

AGENT-BASED SIMULATIONS OF SERVICE POLICY DECISIONS

Richard B. Lam

IBM T.J. Watson Research Center
1101 Kitchawan Road, Route 134
Yorktown Heights, N.Y. 10598, U.S.A.

ABSTRACT

During service engagements, project managers frequently encounter resource constraint issues. For each resource shortfall encountered, a project manager must decide among a narrow set of alternatives, weighing the resulting effects on project schedule, cost, and customer satisfaction. If a project is part of a larger collection of similar service engagements, it is less clear what the optimal strategy across all projects should be in deciding between alternatives. This paper describes an agent-based simulation environment to explore decision-making policies for hypothetical service business models using different agent-policy combinations. Results suggest advantages for maintaining flexibility in handling resource shortfall actions.

1 INTRODUCTION

In some sense, project management is essentially the handling of perturbations. A vast literature describes the methodologies associated with project management, including proper planning, budgeting, contracting, risk mitigation, and execution. However, most projects deviate in one or more respects from their detailed schedules despite the comprehensiveness or foresight reflected in the best-laid plans (see, for example, Cole, 1995). Communication breaks down, equipment is not available, software is not delivered on time, costs change, contractors cannot be hired, etc.

In this paper, we consider one particular type of perturbation to project plans – resource shortfalls. During a service engagement, there may be many reasons a project suddenly experiences resource shortfalls. Perhaps insufficient resources were identified at the beginning of the project. Resources could be unavailable due to sickness, accident, injury, vacation, training, emergency, geographical location, travel delay, etc. Activities may require different numbers of resources or skill sets due to unforeseen complexity, added scope, incomplete information, performance, or changing requirements. However, the point is not to model each potential reason for encountering a resource

issue – we are concerned only with handling a generic resource issue, which is a proxy for any of the resource availability causes listed above.

The project type modeled here is an outsourcing engagement that is part of a collection, or business tower, of similar outsourcing deals. In this model, an outsourcing vendor sets up a business to deliver a specific service, such as outsourcing of a human relations or accounting function, where the service is based on a common asset or set of assets. Thus, each client receives approximately the same solution, rather than one completely customized. The business tower also includes a pool of shared resources that are used in the transition phase of each service engagement. Thus, the projects are modeled as a collection of similar outsourcing deals, with similar resource requirements, schedules, budgets, profit, and overhead costs.

The problem in this space is to identify the appropriate response or responses to a resource shortfall during an engagement. These responses are modeled as to their effects on two engagement parameters – customer satisfaction and gross profit. The objective is to determine an optimal strategy or policy for deciding among alternatives on handling resource shortfalls. The policy should try to optimize both customer satisfaction and gross profit across all of the ongoing engagements of the business tower.

2 GAME THEORETIC ANALYSIS

As a project manager confronted with a missing or needed resource during an engagement, you have three typical responses:

1. Wait for a suitable resource to become available in the resource pool.
2. Borrow a resource from another ongoing project within the business tower.
3. Hire a resource to add to the resource pool for the business tower.

A project manager can choose to act on either option 1 or 3 alone, but option 2 requires the cooperation of another project manager. This provides a framework for a two-

player game. Figure 1 illustrates the choices of the players using an extensive view of the game (Rasmusen 2001).

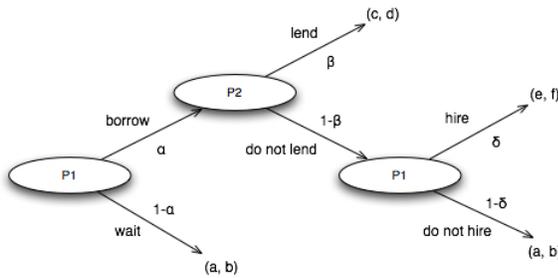


Figure 1: Extensive form of resource sharing decisions.

