implied by the decision theory view. To see why we need to make an additional assumption to justify this interpretation, consider the following example taken from Heckman, Heinrich and Smith (2011). Suppose respondents focus on earnings, behave according to the second decision theory view, and that  $k_i = k = 0$  in (2.6), so that respondents give a positive evaluation only if they expect a positive earnings impact. Now consider a population composed of just two groups. In group one, 10 percent of the individuals expect a positive \$1000 impact while the rest expect a zero impact: The mean group one impact is \$100 and the fraction giving positive evaluations is 0.1. In group two, 20 percent of the individuals expect a \$400 impact while the rest expect a zero impact: The mean group two impact is \$80 and the fraction giving positive evaluations is 0.2. This example shows that subgroup mean impacts could vary inversely with the fraction of respondents reporting a positive impact even if (2.7) holds at the individual level. When we interpret  $\text{cov}(\hat{\Delta}_S(X_{S(i)}), R_i) \leq 0$  as evidence against the decision theory view we assume that this does not occur in our data.

#### *4.2 Predicted impacts: quantile differences*

The second econometric method again uses the experimental data, but adds an additional non-experimental assumption. The recent literature (e.g. Djebbari and Smith 2008) calls that assumption “rank preservation”, while Heckman, Smith and Clements (1997) call it “perfect

---

<sup>15</sup> Our subgroup impacts have standard deviations that range from \$840 to \$2600 depending on the demographic group and set of covariates. The quantile treatment effect standard deviations range between \$257 and \$477.

positive rank correlation”. Rank preservation assumes that the counterfactual for an individual at a given quantile of the treated outcome distribution is the same quantile of the control outcome distribution, and vice versa. Thus, under rank preservation, quantile treatment effects (QTEs) represent treatment effects both *on quantiles* and *at quantiles*.

Formally, we estimate the impact for the treated individual whose outcome falls at percentile “ $j$ ” of the treatment group outcome distribution as

$$(4.3) \quad \widehat{\Delta}_Q(j) = \widehat{Y}^{(j)}(1) - \widehat{Y}^{(j)}(0),$$

where the superscript “ $(j)$ ” denotes the percentile. In words, we estimate the QTE for a particular percentile of the outcome distribution as the difference in outcomes at that percentile of the treatment and control outcome distributions, and then interpret the QTE as the impact of treatment on an individual  $i$  who falls at percentile  $j$  of the treated distribution. Unlike the subgroup impact estimator defined in the preceding section, this estimator yields predicted impacts that vary among individuals with the same observed characteristics; as a result, it may capture some of the underlying variation in impacts that the subgroups miss.

We do not think rank preservation holds exactly, but it may provide a reasonable approximation, particularly in cases, such as the JTPA program, that correspond to treatments of modest intensity that we would expect to yield only modest changes in individuals’ relative labor market performances.

#### 4.3 Predicted impacts: estimation

We can examine relationships between predicted impacts (based on either subgroup variation or rank preservation) and participant evaluations by simply regressing one on the other. In particular, we estimate this equation:

$$(4.4) \quad \widehat{\Delta}_i = \alpha_0 + \alpha_1 R_i + e_i,$$

where the hat on the econometric impact on the left-hand side denotes an estimate and where  $e_i$  includes all of the unobserved factors that affect the predicted impact, including the estimation error in the predicted impact and any approximation error due to inappropriate linearization.

When examining subgroup impact estimates constructed as in (4.2),  $\widehat{\Delta}_i = \widehat{\Delta}_S(X_{S(i)})$  where  $S(i)$  is the subgroup of respondent  $i$ . When we examine quantile treatment effect estimates constructed as in (4.3),  $\widehat{\Delta}_i = \widehat{\Delta}_Q[j(i)]$  where  $j(i)$  is the percentile of the treated outcome distribution at which respondent  $i$  is located. In both cases, we take non-positive estimates of  $\alpha_1$

as evidence against the covariance property of the second decision theory view summarized in (2.7).

Despite its simplicity, three issues regarding equation (4.4) warrant discussion. First, we seek to measure association, not causation. This follows immediately from the fact that we seek evidence on the covariance property in (2.7); put simply, we want to know whether predicted impacts based on our two econometric approaches covary with the participant evaluations in the manner predicted by one of our three viewpoints.

Second, we made the econometric impact estimate the dependent variable rather than the independent variable for a reason: we know that it embodies non-trivial estimation error and likely other types of measurement error (i.e. due to recall bias) as well. In the linear regression model, putting a variable with measurement error on the right-hand-side leads to biased and inconsistent estimates while, under certain assumptions, putting it on the left-hand side does not. In an attempt to avoid this bias, we make the econometric impact estimates our dependent variable.<sup>16</sup>

Third, and finally, we include no additional covariates on the right-hand side because one of our two econometric impact estimators, namely that based on subgroup variation in the experimental impact estimates, consists of a linear combination of the variables that we would otherwise include as regressors. We could include covariates when using the percentile differences as the dependent variable, but omit them to make the two analyses symmetric.

We do not believe that we can bring the subjective rationality view to any strong test. Yet a weak estimated relationship in equation (4.4) might mean that impacts captured by the econometric estimates (subgroup impacts and quantile treatment effects) are based on outcomes that are of relatively small account to participants. Of course we should not forget what lies in the error term. Among the items in the error term, we would expect impacts in the period after that captured in the predicted impact to correlate positively with short term impacts; in contrast, impacts on leisure likely correlate negatively with impacts on labor market outcomes prior to the

---

<sup>16</sup> Two arguments run counter to our decision to make the econometric impact estimates the dependent variable. First, in the absence of causal concerns, putting the variable with the largest variance on the right-hand side of a simple linear regression minimizes the variance of the resulting slope coefficient. In our context, given a binary participant evaluation measure, this argument militates in favor of putting the predicted impacts on the right-hand side. The second argument questions the implicit assumption we make that the participant evaluation contains no measurement error. In fact, it may well contain substantial measurement error, but the literature provides no guidance, in the form of repeated measures taken a relatively short time apart, on its nature and extent.

survey. A weak relationship in (4.4) could thus also result from a combination of a positive direct correlation between the econometric impacts on labor market outcomes (such as employment or earnings) and a negative indirect correlation with leisure, working through the correlation between the omitted impact on leisure and the included impact on employment or earnings.

Only non-positive estimates of  $\alpha_1$  run counter to the decision theory view when there is substantial variation in program impacts. If the decision theory views appear to be empty in an explanatory sense, the lay science view gives us guidance for exploring the actual correlates of participant evaluations. It suggests variables that participants might use to create their estimates, such as outcomes or before/after outcome differences (in the case of lay empiricists) or measures of program inputs (in the case of lay theorists depending on a theory of positive marginal products of training and education inputs). We discuss our econometric framework for examining those potential relationships in the next subsection.

