Planning for success: The interdisciplinary approach to building Bayesian models

Alex Dekhtyar , Judy Goldsmith , Beth Goldstein Krol Kevin Mathias Cynthia Isenhour

Abstract

This paper describes a process by which anthropologists, computer scientists, and social welfare case managers collaborated to build a stochastic model of welfare advising in Kentucky. In the process of collaboration, the research team rethought the Bayesian network model of Markov decision processes and designed a new knowledge elicitation format. We expect that this model will have wide applicability in other domains.

1 Introduction

This paper describes part of a research project devoted to studying stochastic planning with constraints in the context of welfare-to-work system. The project involved computer scientists, anthropologists and welfare case workers. We discuss the social interactions that led to a new way to model stochastic actions and briefly introduce the formal model and its elicitation process. The story as told centers on solving the computer science research problem. The paper is informed by ideas from within social science research on technology,

specifically the social construction of technology theory (SCOT) which examines the social factors that influence the construction and use of technology. [26,17].

This work grows out of the research program to build decision support for advising scenarios. In the context of this paper, the term "advising scenario" refers to an interaction of two people, in which one person, the advisor, is assigned the task of helping the other person, the advisee, achieve a specific goal by suggesting a series of possible actions/steps that the advisee should take. In our project we considered all such advising scenarios to involve uncertainty: when advice to take a specific action is given, the outcome of the advice is not determined. The advisee might act on the advice, or might ignore it. If she acts on it, her actions may succeed, with a variety of possible effects, or may fail, with equally undetermined effects. We model this by associating probabilities with possible outcomes of each action.

We model advice-giving in terms of factored Markov decision processes (MDPs), where decision variables represent the advisee and actions represent what we advise them to do. In 2000, the first two authors of this paper formulated the MDP-based advising scenario, using academic advising as the example. In 2003, in conversations with the third author it became apparent that the Welfare-to-Work assistance programs 1 present a more complex scenario. The main challenge lies in the need to combine stochastic planning with multiple constraints dictated by the rules and regulations of the program.

Under Welfare-to-Work assistance programs, a recipient, or client, meets with an assigned case manager, who negotiates a contract between them intended to move the client from welfare support into independent, paid employment. The client agrees to participate in certain activities, and the case manager authorizes support in various forms, including healthcare, childcare, transportation, school or training, and a stipend, with a 60-month lifetime cap on the total time a client may receive such support.

Since 2003, a team of computer scientists and anthropologists has worked together with case managers to build a formal stochastic model of the advising process in the Welfare-to-Work system. This paper details the modeling process and concentrates on the following aspects of it:

• The model elicitation process: in particular, the interactions between anthropologists and case managers, and between computer scientists and anthropologists.

When we use the phrase Welfare-to-Work in the paper, we refer to the U.S. welfare system's formal program Temporary Assistance to Needy Families (TANF). This federal program provides block funds to individual states with which to implement welfare assistance programs to families in need. The goal is to move parents in these households into independent, paid employment and thus off welfare assistance.

pologists.

• The model evolution process: our original approach to model advice in terms of 2-phase temporal Bayes nets (representing, in turn, factored MDPs) was shown to be inconsistent with the perceptions of the case managers about their advising process. As the elicitation progressed, the case managers' responses led the computer scientists to reformulate the model from 2-phase temporal Bayes nets to a new model that we call Bowtie Bayes net fragments. The latter model reflects directly the assertion by case managers that in order to predict changes in the state, they need first to know whether the action taken by the client has succeeded or failed. (It is worth noting that bowtie fragments are already in use in other contexts. In particular, Almond has used them in models of education [1].)

These two dimensions of this project's modeling process illustrate the social construction of technology (SCOT). According to SCOT theory, "technology design is an open process that can produce different outcomes depending on the social circumstances of development." [17]. Technology development is shaped by the particular people and groups of people engaged in its design: (1) through their different knowledge, meanings, ways of doing and values and (2) through the asymmetrical relations of power among these people. Using the SCOT lens to understand the modeling process enables us to recognize an emergent design process that is recursively informed by the multiple perspectives of a research team comprised of computer scientists, anthropologists and case managers. In other words, we can see how the technology that emerges happens as a result of negotiations across different ways of working and of understanding the work at hand. SCOT theory also asks that we attend to the consequences of a design team in which these three groups of people are not equal partners in the development process. Ultimately, the use of project resources, its rhythm and its eventual outcomes are driven by the computer scientists, given their more powerful location within the structures that inform this particular project (through virtue of the NSF funding designated specifically to solve a decision-theoretic research problem using MDPs and constraint solvers).

The original plan for the research project was both overly optimistic and computer science-centric. The computer scientists intended to have the anthropologists elicit dynamic Bayes nets, which the computer scientists would parse into the desired format. The anthropologists were invited to elicit client state variables from the case managers. The computer scientists' research agenda was to develop fast MDP solvers that bootstrapped on constraint solvers. With the MDPs, constraints, and solvers, software would be able to generate plans of action for case managers to share with their clients as ready-made plans or as choices among different options created by varying client preferences.

Nothing in the original research plan turned out to be easy. The research

team is still working on MDP policy displayers that have actual explanatory power to their users, who are not specialists in MDPs. The team learned that all elicited information rules and regulations, available actions and resources, client and case manager preferences are subject to change. The computer scientists had to learn to listen to the anthropologists and through them to the case managers. The anthropologists had to learn to translate language and knowledge between case managers and computer scientists. And the research team had to learn together how to negotiate, problem-solve across different modes of reasoning, and resist premature closure to the emergent technologies. What follows presents the challenges and ah ha moments of engaging technology as socially constructed mainly from the perspective of the computer scientists.

Section 2 provides necessary decision-theoretic background. In Section 3 we briefly outline the Welfare-to-Work domain, and our initial approach to decision-theoretic planning in it. Section 3.3 describes the first round of the elicitation process, in which the case managers have rejected our preconceived model of their advising. Section 4 focuses on the subsequent revision of the model and the second round of elicitation, which, in our view, succeeded. We offer qualitative comparisons between the results of two rounds, because we lack data quantifying the first, unsuccessful round. Our qualitative analysis, however, documents a change in both the attitude of the case managers to the process and the tasks they were asked to perform, and the nature of the outcomes in the two rounds. We offer these experiences as the main contribution of the paper.

