

An agent-based model of innovation in organ transplant data

David S. Dixon and Sarah See Stith

*Department of Economics
University of New Mexico*

Abstract

In the face of government and professional regulation, insurance underwriters, public perception, and the Hippocratic Oath, the medical profession is naturally risk-averse. Simultaneously, the profession is expected to provide state-of-the-art care, especially in rapidly innovating fields. This expectation - often also coming from regulators, insurance providers, the public, and ethical considerations - amount to an incentive to be early adopters of innovative practices. This paper examines the diffusion of innovative medical practices with a network game theory model in which players face strategic complement payoffs. That is, players receive a payoff that depends on the number of their peers playing the same strategy. Innovations diffuse through the network as the marginal benefit exceeds the marginal cost given each player's risk aversion. A specific application to kidney transplants in the U.S. is presented. The network topology inferred from historical trends in the diffusion of immunosuppression innovations is examined for underlying links such as common ownership, common regulatory authorities, movement of staff, and medical pedigree of staff.

Background

Helderman et al. (2003) present a history of organ transplantation in terms of periods of innovation in immunosuppression. They refer to the period from 1954 to 1962 as the *experimental era*, when the only routine procedure was kidney transplants between identical twins. Azathioprine was the first major drug innovation in immunosuppression, cited by the Nobel committee when awarding the 1988 prize in physiology and medicine to Gertrude

Elion and George Hitchings.¹ Helderman et al. report that the *azathioprine era*, beginning 1962, saw fifty percent survival rates in kidney grafts and experimentation in liver and heart transplantation. They note that in the third era, kidney transplantation became routine with high survival rates, and success was common in liver, pancreas, heart, and lung transplants. This era is named for the immunosuppressant cyclosporin A.

Beginning in 1986, researchers were reporting success with the new immunosuppressant tacrolimus.(Fung, 2004) This was not the revolution in immunosuppression that was ushered in by cyclosporin, but a more subtle innovation. Clinical trials(Wallemacq and Reding, 1993) showed that tacrolimus was successful with higher-risk populations. The new immunosuppressant was made available routinely in Japan beginning 1993 and in the United States in 1994. Subsequent clinical studies (e.g. Knoll and Bell, 1999;Margreiter, vs Ciclosporin Microemulsion Renal Transplantation Study Group et al., 2002;Vincenti et al., 2002;Webster et al., 2005) showed both improved efficacy and reduced side effects of tacrolimus over cyclosporin.

The innovation of tacrolimus was unlike those of azathioprine and cyclosporin in that there was no obvious increase in marginal benefit, at least initially. This makes the early adoption of tacrolimus an indicator of the relative marginal costs of adoption among the transplant centers. The market for kidney transplantation is not a market in any of the conventional senses, aside from the reality that a) demand far exceeds supply at any regulated price, and b) market equilibrium is achieved on an organ-by-organ basis.

Given that none of the conventional measure of marginal cost and marginal benefit are appropriate, in this paper we use the term *risk aversion* to refer to the marginal cost to a transplant center for adopting an innovation. Marginal benefit of adopting an innovation for a transplant center is inferred from the number of other centers already adopting the innovation. The clearest analog in reality is network externality, but it may also encompass learning by doing, economies of scale, and increased human capital.

Data

Center-identified data on transplanted patients were obtained under a confidentiality agreement from the United Network for Organ Sharing. These data cover the period from 1 October 1987 through 30 September 2013.

The data provide partial records for 387,021 transplants and 2,636,830 follow-up examinations. There were many immunosuppressant regimes administered over the data period. The two of interest to this paper are cyclosporin and tacrolimus, the replacement of the former with the latter being the innovation under examination. The annual frequency for each is shown in Figure 1.

¹http://www.nobelprize.org/nobel_prizes/medicine/laureates/1988/press.html accessed 26 June 2015.

Network game theory model

When an experimental practice has deleterious effect, it is abandoned. Otherwise, it is an innovation. If the innovation is an unqualified improvement to the state of the art, like the immunosuppressants azathioprine and cyclosporin, it is adopted as quickly as possible. Not to do so would constitute failure to provide due care.

If an innovation is an incremental improvement, however, its adoption rests on the marginal costs and marginal benefits to the transplant center. Immunosuppressant therapy has a complex fitness space, and an innovation may be an improvement in some fitness aspects while imposing increased risks in others. A therapy with increased efficacy for higher-risk transplants which also exhibits higher incidence of negative side effects is such an example. Importantly, although initial trials may have demonstrated a necessary level of safety, there is high uncertainty in many fitness trade-offs early in the adoption of an innovation. Marginal cost and marginal benefit may be difficult to ascertain, may change with patient profiles, and may vary with changes in practices, personnel, or management of the center.

In the case of immunosuppression with tacrolimus, its dual benefit of increased efficacy and reduced side effects were not well established for several years after its introduction. Early adopters of tacrolimus were less sensitive to uncertainty in its fitness. This may be because they were research hospitals, served a higher-risk patient community, or experienced a wide risk envelope simply because they performed large numbers of transplants. For simplicity, we consider these early adopters as less risk averse, and that the marginal cost of adoption is proportional to risk aversion.