#### *4.4 Determinants of positive participant evaluations*

To explore the lay scientist view described in Section 2 – that participants’ evaluations depend on relatively crude proxies for impacts – we need to examine relationships between participant evaluations and these potential proxies. As we examine only binary participant evaluation measures, we use standard logit models to study these relationships. In formal notation, we estimate versions of

$$(4.5) \quad R_i = 1(\gamma_0 + \gamma_1 proxy[\mathcal{Y}_i(1) > \mathcal{Y}_i(0)] + \gamma_X X_i + v_i > 0),$$

where  $proxy[\mathcal{Y}_i(1) > \mathcal{Y}_i(0)]$  indicates one (or, in some cases, a vector) of observed proxies for impacts,  $X_i$  denotes a vector of observed characteristics with corresponding coefficients  $\gamma_X$ , and where we assume  $v_i$  has a logistic distribution. Following the notation in Section 2, the script “ $\mathcal{Y}$ ” in the notation for the impact proxies signals that we view them as proxies for impacts on the outcomes as the respondent conceives them. As the standard logit model identifies the coefficients only up to scale, we report mean derivatives rather than coefficient estimates (or derivatives evaluated at the mean of the  $X$ ).

Astute readers will have noticed that the participant evaluations have moved from the right-hand side of equation (4.4) to the left-hand side of equation (4.5). We do this because we think of both objects in equation (4.5) as measured more or less without error, in which case the

argument recounted above regarding relative variances points toward putting the variables with the larger variances, in this case the impact proxies, on the right-hand side. As discussed earlier, one can reasonably worry about the actual extent of measurement error in the participant evaluations. The covariates  $X$  in equation (4.5) soak up residual variance, yielding more precise estimates, and also help to clarify the interpretation of the effects of the impact proxy.

We end with another reminder that we do not estimate the relationships in (4.4) and (4.5) to obtain causal effects. Rather, we seek “merely” to examine the relationships between variables, and use the linear regression framework as a convenient tool to accomplish that goal.

## **5.0 The relationship between econometric impact estimates and participant evaluations**

### *5.1 Regression results for experimental subgroup estimates*

We first consider evidence from regressions of experimental subgroup impact estimates on participant evaluations. In terms of our earlier discussion, we report estimates of  $\alpha_1$  from (4.4), where the dependent variable is obtained via (4.1) and (4.2). Table 1 shows the results. Each entry in the table shows  $\alpha_1$  and its standard error from a different regression. Each row shows all of the regression results using impacts on one of the eight labor market outcomes as the dependent variable. Columns are grouped into four pairs. Each pair corresponds to one of the four demographic groups; within pairs, the columns headed by (1) contain the estimates using the covariate set chosen by the stepwise procedure, while the columns headed by (2) contain the estimates based on the covariates from Heckman, Heinrich and Smith (2002). The final two rows of the table summarize the evidence in each column; they give the numbers of positive and negative estimates and, within each category, the number of statistically significant estimates at the five and ten percent levels.

The regression evidence in Table 1 suggests little, if any, relationship between the impact estimates and the participant evaluations. While the estimates lean negative in the aggregate, only a handful of the estimates reach conventional levels of statistical significance (and not all of those fall on the negative side of the ledger).

## 5.2 Results based on quantile treatment effect estimates

This section presents evidence on the relationship between impact estimates constructed under the rank preservation assumption described in Section 4.2 and participant evaluations. We focus on one particular outcome in this analysis: the sum of self-reported earnings over the eighteen months after random assignment (quantiles of employment provide little in the way of insight for obvious reasons). Figures 1A to 1D correspond to the four demographic groups. The horizontal axis in each figure refers to percentiles of the untreated outcome distribution. The solid line in each graph presents impact estimates at every fifth percentile (5, 10, 15, ..., 95) constructed as in (4.3). The broken lines represent estimates of the fraction with a positive participant evaluation at every fifth percentile. For percentile “ $j$ ”, this estimate consists of the fraction of the treatment group sample members in the interval between percentile “ $j-2.5$ ” and percentile “ $j+2.5$ ” with a positive participant evaluation.

Several features of the figures merit notice. First, in the lowest percentiles in each figure the econometric impact estimate equals zero. This results from the fact that individuals in the lowest percentiles in both the treated and untreated outcome distributions have zero earnings in the 18 months after random assignment, implying an impact estimate of zero. Surprisingly, more than half of the treatment group members with zero earnings provide positive participant evaluations in all four demographic groups. This could mean that respondents view the question as asking about “employability” (rather than just employment) so that they respond positively if they think the program has improved their future employment chances. It could also mean that respondents are acting as lay theorists (more on this shortly).<sup>17</sup>

Second, the fraction with a positive participant evaluation has remarkably little variation across percentiles of the outcome distribution. For all four demographic groups, it remains within a band from 0.6 to 0.8. For the adults, the mean increases with the percentile; for the youth, the data fail to reveal a clear pattern. Third, the figures do not reveal an obvious relationship between the two variables other than for adult females: for them, both variables increase with the percentile of the outcome distribution. More specifically, for adult women, both variables have a higher level for percentiles where the impact estimate exceeds zero. Within the two intervals defined by this point, both variables remain more or less constant.

---

<sup>17</sup> We obtain qualitatively similar findings when using either UI earnings over the six quarters after random assignment or the average of the two earnings variables as the outcome variable in the analysis.



Table 2 presents some of the numbers underlying the figures. The first five rows present the values for the 5<sup>th</sup>, 25<sup>th</sup>, 50<sup>th</sup>, 75<sup>th</sup> and 95<sup>th</sup> percentiles. The last two rows give the correlation between the quantile treatment effects and the fraction with a positive participant evaluation (and the corresponding p-value from a test of the null of a zero correlation) along with the estimated coefficient from a regression of the quantile treatment effects on the fraction with a positive participant evaluation (and its standard error). The estimates in Table 2 quantify and confirm what the figures indicate: a strong positive relationship for adult women, a weak and statistically insignificant positive relationship for adult men, a moderately strong and negative relationship for male youth and a similar, but not statistically significant, relationship for female youth. Although we find a bit more here than in the estimates that rely on subgroup variation, once again the data do not suggest a strong, consistent relationship between the econometric impact estimates and the participant evaluations.

To sum up, we find little evidence of positive relationships between participant evaluations of JTPA and individual impact estimates based on subgroup variation in the experimental impacts or on the rank preservation assumption of the quantile estimator. We also find a majority of treatment group members have positive participant evaluations, including many with zero earnings in the 18 months (or six quarters) after random assignment. Participants are not using the same information in the same way as econometricians to attempt to construct the same impacts.

## **6.0 Relationships between positive participant evaluations and impact proxies**

### *6.1 Motivation and caveats*

The lack of a strong positive relationship between the estimated impacts and the participant evaluations indicates that program participants are not behaving as decision theorists.<sup>18</sup> We next examine the relationship between participant evaluations and impact proxies. As noted above, while the subjective rationality interpretation does not suggest a relationship between participant evaluations and impact proxies, the lay scientist view does.

---

<sup>18</sup> Another possible interpretation is that our econometric impacts estimates contain too much noise, as a result of being imprecisely estimated. As a crude test of this view, we considered individuals with estimated impacts in the top and bottom five percent of the distribution of estimated impacts for each of the two identification strategies and found that these individuals did not have noticeably higher or lower rates of positive participant evaluations. Our findings are also consistent with a combination of low participant effort and a desire to please the interviewer or reward the program, though this possibility does not suggest a relationship with impact proxies.