P1 represents the project manager confronted with the resource issue. The initial decision is basically whether to simply wait for a resource to become available (option 1 above) or to attempt to borrow a resource (option 2). P2 represents any other project manager within the business tower who may or may not choose to lend a resource to the requesting project manager. In the event a resource cannot be borrowed, P1 must decide if hiring (option 3) is appropriate.

Separating the decisions in this pair-wise manner where hiring is considered last introduces an assumption that hiring a resource is more of a last resort. This seems reasonable, as typically hiring is more expensive (in terms of initial and ongoing costs and hiring lead time) to the business as a whole rather than simply waiting or borrowing. However, if project managers are empowered to make their own localized hiring decisions, they may override the preference of the business as a whole.

The payoffs for each player in Figure 1 are given in (x, y) form, with the payoff to P1 denoted by x and the payoff to P2 by y . Note that the payoffs for both players are the same if P1 chooses to wait or chooses not to hire, as not borrowing and not hiring forces P1 to wait in this scenario.

Also note that Figure 1 does not specify the actual payoff amounts yet. Let's postpone that discussion briefly, and write out expressions for the expected payoffs for P1 and P2 (π_{P1} and π_{P2} , respectively) assigning probabilities to each decision as shown in Figure 1.

$$\pi_{P1} = (1-\alpha)a + \alpha\{\beta c + (1-\beta)[\delta e + (1-\delta)a]\}$$

$$\pi_{P2} = (1-\alpha)b + \alpha\{\beta d + (1-\beta)[\delta f + (1-\delta)b]\}$$

In these two payoff equations, the probability of P1 selecting to borrow is given by α and the probability that P1 will wait is thus $1-\alpha$. The probability of P2 deciding to lend is given by β (and thus not lending by $1-\beta$), and the probability of P1 hiring is denoted by δ (and not hiring by $1-\delta$).

2.1 Payoff metrics

Two important project measures we wish to model in this paper are customer satisfaction and gross profit margin. For implementing a gross profit measure, we specify the opportunity costs (in dollars) for hiring, borrowing, and lending resources. We implicitly assume there is no effect on gross profit for an engagement if an agent chooses to wait (or not hire). For the hypothetical business tower considered here, we will set the cost to a borrower for a resource at \$1K, and thus the advantage for a lender to -\$1K (i.e., borrowers incur the expense of borrowed resources against their own projects, while lenders' expenses on their projects are reduced). We also distribute the opportunity cost (\$100K) for hiring a new resource across the total number (N) of ongoing engagements, as the business tower is assumed to share a pool of resources and hiring increases the pool size and thus increases overall business expenses.

To implement a customer satisfaction measurement, all engagements are assumed to start with a nominal customer satisfaction score (although realistically, many projects may be set up to fail from the start in terms of customer satisfaction through poor expectation setting – we assume this is not the case here). Strategy choices of waiting or lending are assigned negative customer satisfaction effects, reflecting a decrease because either the waiting project or the lending project may now fall behind schedule. Strategy choices of borrowing or hiring in response to a resource shortfall are assigned positive values, as they represent proactive risk mitigation in response to resource issues.

Typical assigned values for customer satisfaction effects are: -2 if waiting for a resource due to a resource shortfall; -1 if lending (presuming lending a non-critical resource has less effect on a project schedule than waiting for a resource to become available when one is needed); and +2 if able to borrow or hire.

2.2 Strategy choices

Using values for the payoffs a through f based on gross profit only, we can analyze the payoff equations to determine if there are equilibrium strategies for P1 and P2. There is no effect on gross profit for waiting, so $a=b=0$. There is a negative effect on gross profit for borrowing and incurring the extra expense, so let $c=-1$ for the borrower and, correspondingly, $d=1$ for the lender. If P1 hires, all projects incur the overhead, so $e=f=-100/N$. Thus, the payoff equations (for $N=100$ projects) reduce to:

$$\pi_{P1} = \alpha[-\beta - (1-\beta)\delta]$$

and

$$\pi_{P2} = \alpha[\beta - (1-\beta)\delta].$$

Analyzing the game by looking backward from the last decision (i.e., P1's decision to hire or not), the payoff to P1 is either $a=0$ or $e=-100/N$, so P1 will prefer not to hire ($\delta=0$). Given this choice, P2 should rationally choose between the payoffs for lending ($d=1$) and for not lending ($b=0$, which will be P2's payoff based on P1's decision). Thus, P2 will prefer to lend ($\beta=1$). Finally, P1 must choose between waiting ($a=0$) or borrowing ($c=-1$), so the rational decision is to wait ($\alpha=0$). Given P1's decision to wait, there is no advantage to P2 for changing strategies, so this is an equilibrium – plugging the strategy values into the equations yields payoffs of zero for both players.

If we examine the game solely from a customer satisfaction standpoint, $a=-2$ if P1 waits, and $b=f=0$ as there is no difference in customer satisfaction for P2 if P1 waits or hires. And for P1, $c=e=2$ if a resource is borrowed or hired, while $d=-1$ penalizes P2 for lending. The payoff equations then reduce to:

$$\pi_{P1} = -2(1-\alpha) + \alpha\{2\beta + (1-\beta)[-2(1-\delta) + 2\delta]\}$$

and

$$\pi_{P2} = -\alpha\beta.$$

In analyzing the game tree for this case, P1 will compare $a=-2$ to $e=2$ and prefer to hire ($\delta=1$). Next, P2 will choose between the payoff for lending ($d=-1$) and for P1 hiring ($f=0$), so P2 will choose not to lend ($\beta=0$). Then P1 will choose between waiting ($a=-2$) and the expected payoff for trying to borrow ($e=2$) and will choose to try to borrow ($\alpha=1$). Again, there is no reason for P2 to change strategies, so this is a different equilibrium yielding payoffs of two for P1 and zero for P2.

To compare these contrasting strategy decisions, we need to look at how project managers are evaluated. If gross profit is important or if the business assigns strong incentives for supporting fellow project managers, then lending is a favorable alternative. However, if the business rewards a project manager solely on the basis of their achieved customer satisfaction, then holding onto resources may best achieve those aims. If we take the latter view, then P2 will choose never to lend.

Even in the absence of business-sponsored incentives though, altruistic motives may entice project managers to lend resources. They may then benefit from reciprocal behavior when, as is often the case, they find themselves in the position of needing additional resources. Implicit assumptions here are that project managers are empowered to make such resource-sharing decisions without interference from business policies, and that resources are interchangeable among projects (i.e., the skill sets and experience levels of shared resources are similar).

It seems clear that the strategies for P1 and P2 vary considerably, dependent on the metric selected to deter-

mine relative payoffs for the decisions. This situation becomes more complex for an N-player iterated game, where at each iteration of the game, all players run out of a resource and thus act as P1 once and act as P2 up to N-1 times during an iteration (i.e., the other N-1 project managers run out of a resource, so the decision for a given player to lend or not may be made as many as N-1 times for each iteration). To examine these cases, we need a simulation-based approach to examine project manager behaviors over a large number of decisions.

3 AGENT-BASED SIMULATION APPROACH

The preliminary results of the game theoretic analysis led to developing agent-based simulation models (Bonabeau, 2002) to explore the probabilities of selecting preferred resource sharing strategies depending on relative payoffs. NetLogo (NetLogo 2007) was selected as the agent-based simulation platform for this work. This section discusses the model parameters and metrics used to determine optimal strategies based on the overall distributions of gross profit and customer satisfaction for a hypothetical services business tower.

3.1 Business model definition

The NetLogo simulation defines two types of agents – expert project managers and novice or learning project managers. Expert agents in the simulation are programmed to follow a rational strategy with respect to the payoff metrics chosen, while learners follow probabilistic strategies set by the simulation user. Each project manager/agent belongs to a single services business tower specifically designed in this hypothetical model to represent a set of engagements designed for small and medium business outsourcing deals.