2 Background: Mathematical Models for Decision-Theoretic Planning with Uncertainty

A Markov decision process (MDP) [2] is a model of a controlled stochastic process $M = \langle S, A, t, r \rangle$, where S is a finite set of states, A a finite set of actions, $t: S \times A \times S \to [0,1]$ a transition function, where t(s,a,s') is the probability that, if action a is taken in state s then the process will be in state s' at the next time step. The utility function r can be defined as $r: S \to \mathbf{R}$ or $r: S \times A \to \mathbf{R}$, or even $r: S \times A \times S \to \mathbf{R}$. The latter would assign a utility to being a state s, taking action a, and ending up in state s'. The first two definitions of r are more common.

Note that not all definitions of MDPs assume that S is finite or that time is discrete.

MDPs are used to model planning under uncertainty. The "solution" of an MDP is a policy that specifies what action to take under any circumstance.

That could be a function from states to actions, or perhaps from states cross time to actions. A policy is optimal if it has the best expected utility over time. There are algorithms that run in time polynomial in |S| + |A| for finding optimal policies if the system will run forever but future utility is discounted relative to current utility. If the process will run for a fixed, finite number, h of steps, there are algorithms for finding the optimal policy that are polynomial in |S| + |A| + h (assuming h is written in unary).

Many interesting controlled stochastic systems have enormous state spaces. One way to represent the state spaces is by factoring it, so that $S = \Pi_i S_i$, where the S_i s are salient features of the states. For instance, the features of a welfare client include gender, age, number of dependents, education level, job readiness, etc. One can observe that a given action will deterministically affect certain characteristics (age, for instance), stochastically affect others, and not affect many (such as gender, usually).

A Bayesian network, or *Bayes net* [25], is a directed acyclic graph, where each edge represents a dependency, and each node in the graph has an associated probability table. If a node has in-degree 0, then the probability table is a simple probability distribution over possible values the feature represented by that node can take. If the node has parents, then it has a conditional probability table, where its probabilities are conditioned on the values of the parent nodes.

A Bayes net can be used to represent a factored MDP in several ways. The most common is a 2-phase temporal Bayes net (2TBN) [3], which is a particular type of dynamic Bayes net. For a factored MDP with n state features $\{s_1, \ldots s_n\}$, the corresponding 2TBN has 2n nodes, $\{s_1, \ldots s_n s'_1, \ldots s'_n\}$, representing the features at times t and t+1, respectively. An edge from s_i to s'_j represents the fact that the value of feature s_j at time t+1 can be predicted stochastically from the value of s_i at time t, as well as from the values of its other parents. An example can be seen in Figure 1.

In fact, a 2TBN model of a MDP consists of one 2TBN for each action. These can be graphically combined into one decision diagram, at the cost of losing visual interpretability.

We next present the formal definition of the new dynamic Bayes net model introduced in this paper. Subsequent sections of the paper describe the development of the model and its use in information elicitation.

A bowtie fragment is a 3-phase temporal Bayes net, where the middle phase consists of a single node representing the outcome of an action. The outcomes discussed in this paper are "success" and "failure," but the model can be used with multivariate outcomes as well.

There are two graphs associated with a bowtie fragment: the implicit and the explicit graph. The explicit graph is used for elicitation, and shows both the variables that influence the probability of the different action outcomes and the variables affected by the outcome. Figure 2 gives such a graph (discussed in more detail below). The effects of the success or failure of an action are described in the elicitation in terms of values such as self-esteem or literacy increasing or decreasing, implicitly referencing the previous value.

The implicit graph is the actual dependency structure for a Bayesian network based on the first graph; the implicit graph shows the connection between a variable at time t and its representation at time t+1, even if that variable is not thought to influence the success of the action.

We overload the term "bowtie fragment" to indicate both the semi-qualitative networks we have elicited and their quantitative counterparts. The quantitative counterparts have conditional probability tables computed from the weights along edges 2 .

The quantification process involves four steps. The first is to take individual input edge weights and transform them into one-parent conditional probability tables, according to a fixed set of possible tables. Next, the tables are combined to give the conditional probability of the success node. This uses a NOISY MAJORITY function (defined below). The success node's outgoing edge weights are similarly transformed to one-parent CPTs. Finally, for each potential outcome variable in the domain, a *combination method* [19] is specified. Combination methods define how to compute outcomes if multiple actions are taken concurrently, and may include deterministic maximization, for instance for handling a node that represents the passage of time, or average, NOISY MAX, or NOISY MAJORITY.

All Noisy combination methods [24] work in a similar manner. Given a node S in a network, and its parent nodes V_1, \ldots, V_n , the combination method does two things. For each combination s_1, \ldots, s_n of predicted outcomes of the value of S induced by individual influences of V_1, V_2, \ldots, V_n on S, the combination method produces $outcome(s_1, \ldots, s_n)$, the combined predicted value of S. Once the outcome() function is established, the probability $P(S = s|V_1, \ldots, V_n)$ is computed as

$$P(S = s|V_1, \dots, V_n) = \alpha \sum_{outcome(s_1, \dots, s_n) = s} P(S = s_1|V_1) \wedge \dots \wedge P(S = s_n|V_n),$$

The procedure for building conditional probability tables is discussed in detail in [20]. The process involved determining what constraints needed to be satisfied by the CPTs, and then selecting specific probability numbers that satisfied these constraints. The selected CPTs were validated with humans experts with good results.

where α normalizes the probabilities, namely, $\frac{1}{\alpha} = \Sigma_s P(S = s | V_1, \dots, V_n)$.

The more common combination methods, NOISY MAX and NOISY MIN, appear to be ill-suited for our application. NOISY MIN is too pessimistic, as it requires all parents of a node S in a network to agree on a higher— or better—value for S. This seems to underestimate the real probability of success: if four factors out of five predict success, while the fifth factor predicts failure, it seems like a stretch to immediately conclude that the action fails. The flip side of this is the optimism of NOISY MAX, which declares success if at least one of the influences (parents) of a node S predicts success.

The Noisy Majority combination method has been crafted to better balance the failure-to-success ratio by using majority rule to predict the outcome of a node S.

Informally, one can think of each parent, V_i of the success node as flipping a biased coin to decide if the value of S is *success*. If the majority agree on the value *success*, then it is so. The probability of *success*, then, under NOISY MAJORITY is roughly the sum of the probabilities of majority agreement on *success* for all possible ways that the majority can agree. Since there are some cases where there is no majority agreement at all—for instance, if there are two parents and they disagree—so we must normalize the probabilities.