The results strongly suggest lay scientists at work. Consider the finding that participant evaluations are both largely positive and vary remarkably little across groups, subgroups and quantiles of outcome distributions. This suggests that participants' evaluations may be largely theory-driven inferences based on shared folk theories. In particular, we suspect that participants may share the theory that impacts are monotone increasing in inputs (the expense or resource-intensiveness of program services received). To explore this, we estimate relationships between participant evaluations and services received by participants, and expect relatively large positive effects for relatively expensive (resource-intensive) services.<sup>19</sup>

Our results so far could also be explained by some participants acting as lay empiricists, making judgments based on proxy variables that correlate only weakly with true impacts, and perhaps with insufficient notice of potential confounds. If so, their evaluations can be both inconsistent and full of nuisance variation, undermining any relationship between them and consistent impact estimates. The proxies we examine are actual labor market outcome (employment and earnings) levels and simple before-after differences in those outcomes. If respondents really do know the impacts, then such proxies should have little explanatory power except to the extent that they correlate with actual impacts, which, as Heckman, Heinrich and Smith (2002, 2011) show using the NJS data, they largely do not.<sup>20</sup>

## *6.2 Results with service types*

Table 3 presents results from estimates of (4.5) that include indicators for five self-reported service types received by JTPA treatment group members: CT-OS, OJT/WE (almost all OJT), JSA, ABE and "other".<sup>21</sup> Respondents reporting multiple service spells get coded based on the first such spell. The omitted category is no self-reported service receipt. The lay theorist view

---

<sup>19</sup> An obvious interpretive caveat is that different services may have different subjective or direct costs and/or benefits not captured by labor market outcomes. For example, classroom training may be more fun (or more tedious) than, say, job search assistance. Thus, the subjective rationality interpretation also allows for relationships between participant evaluations and service types, though it makes no obvious prediction about the direction of those relationships. The question wording also does not encourage this interpretation.

<sup>20</sup> We discuss this issue in section 5.4 and report results comparing impact proxies to predicted program impacts in the online appendix Table A6. Echoing the results in Heckman, Heinrich and Smith (2002, 2011), Table A6 shows that most of our impact proxies are uncorrelated with predicted impacts.

<sup>21</sup> We obtain qualitatively similar results when using administrative data on service receipt in place of the self-reported service types even though, as shown in Smith and Whalley (2012), the two data sources often disagree.

predicts positive participant evaluations to be relatively more likely for more expensive services such as CT-OS and OJT.<sup>22</sup>

The logit models also include a variety of background variables; see the table notes for a list (and the on-line appendix for a full set of results). These variables pick up parts of the overall impact of participation unrelated to the labor market outcomes we examine. For example, the site indicators pick up differences in the friendliness and efficiency of site operation as perceived by the respondents. The variable “work for pay”, which is an indicator variable for whether or not the respondent has ever worked for pay (as of random assignment), relates to the opportunity cost of participation, as does the variable for having a young child. The AFDC receipt at random assignment variable captures variation in the cost of classroom training due to the availability of an income source not tied to employment. The background variables also soak up residual variance, thereby increasing the precision of the estimates on the service type indicators.

The top panel of Table 3 displays the mean derivatives for the service type indicators. The bottom panel displays test statistics and the associated p-values from tests of the nulls that specific groups of variables all have zero coefficients. The columns correspond to the four demographic groups. With the sole exception of CT-OS for female youth, for all four demographic groups the more expensive services, CT-OS and OJT/WE, clearly dominate inexpensive JSA. It is less clear what to make of the results for ABE, which has few participants, and “other”, which includes both expensive and inexpensive services.

Turning to the other variables in the bottom panel of Table 3, the site variables have a strong and statistically significant effect on the probability of a positive participant evaluation. Respondents may take account of non-pecuniary aspects of their JTPA experience, such as the staff or the office, even when responding to a question nominally about jobs. Variation in local conditions across sites, such as hiring opportunities, also might affect respondents’ evaluations through an influence on outcomes. Although there might also be site differences in program impacts, this seems less likely given the findings in Section 5.1.<sup>23</sup>

---

<sup>22</sup> Keep in mind that in the JTPA study, access, not service type, was randomly assigned. The exact set of services an individual receives depends on many factors, including the preferences of the participant and the caseworker, their beliefs about service effectiveness, the site budget, and so on. Thus, the estimated effects of the service type indicators on the probability of a positive participant evaluation need not have a causal interpretation. Instead, they likely reflect both the sorting of individuals who believe a particular service to be effective into receipt of that service and the effect of particular services on participant beliefs about program effectiveness.

<sup>23</sup> Bloom et al. (1993) and Orr et al. (1994) also do not find much evidence of site-level differences in impacts.

With the exception of age for adults, race for youth, and age and education for female youth, the demographic variables play surprisingly little role in determining the probability of a positive participant evaluation. Among adults, age has a strong negative effect on the probability of a positive evaluation, while black male youth and Hispanic male and female youth have higher probabilities of a positive response. The limited role played by background characteristics in the analysis surprised us.

### *6.3 Results with labor market outcomes*

Table 4 reports results from estimating versions of (4.5) that include the same background variables as the models in Table 3, but add various versions of  $Y_1$ , the labor market outcome in the treated state. Acting as lay empiricists, respondents may be relatively more likely to infer a positive program impact if they have done well in the labor market between random assignment and the survey, or if they are doing well around the time of the survey.

Table 4 summarizes the evidence. As in the bottom panel of Table 3, the summary takes the form of chi-square statistics, and their p-values, for tests of the null hypotheses that a coefficient (or all coefficients) on a specific labor market outcome measure (or vector of outcome measures) equal(s) zero.<sup>24</sup> The relationships tend to be statistically stronger for adults than for youth, and for measures based on self-reported data than on UI data. Also, earnings measures tend to yield more statistically significant relationships than employment measures, especially for outcomes at or just around the time of the survey.<sup>25</sup>

Overall, we find strong evidence that participants use labor market outcomes as proxies for impacts. This is consistent with lay science, but perhaps with real science too as outcomes may be correlated with actual impacts. Indeed, some fraction of the treated group would have zero counterfactual (untreated) earnings outcomes after random assignment; for them, outcomes and impacts perfectly coincide. We cannot completely rule out this possibility by an appeal to our results in Section 4 because the measurement error in our own impact estimates may be quite large.

---

<sup>24</sup> See the on-line appendix for the full set of results

<sup>25</sup> The results are not sensitive to coding employment as earnings above \$400 rather than as non-zero earnings.

#### *6.4 Results with before-after comparisons of labor market outcomes*

This section explores the relationship between participant evaluations and before-after differences in employment and earnings. The cognitive appeal and simplicity of before-after comparisons as an estimator of impacts are undeniable. Moreover, despite their simplicity, before-after comparisons are consistent impact estimates in the absence of confounds, that is, if there is no change in any outcome-relevant factor over the period between the two measurements, as the “before” outcomes will then consistently estimate the “after” outcome that would have occurred without treatment.