When the fitness of an innovation is mixed or uncertain, the benefits for a transplant center are difficult to assess. Each center experiences positive externalities from adoption by others, however. There are network externalities in that increased use of an innovation reduces uncertainties and promotes the perception of it as state-of-the-art. There is also the possibility for learning by doing², and the human capital benefit from the availability of a larger pool of medical personnel experienced with the innovation. There may also be economies of scale in drug or apparatus production.

The adoption of the tacrolimus innovation is a study at the margin, where early adoption may indicate low marginal cost, since the marginal benefit (adoption by peers) is low. Our basic model comes from network game theory.(Galeotti et al., 2010) The payoff in this game is a strategic complement: the payoff for playing a strategy increases with the number of network neighbors also playing that strategy. The cost of playing the adoption strategy is inferred as risk aversion.

Lamberson (2011) shows that a network game with simple two-strategy strategic complement payoffs reaches an equilibrium outcome that all players play the same strategy. Which is the final equilibrium strategy depends on the initial conditions and the game payoff, as seen in Figure 2. Dixon (2015) shows that this is a static equilibrium: the network reaches a state wherein all players are at Nash equilibrium. Which is the final equilibrium

²Stith, however, finds no evidence of learning by doing where patient outcomes are concerned.

strategy depends on the *critical fraction*, the fraction of other players initially playing a specific strategy. All states beginning above the critical fraction end at all players playing that strategy, and all states beginning below the critical fraction end at all players playing the other strategy. The critical fraction depends on the relative payoffs.

The outcome equilibrium and the critical initial strategy are sensitive to network topology.(Dixon, 2011) That is, outcomes and critical fractions for a network in which all the players have the same number of connections to all the other players - a *regular* network - are different than for a network with a random distribution of connections between players (a *Bernoulli random* network).

Agent-based model

The agent-based model is implemented using the network extensions in NetLogo.(Wilensky, 1999) The initial network is an undirected $(N-1)$ -*Regular network*. That is, it has N nodes each with the same degree (the number of connections to other nodes) equal to $N - 1$. This is also termed a *unit density graph*, meaning that all nodes are connected to all other nodes.

In game-theoretical terms, the nodes are players, which are transplant centers in this model. The unit density graph is the zeroth-order assumptions that transplant centers are a tight-knit community in which the practices of all other centers are important information to each center. Subsequent studies will examine other network topologies in which clustering occurs due to communality among subsets of centers.

The centers are the only agent-type in the model. Agents are initialized from a database with location, number of total transplants, number of tacrolimus transplants, time step of first and last transplants, and mean and standard deviation for risk aversions. The term *risk aversion* is chosen to reflect an abstract overall cost of adopting an innovative practice. It is an empirical value determined through simulated annealing of simulation outcomes as described below. How this value relates to known risk factors for transplant centers is the topic of ongoing research. For the purpose of discussing this model, the term risk aversion refers always to this parameter and not to any direct measurement of an individual center's risk aversion. In simulation, risk aversion is stochastic, with a small random increase or decrease each time it's computed. This is intended to capture variability in information about patients, donors, and the state of the art, as well as stochasticity in center practices.

The benefit to a center for adopting an innovation is assumed linear in the number of other centers having already adopted the innovation. This is considered an abstract measure of the externalities discussed in the previous section. The payoff is a Heaviside function of the difference between benefit and cost. A center utilizes the innovation when the benefit exceeds the cost. Because of stochasticity, a center may utilize the innovation one time and not the next. For outcome assessment purposes, however, the center is considered as having adopted the innovation at the time of its first use.




Figure 1: Comparison of cyclosporin and tacrolimus immunosuppression regimes.




Figure 2: Outcomes with a strategic complement payoff given different initial conditions. The payoff is a Heaviside function at $k = 4$, where k is the number of other players playing strategy x . Shown are the equilibrium outcomes resulting from varied initial fractions of players playing x . (From Lamberson, 2011)

The first goal of simulation is to estimate a relative risk aversion for all centers. This is done through a process of simulated annealing (SA). In SA, the term *temperature* refers to the stochasticity of the value to be estimated, in this case, risk aversion. In the first round of Monte Carlo simulations, all agents are initialized with the population mean and standard deviation for their overall ratio of tacrolimus transplants to all transplants.

The time span of each simulation is 9550 days, the time span of the transplant data. At each time step, all 304 centers draw from a probability distribution determined by the date first and last days of transplants, and the total number of transplants they performed. Centers appearing in the the first day of data has a first day of zero, and centers appearing the last day of data have last days of 9550. Between the first and last days each center has a probability $1/N$ of performing a transplant, where N is the total number of transplants performed by that center.

If a random draw determines that a center performs a transplant on that day, the benefit and cost of employing tacrolimus immunosuppression are computed as described above. If the benefit exceeds the cost, tacrolimus is employed for this transplant. The result of each transplant decision is written to file for post-processing. The risk aversion estimates are refined in post-processing and used as input data for a subsequent set of simulations, where risk aversion is varied randomly with one-half the standard deviation used in the previous iteration.