More formally, consider each parent, V_i , as an expert, and the (one-parent) conditional probability of success as their prediction, given their current value, v_i . Then we can define, for a given set of input values, the probability that a majority of the experts/parents predict the value success".

Consider a node S with $dom(S) = \{failure, success\}$ which has three parents V_1, V_2, V_3 , each with domain $\{a, b, c\}$. Consider the state where $V_1 = a$, $V_2 = b$ and $V_3 = c$. We are interested in constructing P(S = success|a, b, c) from $P(S|V_1 = a)$, $P(S|V_2 = b)$ and $P(S|V_3 = c)$. NOISY MAJORITY will predict that S = success as long as at least two out of three individual predictors predict S = success. There are four such possible predictors: (success, success, success), (success, success, failure), (success, failure, success) and (failure, success, success). The overall probability P(S = success|a, b, c) is then computed as follows.

```
P(S = success | a, b, c) = \\ Pr(S = Success | V_1 = a) \cdot Pr(S = Success | V_2 = b) \cdot Pr(s = Failure | V_3 = c) \\ + Pr(S = Success | V_1 = a) \cdot Pr(S = Failure | V_2 = b) \cdot Pr(S = Success | V_3 = c) \\ + Pr(S = Failure | V_1 = a) \cdot Pr(S = Success | V_2 = b) \cdot Pr(s = Success | V_3 = c) \\ + Pr(s = Success | V_1 = a) \cdot Pr(s = Success | V_2 = b) \cdot Pr(s = Success | V_3 = c) \\
```

Note that this is well-defined for any number of inputs.

Although a bowtie can be transformed into a 2TBN, a straightforward transformation loses the correlation between the different effects of the success node. Thus, the bowtie models appear to call for new MDP planners that can take advantage of the extra information inherent in the bowtie knots. (Almond [1] shows ways in which this information can be used.)

3 Decision-theoretic Planning for Welfare-to-Work

The project described in this paper originated with an observation by the AI group at the University of Kentucky, that, individually, stochastic planning and constraint satisfaction are well-studied topics, but stochastic planning in the presence of constraints on the domains and actions is an open area of investigation.³

As a domain for stochastic planning with constraints we have considered advising settings. In such a setting one human agent, the advisor, is charged with suggesting to another human agent, the advisee, a plan of actions. In coming up with a long-term plan, the advisor has to base her decisions on three sets of criteria: (a) the perceived stochastic effects of the actions taken on the "state" of the advisee; (b) constraints on which actions, or action combinations, can be taken under which circumstances, and (c) the preferences stated by the advisee.

3.1 Introduction to the Welfare-to-Work domain

In 1996 the US legislature signed the Personal Responsibility and Work Reconciliation Act (PRWORA). This act, heralded by then-President Bill Clinton as the program that would "end welfare as we know it" was built upon the popular assumption that welfare programs had enabled apathy, dependence, a poor work ethic and abuse of the system among the nation's poor. ⁴ PRWORA

³ Some MDP solvers do handle constraints. These include linear programming-based solvers that can include any linear constraints (see, for example, [7] for an example solver for factored MDPs). There are solvers which directly convert the MDP and its constraints into a constraint satisfaction problem (see [5] for an example). However, we expect that solvers for MDPs with constraints will be a growth area in AI in the next few years.

⁴ The overwhelming majority of welfare recipients are single mothers with sole responsibility for raising their children and maintaining households. Unpaid domestic work including dependent care is not considered "work" by most states under this

placed the responsibility for poverty on individuals rather than on structural, social and economic inequalities [11,6,15]. To alleviate the perceived substandard work ethic among the poor, PRWORA set a 5 year lifetime limit on welfare benefits for all recipients and mandated that welfare clients work or participate in work readiness, education or training programs in order to receive benefits. (There are exceptions to these rules for individuals in extreme circumstances.)

These changes in welfare legislation significantly restructured the work of welfare [23]. Case managers who were once responsible for determining eligibility and processing cash assistance payments by means of established formulas, under TANF suddenly became accountable for informing clients about work and work-related program requirements, assisting clients in the discovery and/or definition of career goals, and helping individuals to match their preferences, abilities and goals to a long list of "countable" activities. ⁵

Case managers in Central Kentucky, where our research program is based, must process information about myriad training, support, employment and educational programs available to their clients including information about prerequisites, schedules, locations and content. In the central Kentucky city of Lexington alone, more than 200 agencies offer support services to welfare clients. Case managers develop action plans that fit client needs for these services and suggest agencies that are appropriate in terms of location, schedule, etc.

In order to help clients fulfill their needs, case managers must also be familiar with the goals, preferences, abilities, constraints and interests of their clients. The case managers from whom knowledge of the WtW process was elicited handled between 40–80 active cases at a time. Even with rigorous documentation and a good memory, it is difficult for the case managers to recall the unique needs and characteristics of each client. In addition, case managers must stay abreast of changes in welfare regulations, policies and rules as shifting budget allocations dictate changes in policy execution, as well as service availability.

While the 28 case managers participating in our study express a strong desire to do their absolute best for their clients, many admit that it is difficult to keep up with the interests, abilities and preferences of a constantly changing list of unique clients while simultaneously managing information about dynamic

legislation.

⁵ Countable activities are those which can be counted towards the fulfillment of federal TANF participation quotas.

⁶ One of the less research-oriented, but more tangible, contributions of our project to the work of the case managers has been the development and deployment of a database application to track available services.

services and policies. In this context, our research team of computer scientists, social scientists and domain experts (case managers) is working to build decision support software for the welfare domain. These efforts can help to manage the information load carried by case managers. By improving client services, the work may also help the nation's poor and their case managers develop plans for self-sufficiency that more accurately build upon client interests.

3.2 Planning for Welfare-to-Work domain

The increased burden on the case managers can potentially be reduced by introducing decision support/planning software into their work. As mentioned in Section 3.1, a key feature of the TANF system is the focus on providing services to clients to support their move into employability. These services can be roughly partitioned into two categories. Services in the first category are performed to alleviate barriers a client might have, preventing her from participation in the rest of the program. Such services may include subsidized housing, health care, child care, transportation allowances, help with basic coping strategies, etc. The services in the second category are the activities for a client to (a) remain eligible for TANF assistance and (b) become employable and employed. Such activities include volunteering opportunities, literacy training, high school equivalency or college classes, professional training, English as a second language classes, job search and interview preparation seminars and more.