The literature provides dramatic evidence of the importance of confounds in this context under the heading of “Ashenfelter’s dip”. The dip refers to the decline in mean earnings and employment commonly observed for participants in employment and training programs in the period prior to participation. The dip reflects selection into programs on the basis of transitory labor market shocks. As Heckman and Smith (1999) show using the experimental control group data from the NJS, the labor market outcomes of participants would improve in the “after” period even in the absence of participation. Remembering that lay empiricists may fail to correct for non-salient confounds, participants making judgments on the basis of before-after comparisons may well fail to appreciate that they would have found a job even without participating in JTPA, particularly when the survey question does not push them to think about counterfactuals.

Given that one of our survey questions asks directly about finding a job, in Panel A of Table 5 we first consider the relationship between participant evaluations and before-after changes in employment. We coded the before-after employment status variable based on employment at the date of random assignment and 18 months after random assignment. This yields four patterns. We include dummy variables for three of the four patterns, with employed at both points in time as the omitted pattern. The findings here are, perhaps, less strong than expected. In general, relative to the always employed, those who are never employed or who lose a job tend to have less positive participant evaluations. Among adults, those who gain a job tend to be somewhat more positive. However, only a handful of the differences achieve statistical significance. Measurement error in the “after” employment status may account for our weak results. By looking at employment around the time of the survey, we have given the respondents plenty of time to lose jobs that JTPA helped them find and, in the case of program dropouts, to find jobs without the help of JTPA. However, the joint tests for the employment change variable

look stronger than the individual t-tests, which is consistent with our omitted group lying in the middle of the categories in terms of its effect on the participant evaluation measure.

Table 5 Panel B presents estimates of versions of (4.5) that include before-after earnings changes as independent variables, along with the usual background variables. The earnings change measure consists of the difference in average self-reported monthly earnings between the 12 months before random assignment and the 18 months after random assignment. We can use only the 12 months before random assignment due to the limitations of the survey data on pre-random assignment earnings for the treatment group. We let the data speak to the functional form by including indicator variables for quintiles of the before-after difference.

We find evidence that before-after differences in labor market outcomes predict participant evaluations. The relationship is clearest for adult females and male youth, where the estimated coefficients increase monotonically (or almost so) and are statistically and substantively significant for the upper quintiles. Overall, the findings in this section lend support to the view that respondents implicitly or explicitly use natural and cognitively simple (but nonetheless quite biased) before-after comparisons in constructing their evaluations.

There are two potential reasons why some impact proxies correlate more strongly with participant evaluations than others. First, individuals may focus particular attention on (for instance) employment status changes when attempting to construct counterfactuals as lay scientists. Employment status changes then correlate with participant evaluations for this reason. Alternatively, some “poor” proxies may be less poor than others. Whether the particular proxies individuals use to form their survey responses better predict program impacts than the ones they do not has important implications for our lay scientist interpretation. In an analysis not reported in detail here, we find that the proxies that best predict impacts do not best predict participant evaluations, which is consistent with participants acting as lay scientists.<sup>26</sup>

## **7.0 Conclusions**

Broadly speaking we have two main findings. The first is that participant evaluations by treatment group members from the JTPA experimental evaluation have, in general, little if any relationship to either experimental impact estimates at the subgroup level or to what we regard as relatively plausible econometric impact estimates based on percentile differences. The second is

---

<sup>26</sup> See the online appendix for full details on the analysis and its results.

that the participant evaluation measures do have consistent relationships with crude proxies for impacts, such as measures of service type (a proxy for the resources devoted to the participant), labor market outcome levels (which measure impacts only if the counterfactual state consists of no employment or earnings, which it does not for the vast majority of our sample), and before-after comparisons (which measure impacts only in the absence of the “dip”).

Taken together, these two findings provide strong support for the view that respondents avoid the cognitive burden associated with trying to construct (implicitly or explicitly) the counterfactual outcome they would have experienced had they been in the control group and thus excluded from JTPA. Instead, they appear to act as lay scientists, using readily available proxies and simple heuristics to conclude, for example, that if they are employed at the time of the survey or if their earnings have risen relative to the period prior to random assignment, that the program probably helped them find a job. At the same time, our evidence does not rule out the view that respondents consider factors in their answers not captured in our experimental and econometric impact estimates, such as expected impacts in later periods (“employability”) or subjective and/or direct costs and benefits associated with the services they received. The proxy variables still leave much variation in the participant evaluation measure to be explained by other factors.

We borrow our “lay science” interpretation of our results from a large literature in social psychology on the fallibility of self-reports. The “study skills” experiment of Conway and Ross (1984) is the most parallel study we know of. Conway and Ross recruited subjects from one large introductory psychology course for a three-week study skills class, and randomly assigned them to either the class (treatment) or a waiting list (control). Both groups gave self-reports on their own study skill proficiency both before and after the three-week class. Since subjects came from one course, and the experiment took place between the midterm and final in that course, comparable outcome measures (in the form of grades on the midterm and final in the same class) were available to Conway and Ross, as were overall semester grades collected from registrar records. The objective measures confirmed what professional evaluators of such classes, e.g. Gibbs (1981), have found: the class had no substantively or statistically significant effect on outcomes. Yet treatment subjects reported greater improvement in their study skills, and expected better grades, than did control subjects. Conway and Ross interpret these results as

showing that their subjects act as lay theorists and rely on a theory that a study skills class will improve study skills.

Our findings indicate that participants behave as lay scientists. Byker and Smith (2012) and Calónico and Smith (2012) find similar results using data from other experimental evaluations with similar participant evaluation measures. Despite these disappointing findings, more sophisticated survey questions might do a better job of aligning participant and econometric evaluations. For example, improved performance may come from employing a question that explicitly asks the respondent to think about or construct a counterfactual and which makes the particular outcome of interest more concrete. Smith, Whalley and Wilcox (2012) review the types of participant evaluation questions in extant surveys and make detailed suggestions for improvements along these lines. For the present, we conclude that when they rely on survey measures like the one examined here, participant evaluations do not provide a plausible substitute for econometric estimates and should not be presented as such when evaluating programs. Ross (1989, p. 354) bluntly warns that “when self-reports are a primary indicant of improvement, a conspiracy of ignorance may emerge in which both the helper and the helped erroneously believe in the achievement of their common goal.”