Typical parameters defining the business tower are read from a configuration file at runtime. These parameters define: the average revenue expected from each engagement along with a standard deviation to express the expected spread for the revenue per engagement; standard sales, general and administrative costs (SGA); transition costs for the outsourcing engagement; and an average resource cost and standard deviation to allow for variation in resource expenses per engagement. Note that this model follows the business model assumptions explained in the previous section, with relatively similar costs and resources associated with each engagement.

Typical values for the simulations are \$1M revenue per engagement, with a standard deviation of \$10K, SGA of \$200K, transition costs of \$150, and resource costs of \$500K with a standard deviation of \$10K. The opportunity costs and customer satisfaction scores discussed in Section 2.1 are also read from the simulation's configuration file.

3.2 Simulation setup

Simulations are initialized by reading in model parameters from the configuration file and assigning values to each agent based either directly on the parameters or on their average values and distribution (if specified). The numbers of agents of each type (expert or learner) are controlled through sliders provided on the NetLogo user interface. These numbers reflect the total number of simultaneous ongoing engagements within the business tower.

Figure 2 shows a typical NetLogo “world” after initialization, where each point represents one project manager/agent, plotted to reflect the starting customer satisfaction (y-axis) and gross profit (x-axis) for that agent. The agent types are distinguished by rendering the points in different colors (dark for expert, light for learner).

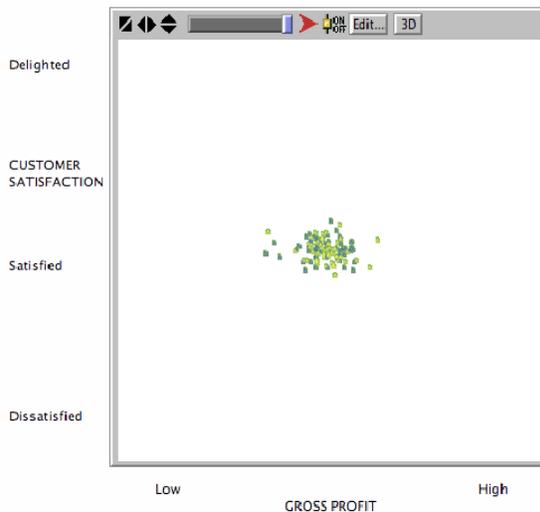


Figure 2: Example NetLogo simulation after initialization.

Customer satisfaction is set up on a scale from -10 (dissatisfied) to +10 (delighted). Starting values are assigned from a distribution with a mean of 0 (satisfied) and a standard deviation of 0.5.

3.3 Strategy decisions

Simulation runs begin by allowing the user to specify the number of expert and learner project managers. As a simulation proceeds, each project manager in turn faces a resource shortfall at each iteration. Based on the agent type, a strategy is followed to decide on a course of action. The action is taken and the consequences in terms of effect on customer satisfaction and new gross profit are calculated and used to update the agent’s position on the simulation world. The simulation terminates if an agent’s position falls outside the world boundaries (very low or high customer satisfaction, and no profit or profit in excess of 30%).

Referring back to the game analysis in Section 2.2, the rational strategy for maintaining or maximizing customer satisfaction is $\{\alpha=1, \beta=0, \delta=1\}$. And the rational strategy for maintaining or maximizing gross profit is $\{\alpha=0, \beta=1, \delta=0\}$.

Thus, the strategy chosen for a given situation such that in an iterated game, both overall gross profit and customer satisfaction are optimized, varies depending on the agent’s current position in the simulation world. If customer satisfaction is low and gross profit is high, an agent can choose to try to borrow or hire and refuse to lend, whereas if customer satisfaction is high and gross profit is low, an agent can choose to wait and offer to lend. This combination of strategies is automatically selected for the expert project manager/agents in the simulations.

For the learner agents, the user inputs the frequencies to use for borrowing, hiring, and lending. As a simulation proceeds, the interactions among the expert and learner agents results in a distribution of outcomes for customer satisfaction and gross profit across all of the engagements.

3.4 Simulation runs and results

Figure 3 shows sample results of a simulation run with the parameters discussed above, assuming the interaction of 100 expert project managers. Overall results of simulation runs using experts only are summarized in Table 1.