Case managers are entrusted with advising their clients on activities, which, in their opinion, advance the client towards the general goal of employment. Each such action has the potential to change the client's state. These changes are uncertain and can be modeled stochastically. The action space and the current information available about the client are factored.

The Welfare-to-Work system operates under a wide array of federal, state and local rules, regulations and resources. These supply a rich set of constraints, from the 60-month limit on benefits over an individual's lifetime to soft constraints on "countable" activities (those that go toward meeting the case manager's and the agency's federal participation quotas) and "allowable" activities (those permitted by the state but not counted in the federal quotas). There are also logistical constraints. For instance, a client who relies on public transportation must begin and end activities while public transit is running, and must be able to reach those activities using public transportation.

Client preferences also play a role in determining courses of action. Even if certain activities may be beneficial to a client, she may want to forego them (e.g., a client has the potential for a career in health care, but has a strong

aversion to blood).

We illustrate the case manager-client interaction on the following fictional case.

Example 1 A 21-year-old woman with a 4-year-old son and a 2-year-old daughter has completed 11th grade, lives in a government-subsidized apartment complex, and has been unable to seek work. The barriers to her participation in services are her lack of childcare and lack of transportation—her apartment building is not on a bus route and she doesn't have a car.

The case manager first addresses these barriers by providing transportation to an approved childcare site, and transportation for the client to an adult education center, to allow the client to prepare for her highschool equivalency exam. The long-term goal is a clerical job, with midterm goal of enrolling the client in secretarial school. There are two options for secretarial school. One offers evening courses, which are incompatible with childcare availability. Thus, constraints dictate that she attend the other school.

Another option would be to send this client immediately to car mechanic training. While the training is available and convenient, and this could lead to a high-paying job, the client is unwilling to deal with the prejudice against women she expects to find in the automobile repair world. The case manager determines, therefore, that this option has a significantly lower probability of success. She chooses not to pursue this option for this client.

In order to compare options such as secretarial school and car mechanic training, the case manager must assess the probabilities of each action's success, given the client's state, and the probable effects of both success and failure at each action. As mentioned above, we model these using dynamic Bayes nets.

There are three basic steps to building a Bayes net representation of an MDP model for any application: determining the key components of the domain, translating them into the components—variables, actions, dependencies, and probabilities—of the mathematical formalism, and validating the models.

In order to determine the key components, the anthropologists in our group used open-ended interviewing techniques with welfare professionals. Our translation process has three parts: (a) determining the variables and actions; (b) eliciting qualitative relationships amongst these components, and (c) determining quantitative relationships that are consistent with the elicited information. Validation includes using scenario-based questions to determine both the expected outcomes of actions, and the appropriateness of actions to particular states. Scenario-based questions can also be used to validate MDP planners for the domain.

We model a client's current situation as a factored state, services as stochastic actions, preferences as utility functions over possible states and actions, and regulations and limitations of clients as constraints. In stochastic planning, a policy specifies actions for all possible outcomes or states.

More specifically, we consider the factored MDP states to be formed by a number of *client characteristics*. We have identified a wide range of characteristics, including objective attributes such as the client's age, education level, number of children or disability status, and more subjective ones, such as the client's literacy and numeracy, self-confidence or commitment level. One action in the welfare domain is the client's participation in one of the services/programs, such as GED classes, volunteering or job interview skills seminars. Such participation may affect some of the client's characteristics. For example, a job interview skills seminar affects the client's self-confidence (which can go either up or down, depending on whether or not the client feels she is ready for the rigorous job interviewing process), commitment and work-readiness.

In the welfare domain, gathering empirical data about actual cases is extremely sensitive. Requests for personal information about the nation's most vulnerable populations are not taken lightly. Issues surrounding confidentiality and privacy require the informed consent of all welfare participants before their case records are released for research. Because of these issues, our data collection efforts concentrated on the elicitation of expert knowledge. We relied upon the professional and experiential knowledge of welfare case managers. Of the twenty-eight case managers participating in this project, nearly half of them have five or more years experience in this capacity.

Eliciting data from welfare case managers presents a unique set of challenges. The case managers have been trained on the job to look at their clients as individuals rather than numbers. While decision modeling requires quantifiable data, the case managers with whom we work often vehemently resist our attempts to gather generalized or abstracted data. These women and men consistently insist that it is difficult to generalize about their clients because each one is different, making it equally difficult to generalize about decision making patterns. In one attempt to elicit information about a specific activity (taking GED preparation classes), one case manager expressed her reservations with our efforts:

"I think it's really difficult to think about these issues individually ... it has to be much more holistic. I mean if you look at my list, I've got everything ranked as extremely important. Everything is extremely important and I don't think you can just rank the top five ... In my assessments, I'm

not going to just ask (my clients) for five pieces of information. It seems impossible to isolate these factors or to categorize them." (6/15/2005)

Another case manager, when asked to talk about the characteristics of a successful client, stated,

"I look at them as a whole person, they're all different. They're not a list of characteristics and not a way to build up my participation rate. That's not a good way to do cases ... at least I (stops and reconsiders)... most of us don't think so." (10/23/04).

Case managers prefer to speak in narratives, imparting tacit knowledge through stories of specific clients and their unique circumstances. These perspectives and preferences had a significant impact on our data elicitation methods. We had to be clear about the requirements of data modelling, stressing the need for simplification of a clearly complex decision environment. We also had to frame questions, statements and problems in meaningful language for the case managers. This often involved asking them to consider a specific scenario consistent with cases they have worked.

Our attempts to elicit information in the form of 2-phase temporal Bayes nets still failed. This model, or to be more exact, elicitation procedures based on it, were not intuitive for case managers. In addition to the general reluctance to specify the most important influences, the key problem lay in the fact that the 2-phase TBN model lacked the notion of the action's result, which was pivotal for case managers' understanding of their work.

Through discussions between the anthropologists and computer scientists, and some crucial translation by Russell Almond, we determined that we needed to elicit information from the case managers in terms of the success or failure of actions. From this understanding, we then determined that the best formalization of a success-based model is the MDP with results, also called *bowtie fragments* described in Section 2.