University of Michigan, NBER and IZA

University of California-Merced, NBER and SKOPE

ESI, Chapman University, CEAR, Georgia State University



## References

- Bell, Stephen, and Larry Orr. 2002. "Screening (and Creaming?) Applicants to Job Training Programs: The AFDC Homemaker Home Health Aide Demonstrations." *Labour Economics* 9(2): 279–302.
- Bloom, Howard, Larry Orr, George Cave, Stephen Bell and Fred Doolittle. 1993. *The National JTPA Study: Title II-A Impacts on Earnings and Employment at 18 Months*. Bethesda, MD: Abt Associates.
- Byker, Tanya and Jeffrey Smith. 2013. "Chapter 6: Evidence from the Connecticut Jobs First Program." Unpublished manuscript, University of Michigan.
- Calónio, Sebastian and Jeffrey Smith. 2013. "Chapter 5: Evidence from the National Supported Work Demonstration," Unpublished manuscript, University of Michigan.
- Carrell, Scott and James West. 2010. "Does Professor Quality Matter? Evidence from Random Assignment of Students to Professors." *Journal of Political Economy* 118(3): 409-432.
- Conway, Michael, and Michael Ross. 1984. "Getting What You Want by Revising What You Had." *Journal of Personality and Social Psychology* 47:738-748.
- Devine, Theresa, and Heckman, James. 1996. "The Structure and Consequences of Eligibility Rules for a Social Program," in Solomon Polachek (ed.) *Research in Labor Economics Volume 15*. Greenwich, CT: JAI Press. 111-170.
- Di Tella, Rafael, Sebastian Galiani, and Ernesto Schargrotsky. 2007. "The Formation of Beliefs: Evidence from the Allocation of Land Titles to Squatters." *Quarterly Journal of Economics* 122(1): 209–241.
- Di Tella, Rafael, Sebastian Galiani, and Ernesto Schargrotsky. 2012 "Reality versus Propaganda in the Formation of Beliefs about Privatization." *Journal of Public Economics*, 96(2012): 553-567.
- Djebbari, Habiba and Jeffrey Smith. 2008. "Heterogeneous Program Impacts: Experimental Evidence from the PROGRESA Program." *Journal of Econometrics* 145(1-2): 64-80.
- Doolittle, Fred, and Linda Traeger. 1990. *Implementing the National JTPA Study*. New York, NY: Manpower Demonstration Research Corporation.
- Eyal, Yonatan. 2010. "Examination of the Empirical Research Environment of Program Evaluation: Methodology and Application." *Evaluation Review* 34(6): 455-486.
- Heckman, James. 1997. "Instrumental Variables: A Study of the Implicit Behavioral Assumptions Used in Making Program Evaluations." *Journal of Human Resources* 32(3): 441-452.

Gibbs, Graham. 1981. *Teaching Students to Learn*. Milton Keynes, UK: Open University Press.

Heckman, James, Carolyn Heinrich, Pascal Courty, Gerald Marschke and Jeffrey Smith (eds.), *The Performance of Performance Standards*. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Heckman, James, Carolyn Heinrich, and Jeffrey Smith. 2002. "The Performance of Performance Standards." *Journal of Human Resources* 37(4): 778-811.

Heckman, James, Carolyn Heinrich and Jeffrey Smith. 2011. "Chapter 9: Do Short Run Performance Measures Predict Long Run Impacts?" in Heckman, James, Carolyn Heinrich, Pascal Courty, Gerald Marschke and Jeffrey Smith (eds.), *The Performance of Performance Standards*. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. 273-304.

Heckman, James, Neil Hohmann, Jeffrey Smith, with the assistance of Michael Khoo. 2000. "Substitution and Drop Out Bias in Social Experiments: A Study of an Influential Social Experiment." *Quarterly Journal of Economics* 115(2): 651-694.

Heckman, James and Jeffrey Smith. 1998. "Evaluating the Welfare State," in Steiner Strom (ed.), *Econometrics and Economic Theory in the 20th Century: The Ragnar Frisch Centennial*. Cambridge University Press for Econometric Society Monograph Series, 241-318.

Heckman, James, and Jeffrey Smith. 1999. "The Pre-Program Earnings Dip and the Determinants of Participation in a Social Program: Implications for Simple Program Evaluation Strategies." *Economic Journal* 109(457): 313-348.

Heckman, James, and Jeffrey Smith. 2000. "The Sensitivity of Experimental Impact Estimates: Evidence from the National JTPA Study," in David Blanchflower and Richard Freeman (eds.), *Youth Employment and Joblessness in Advanced Countries*, Chicago: University of Chicago Press for NBER, 331-356.

Heckman, James, and Jeffrey Smith. 2004. "The Determinants of Participation in a Social Program: Evidence from the Job Training Partnership Act." *Journal of Labor Economics* 22(2): 243-298.

Heckman, James and Jeffrey Smith, with the assistance of Nancy Clements. 1997. "Making the Most Out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts." *Review of Economic Studies* 64(4): 487-535.

Heinrich, Carolyn, Gerald Marschke and Annie Zhang. 1998. "Using Administrative Data to Estimate the Cost-Effectiveness of Social Program Services." Technical report. The University of Chicago.

Holt, Charles and Susan Laury. 2002. "Risk Aversion and Incentive Effects." *American Economic Review* 92:1644-1655.

HRSD-Canada (Human Resources and Skills Development—Canada). 2009. Summative Evaluation of Employment Benefits and Support Measures in the Ontario Region. Canadian government document SP-AH-933-01-10E.

Jacob, Brian and Lars Lefgren. 2008. "Principals as Agents: Subjective Performance Measurement in Education." *Journal of Labor Economics* 26(1): 101-136.

Juster, Thomas. 1964. *Anticipations and Purchases*. Princeton: Princeton University Press.

Juster, Thomas. 1966. "Consumer Buying Intentions and Purchase Probability: An Experiment in Survey Design." *Journal of the American Statistical Association* 61:658-696.

Kahneman, Daniel, Paul Slovic and Amos Tversky. 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge University Press.

Kelly, George. 1955. *A Theory of Personality: The Psychology of Personal Constructs*. New York: Norton.

Kelly, Janet. 2003. "Citizen Satisfaction and Administrative Performance Measures: Is There Really a Link?" *Urban Affairs Review* 38(6): 855-866.

Kemple, James, Fred Doolittle, and John Wallace. 1993. *The National JTPA Study: Site Characteristics and Participation Patterns*. New York, NY: Manpower Demonstration Research Corporation.

Kornfeld, Robert and Howard Bloom, 1999. "Measuring Impacts on Employment and Earnings: Do Unemployment Wage Reports from Employers Agree with Surveys of Individuals?" *Journal of Labor Economics* 17(1): 169-197.

Kristensen, Nicolai. 2006. "What Do We Learn from Self-Evaluations of Training? A Comparison of Subjective and Objective Evaluations." Unpublished manuscript, University of Aarhus.

Manski, Charles. 1990. "The Use of Intentions Data to Predict Behavior: A Best-Case Analysis." *Journal of the American Statistical Association* 85(412): 934-940.

Manski, Charles. 1999. "Analysis of Choice Expectations in Incomplete Scenarios." *Journal of Risk and Uncertainty* 19(1-3): 49-65.

Manski, Charles. 2004. "Measuring Expectations." *Econometrica* 72(5): 1329-1376.

Nisbett, Richard and Lee Ross. 1980. *Human Inference: Strategies and Shortcomings of Social Judgment*. Englewood Cliffs, N.J: Prentice-Hall.