Allowing the expert agents to interact over a large number of iterations (i.e., decisions) results in both borrowing and lending at the same frequency, borrowing twice as often as waiting, and never hiring. Note that the strategy decisions for an expert trying to optimize customer satisfaction are set to prefer borrowing or hiring, with the result that borrowing is always tried first (following the game tree), and at least one project manager from a large group of projects can be found willing to lend – thus no hiring is done. However, for small numbers of projects, there will be iterations where no lending is done, and thus hiring is carried out, decreasing the overall gross profit.

Referring to Table 1, the mean customer satisfaction is reduced from a nominal value of 0 (satisfied) to -1.1 (very slightly dissatisfied) for large numbers of projects. Also significant is the increase in the standard deviation of customer satisfaction, which grows from an initial value of 0.5 to a range of 1.8-3.4 during the simulation, indicating that a policy of allowing transfer of resources among projects may introduce a wider distribution of customer satisfaction results across multiple projects. In situations with low numbers of projects, the chances of hiring are increased, but so are the chances of waiting, so the customer satisfaction for these projects is driven very low.



Figure 3. Interaction of 100 expert project manager agents (370 iterations).

Table 1: Summary of mean and standard deviation values for customer satisfaction and % gross profit using varying numbers of simultaneous projects. The results for each table entry are obtained from 100 repeated simulations.

Number of projects	Average customer satisfaction	$\sigma_{\text{cust sat}}$	Average % gross profit	σ_{GP}
10	-6.9	2.5	6.7	1.4
20	-1.9	3.4	14.6	1.3
30	-1.5	2.9	15.0	1.4
40	-1.4	2.7	15.0	1.4
50	-1.1	2.2	15.0	1.4
60	-1.1	2.1	15.0	1.4
70	-0.9	1.8	15.0	1.4
80	-1.0	1.9	15.0	1.4
90	-0.9	1.9	15.0	1.4
100	-1.1	2.0	15.0	1.4

For alternate strategies (followed by learner agents), the user can select the frequencies to use for waiting, hiring, borrowing, and lending resources. As an example, a set of simulations run solely with learners, using a strategy of always borrowing a resource (i.e., never waiting or hiring), was run while varying the number of simultaneous projects and the frequency chosen for lending resources. A 3-D plot of the results for mean customer satisfaction (over 100 runs of the simulation for each [number of projects, lending frequency] setting) is shown in Figure 4. The mean customer satisfaction is greater than the nominal value in all cases due to borrowing, and customer satisfaction for a borrower tends to increase as the lending frequency approaches 1, as there is a progressively greater chance of finding a lender. The downside, however, is that the distribution of customer satisfaction is much wider in all cases

($\sigma_{\text{cust sat}}$ goes as high as 3.6), as lending a resource negatively impacts the lender’s customer satisfaction. The slightly higher customer satisfaction values for low numbers of projects is an artifact of the stopping criterion used for the simulation, where any project hitting a world boundary causes the simulation to halt.

It is more interesting to explore situations where expert and learner agents interact, especially when there is one learner among a group of 100 projects (i.e., 99 experts). In this case, we expect a learner strategy of always hiring and never lending to drive the learner’s customer satisfaction to delighted at the expense of overall gross profit for the collection of projects.

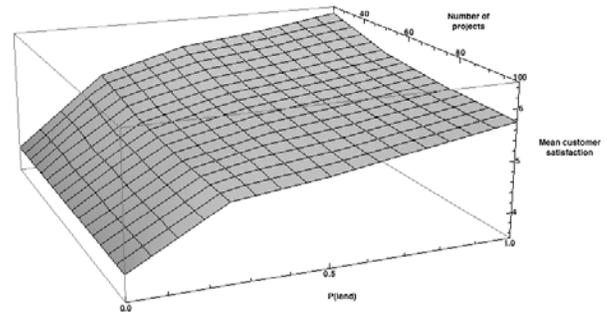


Figure 4. Mean customer satisfaction for learners who choose to always borrow a resource, as a function of the number of simultaneous projects and their chosen frequency of lending.