4 Elicitation of Models

Elicitation of information for construction of Bayesian models of advising in the WtW domain is central to this project. Originally, computer scientists proposed to represent activities (actions) a WtW client can take as two-phase Bayesian network (a 2TBN or DBN) [4]. Each activity, described as a DBN fragment, showed how various client characteristics were likely to change, based on their current state and completion by the WtW client of an action.

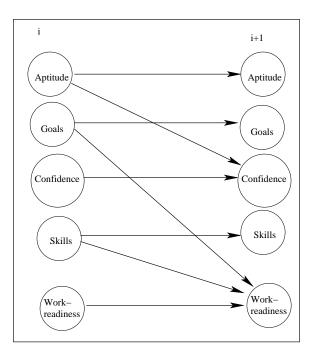


Fig. 1. An action in the Welfare-to-Work domain represented as a two-phase temporal dynamic Bayes net

This approach is illustrated in Figure 1, where a possible two-phase temporal Bayes network is shown for "Volunteer placement", one of the actions in the Welfare-to-Work domain. Here, the five client characteristics used in the fragment are considered to be the most crucially affected by this action. The new value for each of the five characteristics, *Aptitude*, *Goals*, *Confidence*, *Skills*, and *Work-readiness* is affected by its old value and, possibly, by a combination of values of some other characteristics.

From the beginning of the project, anthropologists worked with case managers from three agencies. Through multiple, iterative interviews, they established and conveyed to the computer scientists the main operational procedures and key regulations that guide the WtW program. They also discussed with the case managers their *modus operandi* to establish how, in general, case managers assessed the likelihood of a client's success in different activities. These data laid the groundwork for the elicitation process. Appendix A contains a more detailed description of the pilot elicitation process.

The computer scientists and anthropologists agreed to limit the initial scope of the elicitation process from the case managers to qualitative information: case managers would be asked to help build graphical representations of the stochastic model, but would not be asked direct questions about probabilities. Because the anthropologists observed that the case managers thought in terms of narratives, rather than in terms of statistics, we decided that it was not appropriate to ask them to describe possible changes in terms of probabilities.

The computer science group had the following questions for the case managers:

- (1) What are the various client characteristics that play an important role in your decision-making?
- (2) What are the different activities you recommend to your clients?
- (3) What are the most important client characteristics for each activity?
- (4) How does each activity affect client characteristics?
 - (a) Which characteristics affect which other characteristics?
 - (b) How strong are the individual influences?

The first two questions were designed to establish the basic parameters for the model being built: the domain for the client states and the set of actions used for planning. The third question (actually, family of questions) addressed the computer scientists' need for reasonably small Bayes net fragment representations for the effects of each action. The last set of questions would establish the shape of the 2-phase temporal DBN for each action, and provided information for the eventual automated quantification process.

The preliminary interviews held by the anthropologists with the case managers successfully established the answers to the first two questions. Sixteen actions were elicited from the case managers. One action "Take GED preparation classes" was used in the initial elicitation, and the remaining 15 actions were used in the follow-up elicitation study.

One of the surprises for the computer scientists was the rejection by the anthropologists of the two-layer DBN fragment as the model of actions. When the anthropologists asked the case managers about which client characteristics influence their decision to recommend a specific action, as well as the expected change of client characteristics, case managers refused to answer, explaining their refusal in two ways. First, case managers insisted that they could not responsibly make generalizations based on "generic" clients. Their vivid experiences with clients made it hard to hypothesize about actions and outcomes in the presence of a client described only by a list of characteristics. Asking "what would you advise a 24-year old mother of two who lives in an apartment complex, lacks transportation, has a high-school diploma, but has no work history and has a history of alcohol abuse?" turned out to be a wrong type of question—too decontextualized for case managers to be able to give answers.

In general, case managers agreed that the actions their clients take affect their "state". What they did not agree with was the idea that the mere act of taking an action changes that state, as implied by the DBN model structure. One missing piece, in the opinion of the case managers, was the *result of the action*, i.e., success or failure. A clear outcome from pre-elicitation interviews with case managers was the necessity of representing the success of an activity

explicitly.

The goals and objectives of case managers are tied directly to helping a client succeed in a given action. According to case managers, a client's success or failure in an activity has a profound impact on the client's state. This in turn affects the client's likelihood of success in future actions. For example, if a client succeeds in earning a high school equivalency degree (GED), the client's confidence and motivation for further education will increase, and conversely if she fails at the GED. The DBN model did not represent the transformation between two client states based on the *explicit outcome* of the client's participation in an activity.

To address the concerns of case managers and to facilitate knowledge elicitation from them, computer and anthropologists jointly developed a new class of stochastic models. We call these models *Markov decision processes with actions that have results*, and represent each activity in the model by a *bowtie action fragment* [21]. The new models contain, for each activity, a *success node*, a random variable explicitly quantifying the client's performance in, or level of success in completing, the activity. The success node becomes the central node (the "knot") of the bowtie. Client characteristics from the current state influence the success node. The success node, in turn, influences the client characteristics after completing the activity. This use of a central success node creates the pinched shape that we have read as a bowtie.

We illustrate the use of the bowtie fragments, and the stochastic models they represent, on the following example. Figure 2 shows a bowtie fragment elicited for the Volunteer Placement action. Here, the same five client characteristics as in Figure 1 are used as input nodes. However, unlike the two-phase temporal DBN, in the bowtie fragment, these characteristics affect just one random variable, the success variable for the action (here Volunteer Placement Performance). In this figure, the success variable is shown significantly to affect three client characteristics: Aptitude, Goals and Income. Note that the list of input nodes in the bowtie fragment and the list of output nodes are different. Some input nodes, such as Skills, are assumed to affect significantly the result of the Volunteer Placement activity. However, this result will not have significant effect on the client's skills. On the other hand, while the client's performance in this activity is not affected by her income, her success could lead to a change in the income.

It is important to notice that the bowtie fragment is NOT a full-fledged Bayes network, rather, it is a graphical structure representing a Bayes network fragment. The actual shape of the Bayes network fragment represented by a bowtie is somewhat more complex. Figure 3 shows the Bayes network fragment represented by the Volunteer Placement bowtie fragment from Figure 2. First, we note that a new input node, Income, has been added to the network, but NOT

connected to the *Volunteer Placement Performance* node. Next, we note that according to this Bayes network fragment, the new value of each output node depends on two other nodes: the node representing that variable before the action was taken or attempted, and the outcome of the action. The exact dependence is determined by the weight that was elicited for the corresponding outcome edge. Positive weights scale up the likelihood of a successful action leading to improvement (all variables are assumed to have ordinal domains, with higher values assumed to be better—more education, more self-esteem, etc.) The degree to which values scale up (or down, if the weight is negative or the action fails) depends on the weight for that outcome edge.