- Nisbett, Richard and Timothy Wilson. 1977. "Telling More Than We Can Know: Verbal Reports on Mental Processes." *Psychological Review* 84(3): 231-259.
- Oreopoulos, Philip and Florian Hoffman. 2009. "Professor Qualities and Student Achievement." *Review of Economics and Statistics* 91(1): 83-92.
- Orr, Larry, Howard Bloom, Stephen Bell, Winston Lin, George Cave and Fred Doolittle. 1994. *The National JTPA Study: Impacts, Benefits and Costs of Title II-A*. Bethesda, MD: Abt Associates.
- Philipson, Tomas and Larry Hedges. 1998. "Subject Evaluation in Social Experiments." *Econometrica* 66(2): 381-408.
- Polanyi, Michael. 1964. *Personal Knowledge: Toward a Post-Critical Philosophy*. New York: Harper.
- Ross, Michael. 1989. "Relation of Implicit Theories to the Construction of Personal Histories." *Psychological Review* 96(2):341-357.
- Schochet, Peter, John Burghardt and Sheena McConnell. 2008. "Does Job Corps Work? Impact Findings from the National Job Corps Study." *American Economic Review* 98(5): 1864-1886.
- Simon, Herbert. 1957. *Models of Man*. New York: Wiley.
- Smith, Jeffrey and Alexander Whalley. 2013. "How Well Do We Measure Public Job Training?" Unpublished manuscript, University of Michigan.
- Smith, Jeffrey, Alexander Whalley and Nathaniel Wilcox. 2013. "Chapter 7: Alternatives to Typical Participant Evaluation Questions." Unpublished manuscript, University of Michigan.
- Smith, Vernon. 1982. "Microeconomic Systems as an Experimental Science." *American Economic Review* 72(5): 923-955.
- USDOE (United States Department of Labor). 2005. Evaluation of the Teaching American History Program. USDOE Office of Planning, Evaluation and Policy Development, Policy and Program Studies Service.
- Wilcox, Nathaniel. 1993. "Lottery Choice: Incentives, Complexity and Decision Time." *Economic Journal* 103:1397-1417.
- Wolfers, Justin. 2007. "Are Voters Rational? Evidence from Gubernatorial Elections." Unpublished manuscript, University of Pennsylvania.
- Wood, Michelle. 1995. "National JTPA Study – SDA Unit Costs." Abt Associates Memo to Jerry Marsky [sic] and Larry Orr.

TABLE 1: Regression Results for the Relationship between Predicted Impacts and Participant Evaluations for Eight Outcomes, By Demographic Group

	Adult Males		Adult Females		Male Youths		Female Youths	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
<i>A. Employment:</i>								
Any Employment over 18 Months	-0.91 (0.32)	-0.21 (0.38)	0.10 (0.31)	0.47 (0.36)	-0.29 (0.31)	0.96 (0.60)	-0.54 (0.23)	-0.32 (0.80)
Employment in Month 18	-0.51 (0.31)	-0.15 (0.36)	-0.40 (0.40)	0.19 (0.29)	-0.31 (1.47)	-0.12 (1.12)	-0.20 (0.09)	-0.30 (1.06)
Any Employment (UI) over 6 Quarters	-0.34 (0.19)	-0.24 (0.32)	0.14 (0.28)	0.24 (0.34)	-0.69 (0.71)	0.00 (0.60)	0.29 (0.27)	0.65 (0.64)
Employment (UI) in Quarter 6	-0.39 (0.30)	0.03 (0.45)	0.04 (0.29)	-0.42 (0.32)	0.82 (0.50)	1.88 (1.00)	-0.01 (0.02)	0.72 (0.90)
<i>B. Earnings:</i>								
Earnings over 18 Months	-121.04 (134.85)	48.66 (83.41)	45.71 (85.10)	-16.86 (57.75)	-21.95 (244.01)	273.06 (214.97)	-208.51 (89.36)	24.82 (97.87)
Earnings in Month 18	-21.20 (20.65)	-6.05 (7.73)	2.37 (4.95)	0.97 (3.80)	35.53 (31.26)	-6.67 (16.58)	-2.32 (9.49)	1.54 (10.37)
Earnings (UI) over 6 Quarters	-63.67 (94.48)	-61.17 (85.90)	-85.43 (50.92)	14.51 (35.79)	-71.24 (133.03)	271.47 (134.34)	-103.53 (68.76)	78.86 (68.61)
Earnings (UI) in Quarter 6	-22.56 (23.25)	-14.52 (17.18)	0.73 (10.97)	-10.17 (7.65)	-0.16 (18.90)	80.59 (31.78)	2.60 (13.54)	-13.14 (12.60)
Positive (overall / 0.10 / 0.05)	0/0/0	1/0/0	5/0/0	5/0/0	2/0/0	5/3/2	2/0/0	5/0/0
Negative (overall / 0.10 / 0.05)	8/3/1	6/0/0	2/1/0	3/0/0	6/0/0	2/0/0	5/3/3	3/0/0

Notes: Source: Authors' calculations using the NJS data. Each cell in the table displays the coefficient estimate from a regression with the predicted impacts for an individual (based on their X) as the dependent variable and their participant evaluation as the independent variable. The regression is estimated using the experimental treatment sample. Heteroskedasticity-consistent standard errors appear in parentheses. The values in the bottom two rows are the counts of the

number of cells in the column above which are positive or negative, and counts of those that are significantly different from zero at the 10% and 5% levels. Specification (1) selects the set of  $X$  variables used to predict the impacts for each individual using a stepwise procedure. Specification (2) uses the same  $X$  as in Heckman, Heinrich and Smith (2002) to predict the impacts for each individual. Their  $X$  variables consist of race, age, education, marital status, employment status, AFDC receipt, receipt of food stamps and site.

TABLE 2: Relationship between Quantile Treatment Effects for 18-Month Earnings and the Fraction with Positive Participant Evaluations, By Demographic Group

	Adult Males		Adult Females		Male Youths		Female Youths	
	QTE	Fraction Positive	QTE	Fraction Positive	QTE	Fraction Positive	QTE	Fraction Positive
5 <sup>th</sup>	0 (0.90)	0.51 (0.03)	0 (0.38)	0.57 (0.02)	0 (1.15)	0.56 (0.07)	0 (1.08)	0.68 (0.04)
25 <sup>th</sup>	1233 (452)	0.66 (0.05)	501 (193)	0.63 (0.04)	-516 (515)	0.62 (0.08)	402 (193)	0.72 (0.06)
50 <sup>th</sup>	825 (608)	0.56 (0.05)	747 (416)	0.63 (0.04)	-1161 (681)	0.83 (0.06)	-39 (371)	0.71 (0.07)
75 <sup>th</sup>	8 (590)	0.72 (0.05)	938 (383)	0.78 (0.04)	-1261 (701)	0.68 (0.08)	-479 (566)	0.83 (0.06)
95 <sup>th</sup>	1589 (1323)	0.65 (0.05)	1910 (740)	0.70 (0.04)	-887 (1959)	0.81 (0.07)	-53 (1012)	0.64 (0.07)
Self-Evaluation Correlation	0.08 [0.750]	--	0.77 [0.000]	--	-0.45 [0.045]	--	-0.42 [0.065]	--
Self-Evaluation Coefficient	511 (1686)	--	5489 (1204)	--	-2232 (909)	--	-1576 (931)	--

Notes: Source: Authors' calculations using the NJS data. The values in the left column of the upper panel for each demographic group are quantile treatment effect estimates with standard errors in parentheses for five particular quantiles. The values in the right column of the upper panel for each demographic group are the means of the binary participant evaluation indicator variable for treatment group members in an interval of width 0.05 centered on the indicated quantile of the outcome distribution. The first row of the lower panel contains the correlations between the treatment effect estimates and the fraction with a positive participant evaluation by quantile (where one observation is one of the 20 quantiles); p-values for the correlations appear in square brackets. The second row of the lower panel contains coefficient estimates from regressions with the treatment effect prediction as the dependent variable and the fraction with positive participant evaluations as the independent variable (where one observation is one of the 20 quantiles). Heteroskedasticity-consistent standard error estimates appear in parentheses.