Now consider an altruistic learner, who chooses to always borrow but also to always lend a resource when required. The learner’s customer satisfaction spread is quite wide, and depending on when the simulation stops (i.e., how many decisions before one of the projects hits a world boundary), the learner can have either a higher or lower customer satisfaction than the other project managers. This is shown in Figure 5 for a sample run. Surprisingly, the learner can also end up at a higher gross profit than the other project managers as well. This is because there are many more opportunities to lend than to borrow, as at each iteration, the learner will be in the candidate list to borrow from for each of the other 99 projects and thus will be chosen more often. In a real project, of course, there is a limit to the number of resources that can be loaned, but this is reflected in the model by a large standard deviation in customer satisfaction for the learner.

Compare the altruistic learner to what we refer to as an opportunistic learner. The opportunist chooses never to lend but always to borrow, thereby taking advantage of fellow project managers. Because there is no lending of resources, the gross profit remains in the same range as for the other projects. However, always borrowing a resource results in driving to a customer satisfaction rating of delighted, while the other projects remain near the nominal customer satisfaction.



Figure 5. Sample run for an altruistic learner where both gross profit and customer satisfaction are driven to greater levels than the other projects. Note the learner (shown in the higher color) near the upper right corner of the world.



Figure 6. Sample run allowing interactions among 50 experts and 50 opportunistic learners who choose to always borrow and never lend. The opportunists end up with higher overall customer satisfaction at the expense of the other more cooperative project managers.

The interaction of multiple opportunists with an equal number of expert project managers is shown in Figure 6. The opportunists end up driving to a high percentage of delighted customers at the expense of the experts, who suffer by lending resources when required.

4 CONCLUSIONS

Agent-based simulations can divulge interesting behaviors among sets of entities following different strategies for de-

cision-making. It is difficult to obtain a quantitative understanding or basis for certain resource issues in outsourcing or other types of services engagements. This is particularly the case when soft factors such as customer satisfaction are considered. Qualitative judgments of project managers become the basis for gauging the effects of certain decisions on customer perceptions. Even survey instruments deliver a limited, semi-quantitative view on overall satisfaction, assuming project surveys are done at all and data across multiple services engagements is available.

Still, business executives must make resource decisions despite a dearth of data. In this work, we have attempted to put a framework around a simple set of decisions to discern some reasonable policies or heuristics to use in similar engagement scenarios. For instance, a loose management structure where project managers are free to make their own decisions for sharing resources may result in an unacceptable spread in customer satisfaction scores. Also, if hiring is left to individual project managers when there are low numbers of projects, the overall customer satisfaction of all projects may fall to unacceptable levels. In the case of the altruistic project manager, there may be cases where always lending resources is in the best interest of other projects, as when that customer has less priority than others and it is feasible to transfer multiple resources to other higher priority efforts. Finally, the opportunistic case is interesting as a way of addressing customer satisfaction for those customers who have higher expectations and are thus harder to please than others.

It will be interesting to test these policies against real business towers with a similar structure to the hypothetical model described here.

REFERENCES

- Bonabeau, E. 2002. Agent-based modeling: Methods and techniques for simulating human systems. *Proc. Nat'l Acad. Sciences* 99, suppl. 3.
- Cole, A., 1995. Runaway Projects – Cause and Effects. *Software World (UK)* 26:3-5.
- NetLogo, Available via <http://ccl.northwestern.edu/netlogo/> [accessed May 29, 2007].
- Rasmusen, E. 2001. *Games and Information*. 3rd ed. Malden, Massachusetts: Blackwell Publishing, Ltd.

AUTHOR BIOGRAPHY

RICHARD LAM is a Manager/Research Staff Member in the Mathematical Sciences department at the IBM T.J. Watson Research Center in Yorktown Heights, NY. He is investigating simulation models of service business processes via agent-based and system dynamics approaches.