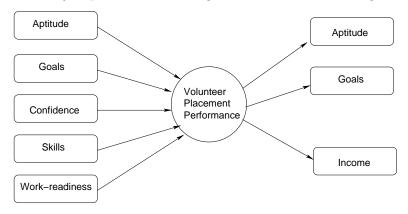


Fig. 2. The bowtie model of the action "Volunteer Placement".

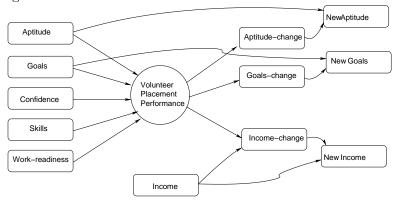


Fig. 3. The Bayes network represented by the bowtie model of the action "Volunteer Placement".

Case manager knowledge was elicited in three stages: (i) a manual pilot study carried out by the anthropologists, (ii) design and implementation of elicitation software, and (iii) software-directed elicitation. We briefly outline these stages below.

First, our team of anthropologists conducted a pilot study to test the bowtie elicitation methodology. In the pilot study, welfare case managers were asked to free-list client characteristics which would affect the client's likelihood of success in the action GET GED, which includes attending preparatory classes and eventually taking the General Educational Development (GED) test.

The welfare case managers listed nearly 30 characteristics that would affect the client's likelihood of success in the action GET GED. These 30 characteristics were then evaluated and ranked by small groups. The five client characteristics cited by case managers as most important were weighted, thus forming the input structure of the bowtie model for the GET GED action.

Based on iterative interviews with case managers and the pilot elicitation, the anthropologists generated a list of 200 client characteristics. In order to get more precise information on the action models, we decided to merge the list of 200 characteristics into a more manageable list of approximately 50. We also loosely categorized these variables as education-related, work-related, and personal characteristics. Within each category, we tentatively outlined subcategories.

The initial elicitation produced a description of the GET GED that corresponded to the anthropologist's understanding of the success-predictors and effects of that action. The computer scientists then built a high-level elicitor (HLE) to elicit the inputs and outcomes for the remaining 15 actions.

The High Level Elicitor (HLE) was designed for the specific use by case managers in a one-time elicitation experiment. Since then, HLE has evolved to include maintenance and management of already elicited bowties as part of its functionality. The version of HLE used in the elicitation experiment presented a clean, straightforward GUI for the bowtie elicitation process. The elicitation process was broken into several steps, and each step corresponded to a new screen, which (a) detailed the work done thus far in the process, (b) outlined the current task for the case managers and (c) provided simple, easy-to-use ⁷ GUI for the task at hand. In the experiment, each case manager was asked to provide information about five different actions. For a single action fragment, HLE elicitation proceeded as follows.

- Step 0. The tool informed the case manager of the action for which information was about to be elicited.
- Step 1. After proceeding to the next screen, the case manager was offered a list of client characteristics and asked to name *five* which in their opinion have the strongest influence on the client's performance in the action.
- Step 2. The next screen asked the case manager to specify the strength of influence for each of the five selected characteristics on the scale from 1 to 4 (1 = "reasonably significant influence", 4 = "extremely significant influence").
- Step 3. The following screen asked the case manager to select how client's *success* in the given action would affect client characteristics. For each client

⁷ Our judgement of the HLE interace as easy to use is based on the success the case managers had in using the tool to record their opinions and impressions, as measured in their satisfaction with the process.

- attribute the case manager could specify one of three choices: "will likely decrease," "will have no effect" and "will likely increase.?
- Step 4. This step repeated the activities of **Step 3**, only for the situation when a client *fails* the action.
- Step 5. The final screen of the elicitation process displayed a short verbal summary of the information submitted by the case manager, asked her to review it and either agree or return back to one or more previous steps and change the information provided.

5 Outcome of Elicitations

We were delighted to see that the case managers substantially agreed with each other. There were, however, a few significant outliers among the case managers. These turned out to be individuals who had worked for a long time within the system, and had longer-term views of the process. They saw success in almost any action as affecting almost all aspects of the client's life. For instance, success at taking the GED (high school equivalency) exam affected marital status. This long-term view of the process is discussed in [14].

When we understood the different interpretations of our questions, we returned to those case managers who had given long term responses and explicitly asked for short-term outcomes. Their responses about the short term proved consistent with those from other domain experts.

Once we had a consistent set of HLE outputs, we used techniques based on the work of Renooij, van der Gaag, Druzdzel, and Henrion [28,27,8,9] to build quantitative Bayes net fragments consistent with those outputs.

Unfortunately, the off-the-shelf factored MDP planner [12] that we expected to use to generate client plans was unable to handle the model that we built, for a variety of reasons. The standard MDP model only allows one action at a time, and the set of all possible combination actions is too large to represent explicitly; the standard MDP solvers expect binary variables, and almost all of our client characteristics have more than two values; the benchmark factored MDPs have very simple utility functions, and the complex utility functions we wished to consider make computations infeasible. We are still working on all these aspects of the MDP solving process.

To computer scientists, a process is a thread, whereas to anthropologists, a process is a weaving together of threads. We entered this project knowing—as individuals—how collaborative research was conducted. Unfortunately, we had no *common* model of the collaborative process. As acknowledged throughout this paper, our work required both technical and social research. Perhaps because of their focus on the technical aspects, the computer science contributors believed they could work on independent technical chunks that would eventually be brought together and retooled to work with each other. The anthropologists, in contrast, saw the social and technical research as interwoven such that one could not proceed completely independent of the other. Cooperation required ongoing conversations yet we did not have one language in common.

The team learned the importance of having social and computer scientists working together on the software development life cycle (SDLC) of software and solutions for the WtW project. Anthropologists were able to translate the needs of computer scientists into a language that made sense to the case managers, given their own perspectives, needs, and interests—and vice versa. This helped immensely in the requirements gathering phase of the SDLC of the model building/elicitation software, i.e., HLE. It was evident to the team that the intended users, the case managers, were more likely to be responsive to software programs that directly addressed their needs and desires. Any software that needed case managers' participation had to be built around their reality; theoretical models that work well with computer science research were not sufficient. Applied anthropology's emphasis on user-centered development programs led the anthropologists to caution the computer scientists frequently not to go too far with their assumptions and objectives until the case managers were involved in both the definition of the problem and the process of imagining possible solutions. We also learned key issues like usability, cognitive overload, and information non-clutter that need to be considered while developing research software.