TABLE 3: Mean Numerical Derivatives and Test Statistics from Logit Models of the Determinants of Positive Participant Evaluation, By Demographic Group

	Adult Males	Adult Females	Male Youths	Female Youths
<i>A. Mean Numerical Derivatives:</i>				
Self-Reported Service: CT-OS	0.151 (0.027)	0.143 (0.022)	0.096 (0.039)	0.031 (0.033)
Self-Reported Service: OJT/WE	0.166 (0.039)	0.123 (0.034)	0.139 (0.057)	0.103 (0.051)
Self-Reported Service: JSA	0.095 (0.041)	0.024 (0.038)	0.014 (0.087)	0.039 (0.066)
Self-Reported Service: ABE	0.083 (0.045)	0.026 (0.037)	0.085 (0.048)	0.091 (0.040)
Self-Reported Service: Other	0.160 (0.053)	0.064 (0.047)	0.090 (0.080)	0.095 (0.060)
<i>B. Test Statistics:</i>				
Site	50.38 [0.000]	51.93 [0.000]	21.98 [0.079]	46.17 [0.000]
Age Category	5.55 [0.062]	23.08 [0.000]	0.25 [0.616]	1.47 [0.225]
Marital Status	1.26 [0.531]	3.49 [0.174]	1.22 [0.544]	7.11 [0.029]
Education Category	3.60 [0.464]	5.66 [0.226]	1.67 [0.645]	7.80 [0.050]
Race	1.11 [0.775]	1.74 [0.638]	8.62 [0.035]	7.87 [0.049]
English Language	1.04 [0.593]	2.93 [0.232]	1.07 [0.301]	0.20 [0.653]
Other Individual Characteristics	7.90 [0.162]	2.62 [0.759]	6.84 [0.232]	1.78 [0.879]
Self-Reported Service Type	39.20 [0.000]	43.36 [0.000]	8.62 [0.112]	7.05 [0.217]

Notes: Source: Authors' calculations using the NJS data. Columns two through five of the table report the results from a logit model where the participant evaluation indicator is the dependent variable and the categorical variables summarized in panel B column one are the independent variables. The values in panel A of the table are mean numerical derivatives, with the standard errors in parentheses, for the self-reported service type variables only. The values in panel B of the table are  $\chi^2$ -statistics for joint tests of the null that all of the coefficients equal zero for a given group of variables, with p-values in square brackets. The models are estimated using the experimental treatment sample. The variables in 'Other Individual Characteristics' are AFDC receipt, child less than six indicator, and worked for pay indicator. Indicator variables for missing values for the independent variables are also included in the regressions.



TABLE 4: Test Statistics from Logit Models of the Relationship Participant Evaluations and Labor Market Outcomes,  
By Demographic Group

	Adult Males	Adult Females	Male Youths	Female Youths
<i>A. Employment:</i>				
Any Employment During 18 Months	13.68 [0.000]	21.66 [0.000]	5.02 [0.081]	4.74 [0.093]
Employment in Month 18	7.93 [0.019]	15.29 [0.001]	5.49 [0.064]	1.33 [0.513]
Employment in the Month of the Survey	31.13 [0.000]	34.41 [0.000]	15.87 [0.000]	8.73 [0.013]
Any Employment During 6 Quarters (UI)	2.46 [0.292]	0.05 [0.973]	0.83 [0.659]	2.54 [0.281]
Employment in Quarter 6 (UI)	1.96 [0.376]	0.29 [0.865]	0.84 [0.360]	0.81 [0.666]
Employment in the Quarter of the Survey (UI)	31.13 [0.000]	28.07 [0.000]	16.42 [0.000]	8.05 [0.018]
<i>B. Earnings:</i>				
Earnings over 18 Months	24.27 [0.000]	42.61 [0.000]	18.39 [0.003]	13.75 [0.017]
Earnings in Month 18	16.81 [0.005]	35.47 [0.000]	8.97 [0.111]	2.24 [0.814]
Earnings in the Month of the Survey	35.19 [0.000]	51.80 [0.001]	23.34 [0.000]	12.33 [0.031]
Earnings over 6 Quarters (UI)	12.06 [0.034]	17.00 [0.005]	6.40 [0.270]	11.90 [0.036]
Earnings in Quarter 6 (UI)	14.07	14.02	4.53	8.41

Earnings in the Quarter of the Survey (UI)	[0.002]	[0.016]	[0.339]	[0.135]
	37.78	36.82	17.60	10.47
	[0.000]	[0.000]	[0.004]	[0.063]

---

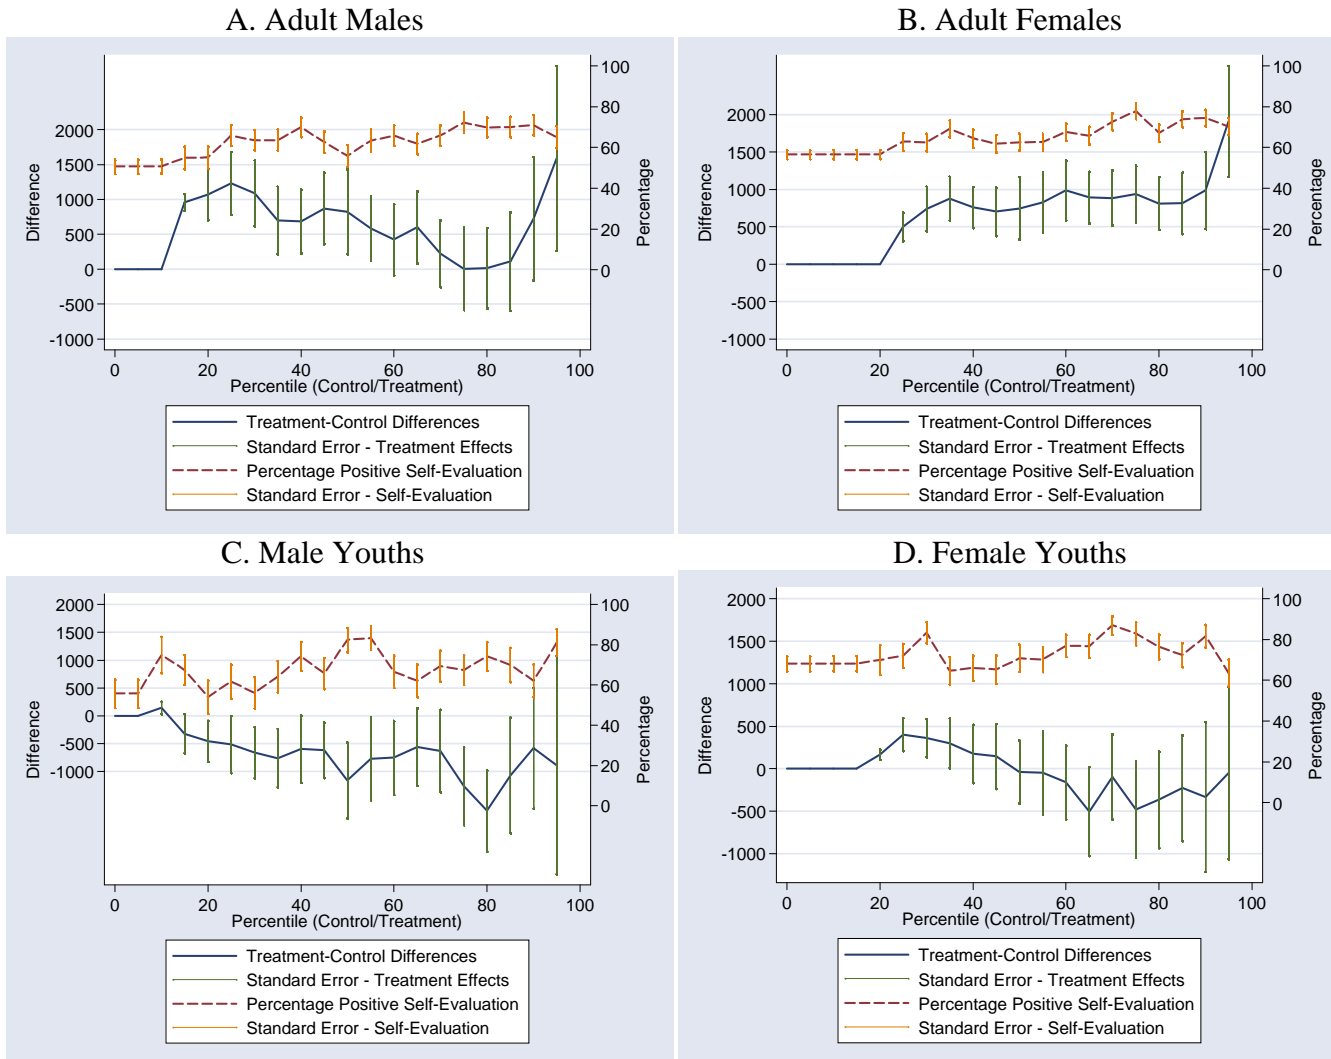
Notes: Source: Authors' calculations using the NJS data. Columns two through five of this table report the results from logit models where the participant evaluation indicator is the dependent variable. Independent variables include the categorical variables listed in Table 3 (with the exception of self-reported service type) as well as the indicated outcome variable. Each cell in the table corresponds to a different specification that includes a different outcome on the right-hand side. The values in the table are  $\chi^2$ -statistics for joint tests of the null that all of the coefficients equal zero for a given outcome, with the p-values in square brackets. The continuous earnings outcomes are coded into indicators for four categories: zero earnings, the lowest quartile of the non-zero earnings distribution, the lower middle quartile of the non-zero earnings distribution, and the upper middle quartile of the non-zero earnings distribution. The omitted category is the highest quartile of the non-zero earnings distribution. For the employment outcomes a binary variable is included indicating whether the respondent was employed or not. The models are estimated using the experimental treatment group sample. All models include indicator variables for missing values for the independent variables.

TABLE 5: Logit Estimates of the Relationship between Participant Evaluations and Labor Market Outcome Changes, By Demographic Group

	Adult Males	Adult Females	Male Youths	Female Youths
<i>A. Before-After Employment Status Changes:</i>				
Employed Before & Not Employed After	-0.084 (0.102) [0.413]	-0.068 (0.093) [0.465]	-0.512 (0.239) [0.032]	-0.272 (0.140) [0.052]
Not Employed Before & Employed After	0.042 (0.029) [0.138]	0.026 (0.024) [0.288]	-0.112 (0.046) [0.014]	-0.023 (0.035) [0.506]
Always Not Employed	-0.134 (0.060) [0.024]	-0.053 (0.042) [0.211]	0.033 (0.142) [0.819]	-0.077 (0.064) [0.231]
$\chi^2$ -Statistic:Joint test coefficients equal zero	12.09 [0.007]	5.21 [0.157]	9.76 [0.021]	4.80 [0.187]
<i>B. Before-After Self-Reported Earnings Changes:</i>				
Before-After Self Reported Earnings 2 <sup>nd</sup> Quintile	-0.035 (0.041) [0.393]	-0.020 (0.036) [0.576]	0.029 (0.058) [0.610]	-0.055 (0.058) [0.343]
Before-After Self Reported Earnings 3 <sup>rd</sup> Quintile	0.039 (0.038) [0.296]	0.054 (0.031) [0.082]	0.108 (0.050) [0.031]	0.030 (0.047) [0.523]
Before-After Self Reported Earnings 4 <sup>th</sup> Quintile	0.053 (0.038) [0.168]	0.079 (0.031) [0.010]	0.113 (0.049) [0.022]	0.071 (0.044) [0.106]
Before-After Self Reported Earnings 5 <sup>th</sup> Quintile	0.027 (0.034) [0.420]	0.113 (0.027) [0.000]	0.137 (0.046) [0.003]	0.052 (0.040) [0.193]
$\chi^2$ -Statistic:Joint test coefficients equal zero	6.23 [0.183]	28.07 [0.000]	10.69 [0.030]	7.96 [0.093]

Notes: Source: Authors' calculations using the NJS data. Employment status changes are based on changes in self-reported employment status measured at the date of random assignment and 18 months after random assignment. The omitted category is always employed. The estimates come from logit models with the participant evaluation indicator as the dependent variable and the before-after employment change variable and the categorical variables listed in Table 3 (with the exception of self-reported service type) as independent variables. The values in the table are mean numerical derivatives, with standard errors in parentheses and p-values in square brackets. The before-after earnings changes enter in the form of indicator variables for being in the 2<sup>nd</sup>, 3<sup>rd</sup>, 4<sup>th</sup>, and 5<sup>th</sup> quintiles of the before-after earnings change distribution. The omitted category is the 1<sup>st</sup> quintile of the distribution. All models are estimated using the experimental treatment group sample and include indicators for missing values of the independent variables.

FIGURE 1: Quantile Treatment Effects and Percentage with a Positive Participant Evaluation



Notes: Source: Authors' Calculations using the NJS data. The outcome used here is self-reported earnings over the 18 months after random assignment.