The anthropologists made the computer scientists aware that their relationship with case managers in the WtW project was very different from the relationship between a development team and clients in the usual software development setting. In most cases, clients come to softward developers with software in mind. In this case, the software developers had a computer science research agenda that was independent of the case managers' work, namely to develop planning algorithms that used both MPD solvers and constraint solvers. In order to build software solutions that would be useable by case managers, the team had to learn to respect the case managers' professional knowledge. The case managers also had to know that their participation in

the research experiments (for example, model elicitation) was critical to the technology's construction.

The team also learned to deal with challenges that arose due to interdisciplinary work. Some of the great challenges emerged not in understanding what esoteric terms like Bayesian network mean, but rather from seemingly simple terms such as "value", "variable", "state", and "utility" [22]. While each of these words are used commonly in the English language, the team found that they have dangerous differences in implication and connotation depending on the academic discipline of the team member. Even subtly different usages of these terms meant that few members of the team were clear about the software being designed and about the type of and format for the information required. Our collective deconstruction of the terms also forced members of the team, often from the same discipline, to rethink assumptions.

The highly contextualized, case-particular reasoning favored by the anthropologists often contrasted with the more positivist approach of the computer scientists. This led to challenges while building the HLE software. The anthropologists placed more emphasis on specific client profiles for eliciting information about different actions in the WtW, whereas the computer scientists wanted to build more abstract models of possible actions. The team learned to merge these contradictory ideas into a single coherent set of requirements for developing the elicitation software.

Reality is complicated. While great things can be done via abstraction, it takes time to figure out how to abstract data in ways that are both valid and reliable. Because building a correct and complete model is a slow process, we first developed a simplified model on which to test our solvers. We presented this model at the 2006 UAI Workshop on Bayesian Applications [10]. The initial, simplified model is sufficient to benchmark algorithms for factored MDPs with constraints, but is not intended to reflect the complex realities of Welfare-to-Work advising.

While working with social scientists and experts can complicate matters for computer scientists in the technology development stages, it ultimately will result in better software, more suited to the needs, worldview and desires of the end users. Isabelle Stengers [29] suggests that a willingness to step outside of our usual scientific norms can create opportunities to unbound disciplinary and even interdisciplinary work. Hunsinger [13] applies this critically to debates within the social construction of technology, suggesting that recognizing the sub-politics served in relation to our ideas should inform the work we do, especially as the social world becomes more complex and volatile. Technology is not neutral, nor merely the resolution of a research puzzle. Its applications have real consequences for real people who therefore are intent on its use, control and affects on their own lives. While computer scientists may be pro-

fessionally rewarded by successfully merging MDP and constraint solvers, case managers feared losing their jobs by being technologically replaced.

As the social construction of technology model predicted, the needs, mental models, and understandings of the project's goals differed among the groups on the team and strongly affected the development of our model-building technology. The academics learned that, when working in an interdisciplinary team, it helps us to be respectful, open and flexible to different ideas and paradigms. We benefited from the continuous scrutiny and reinterpretation from the multiple academic and expert perspectives. By listening to the experts and the anthropologists, the computer scientists were able to recognize and embrace the emergence of a new Bayesian model which more closely resembles the case managers' reality.

7 Acknowledgments

This work was partly supported by NSF grant ITR-0325063. The opinions expressed here are those of the authors and do not represent the Foundation, the University, or any social welfare offices. We thank Russell Almond for enabling communication between the computer scientists and social scientists about bowties, and we thank Joan Mazur for her work on the design of the HLE interface and her discussions of SCOT theory with several of the coauthors.

References

- [1] Russell G. Almond. Cognitive modeling to represent growth (learning) using markov decision processes. *Technology, Instruction, Cognition and Learning* (TICL), 5:313–324, 2007.
- [2] R.E. Bellman. Dynamic Programming. Princeton University Press, 1957.
- [3] C. Boutilier, T. Dean, and S. Hanks. Decision-theoretic planning: Structural assumptions and computationalleverage. *Journal of Artificial Intelligence Research (JAIR)*, 11:1–94, 1999.
- [4] Craig Boutilier, Thomas Dean, and Steve Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. *J. Artif. Intell. Res.* (*JAIR*), 11:1–94, 1999.
- [5] Rina Dechter and Robert Mateescu. Mixtures of deterministic-probabilistic networks and their AND/OR search space. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, 2004.

- [6] Jason DeParle. American Dream: Three Women, Ten Kids, and a Nation's Drive to End Welfare. Viking Adult, New York, 2004.
- [7] Dmitri A. Dolgov and Edmund H. Durfee. Symmetric primal-dual approximate linear programming for factored MDPs. In *Proceedings of the Ninth International Symposiums on Artificial Intelligence and Mathematics (AI&M 2006)*, Florida, January 2006.
- [8] M. Druzdzel and L. van der Gaag. Building probabilistic networks: Where do the numbers come from? *IEEE Transactions on Knowledge and Data Engineering*, 12(4):481–486, 2000.
- [9] Marek J. Druzdzel and Linda C. van der Gaag. Elicitation of probabilities for belief networks: Combining qualitative and quantitative information. In *Proceedings of the 11th Annual Conference on Uncertainty in Artificial Intelligence (UAI-95)*, pages 141–148, San Francisco, CA, 1999. Morgan Kaufmann Publishers.
- [10] Kendra Renee Gehlbach, Brandon Laracuente, Cynthia Isenhour, Judy Goldsmith, Beth Goldstein, and Miroslaw Truszczyński. A benchmark model for decision-theoretic planning with constraints, 2006.
- [11] Lynnell Hancock. Hands to Work: The Stories of Three Families Racing the Welfare Clock. HarperCollins Publishers, New York, 2002.
- [12] Jesse Hoey, Robert St-Aubin, Alan Hu, and Craig Boutilier. SPUDD: Stochastic planning using decision diagrams. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, pages 279–288, Stockholm, Sweden, 1999.
- [13] Jeremy Hunsinger. Broadening possibilities by expanding the theoretical richness of the social construction of technology. *Social Epistemology*, 19(2–3):255–259, 2005.
- [14] Cynthia Isenhour and Beth Goldstein. Welfare case management and competing moralities, 2006. Presented at the American Anthropology Association annual meeting November 2006, San Jose; submitted in 2007 for journal review as "The Welfare Contract and Critical Case Management".
- [15] Michael Katz. The Undeserving Poor: From the War on Poverty to the War on Welfare. Pantheon Books Inc, New York, 1990.
- [16] Catherine Kingfisher. Women in the American Welfare Trap. University of Pennsylvania Press, Philadelphia, 1996.
- [17] Hans K. Klein and Daniel Lee Kleinman. The social contruction of technology: Structural considerations. *Science*, *Technology*, & *Human Values*, 27(1):28–52, 2002.
- [18] Michael Lipsky. Street-level bureaucracy: Dilemmas of the individual in public services, 1980. New York: Russell Sage Foundation.

- [19] Suzanne M. Mahoney and Kathryn B. Laskey. Representing and combining partially specified cpts. In *Proc. UAI '99*, pages 391–400, 1999.
- [20] Krol K. Mathias. Knowledge Elicitation and Management: Backbone of a Decision Support System. PhD thesis, University of Kentucky, Lexington, Kentucky, 2008.
- [21] Krol Kevin Mathias, Cynthia Isenhour, Alex Dekhtyar, Judy Goldsmith, and Beth Goldstein. Planning for success: The social approach to building bayesian models, 2007. Submitted.
- [22] Joan Mazur, Beth Goldstein, and Judy Goldsmith. Social construction of technology in the welfare to work project. In *Proceedings of the Second Bayesian Modeling Applications Workshop (UAI 04)*, 2004.
- [23] Sandra Morgen. The agency of welfare workers: Negotiation devolution, privatization and the meaning of self-sufficiency. *American Anthropologist*, 103(3):747–761, 2001.
- [24] J. Pearl. Fusion, propagation and structuring in belief networks. Artificial Intelligence, 29:241–288, 1986.
- [25] Judea Pearl. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann, 1988.
- [26] Trevor J. Pinch and Wiebe E. Bijker. The social construction of facts and artefacts: Or how the sociology of science and the sociology of technology might benefit each other. *Social Studies of Science*, 14(3):399–441, 1984.
- [27] S. Renooij and Cilia Witteman. Talking probabilities: Communicating probabilistic information with words and number. *International Journal of Approximate Reasoning*, 22(3):169–194, 1999.
- [28] Silja Renooij and Linda C. van der Gaag. From qualitative to quantitative probabilistic networks. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, pages 422–429, 2002.
- [29] Isabelle Stengers. Power and Invention: Situating Science. University of Minneapolis Press, 1997.

A Pilot Elicitiations

Through a series of iterative interviews, our team of anthropologists, computer scientists and domain experts produced a list of the client characteristics which play into a case manager's decision regarding what activity to recommend to her client (including client interests, goals, aptitude, etc.). Before abstraction and consolidation, this list of client characteristics numbered more than 150 traits with multiple values. In order to build Bayes Nets, however, we had

to reduce the list. Therefore, we had to understand which of these variables were most determinative given a set of actions in which welfare clients can participate to fulfill their work requirement.

We began our elicitation process with a pilot elicitation. We utilized a participatory group interview focused on one action fragment: "get GED". By focusing on one action fragment the research team believed we could more accurately determine the effectiveness of the method and could explore the relationship between client variables and perceptions of the likelihood of success in this action. The method was developed to gather three pieces of information: (1) client attributes of state relevant to the case manager's decision to recommend the action "get GED" via free list, (2) The relative importance of each attribute of state listed in predicting success in a GED program via Likert scales, and (3) The optimal value of each attribute of state for predicting success in a GED program via small focus groups.

Twenty case managers from three different agencies were first asked to independently free list the "information you need to know about your client in order to assess whether or not getting a GED is a reasonable short term goal." We found that some case managers had significant problems with this wording, telling us that it is not part of their job to determine whether a goal is reasonable, but only to support clients in their goals. One of the most vocal case managers remarked,

"it's not my concern what the process is of them getting it [a GED]. That's what the adult education workers at the GED center deal with. That's their job. My goal is just to allow them that opportunity to get their GED done and then to tell them what else they have to do for case management to be in cooperation but I'm not going to stop them, I'm not going to determine if it is an appropriate short term goal. That's for the client to choose, that's for the adult education worker to determine if this is going to take three years or six months. That should not... that will not affect my case management at all." 6/15/2005

The case managers insisted on the import of client preferences and goals in the case management decision making process. These case managers privilege the client's preferences and make attempts to the client?s goals rather than passing judgment. While some authors have suggested that this emphasis on the client's directives is a coping mechanism used by case managers to remove themselves from ultimate responsibility for the client's outcomes [18,16], our research also indicates that case managers feel very strongly that the decisions most likely to lead to employment are those which have the full endorsement of the client.

In light of this understanding, we rephrased our question in terms more mean-

ingful for case managers. We now asked, "what information do you need to know about your client in order to determine her likelihood of success in a given action?" Case managers responded to this question more fully as the question's phrasing now paralleled the way they think about client potential success. They emphasize client determination as the client moves toward action plans designed to achieve success in an action, no matter how small the goal.

Next, case managers were asked to create a collective list of significant client attributes, each contributing items from their personal lists. In all, the case managers free listed 37 client variables that might affect a client's ability to succeed in a GED program. Then each case manager was asked to augment her personal list with any additional characteristics listed by the group before rating her own list on a Likert scale from extremely important to not very important. Finally, the group was broken into three smaller focus groups of 6–7 to talk about the five most important variables for determining likelihood of success in a GED program.

In order to compile the results from the pilot, the research team aggregated individual responses by number of mentions and relative importance. The five characteristics considered most important in determining outcomes in a GED program via this method of aggregation were: (1) learning disabilities, (2) last grade completed, (3) access to childcare, (4) age of the client, and (5) client's goals.

The focus groups helped to consolidate the client attributes into composite categories. They also reminded us that we could assume that barriers such as childcare and transportation would be immediately addressed by the agency. Finally, the focus group conversations largely validated the aggregated individual results. All three focus groups listed (1) learning disabilities; (2) educational history (which included highest grade completed, reading level, reason for dropping out of school); and (3) motivation (which included motivation, commitment, goals and resolve) in their five most important characteristics.