In post-processing, R^2 is computed for each simulation run, where $R^2 = ESS/TSS$. TSS is the sum of of the squares of the number of tacrolimus transplants each day (typically but not necessarily one or zero) over the actual transplant data. ESS is the same calculation for each simulation. The results from all simulations are regressed on

$$R^2 = \beta_0 + \sum_{i=1}^{304} \beta_i \mu_i$$

where μ_{1-304} are the risk aversion estimates going into each set of simulations.

Monte Carlo results

The regression results are consistently significant for 68 of the 304 centers in each iteration of simulated annealing. The risk aversion estimates were refined going into the first three phases as shown in Figure 3. In the course of the first three phases, mean R^2 over all simulations ($N \cong 6,000$) improves from 0.49 to 0.74.

Downward revisions of risk aversion are generally smaller than upward revisions, as seen in Figure 3a. The effect of these revisions, however, is not large given that the greatest revisions occurred for some of the smallest centers, as shown in Figure 3b. These results imply that the majority of transplant centers subscribe to the same standard of acceptance for this innovation. While a few of the iterated revisions to the risk aversion parameter are sizable, they appear to only affect the smallest centers.


The next step in the process will be to examine more complex network topologies. It is likely that there are some centers that are more influential than others, and subgroups of centers that are tightly coupled.

The following step will be to introduce autonomous behaviors. Dixon (2015) shows that if even a few players employ a non-self-optimizing strategy, the static equilibrium outcomes seen in Figure 2 are attenuated or even collapse into a single outcome. We propose to examine the effects of certification standards, government regulations, and market pressures in terms of their distorting effects.


References

- Dixon, David S.** 2011. "Preliminary results from an agent-based adaptation of friendship games." 86th Annual Conference of the Western Economics Association International.
- Dixon, David S.** 2015. "Sheep, contrarians and saboteurs: disruptors in friendship games." Eastern Economics Association. 41st Annual Conference of the Eastern Economics Association.
- Fung, John J.** 2004. "Tacrolimus and transplantation: a decade in review." *Transplantation*, 77(9): S41–S43.
- Galeotti, A., S. Goyal, M.O. Jackson, F. Vega-Redondo, and L. Yariv.** 2010. "Network games." *Review of Economic Studies*, 77(1): 218–244.
- Helderman, J Harold, William M Bennett, Diane M Cibrik, Dixon B Kaufman, Andrew Klein, and Steven K Takemoto.** 2003. "Immunosuppression: practice and trends." *American Journal of Transplantation*, 3(s4): 41–52.
- Knoll, Greg A, and Robert C Bell.** 1999. "Tacrolimus versus cyclosporin for immunosuppression in renal transplantation: meta-analysis of randomised trials." *Bmj*, 318(7191): 1104–1107.
- Lamberson, P.J.** 2011. "Friendship-based Games." 37th Annual Conference of the Eastern Economics Association.
- Margreiter, Raimund, European Tacrolimus vs Ciclosporin Microemulsion Renal Transplantation Study Group, et al.** 2002. "Efficacy and safety of tacrolimus compared with ciclosporin microemulsion in renal transplantation: a randomised multi-centre study." *The lancet*, 359(9308): 741–746.
- Stith, Sarah S.** "Organizational Learning-by-Doing: Evidence from Liver Transplantation." Working paper.

- Vincenti, Flavio, Stephen C Jensik, Ronald S Filo, Joshua Miller, and John Pirsch.** 2002. "A long-term comparison of tacrolimus (FK506) and cyclosporine in kidney transplantation: evidence for improved allograft survival at five years1." *Transplantation*, 73(5): 775–782.
- Wallemacq, Pierre E, and Raymond Reding.** 1993. "FK506 (tacrolimus), a novel immunosuppressant in organ transplantation: clinical, biomedical, and analytical aspects." *Clinical chemistry*, 39(11): 2219–2228.
- Webster, Angela C, Rebecca C Woodroffe, Rod S Taylor, Jeremy R Chapman, and Jonathan C Craig.** 2005. "Tacrolimus versus ciclosporin as primary immunosuppression for kidney transplant recipients: meta-analysis and meta-regression of randomised trial data." *Bmj*.
- Wilensky, U.** 1999. "NetLogo." Accessed 31 August 2011.



(a) Sorted by size of adjustment.



(b) Sorted by size of the transplant center.

Figure 3: Refinements in risk aversion over the first two phases of simulated annealing.

ODD

For **An agent-based model of innovation in organ transplant data** by David S. Dixon and Sarah See Stith.

Purpose

The purpose of this model is to estimate the relative risk-aversion for U.S. kidney-transplant clinics based on data from 1987 to 2013.

Entities, state variables, and scales

There are two entities in this model: an agent-entity representing a transplant clinic and a link-entity representing the link between two clinics. Each agent-entity has the state variables in Table 1. The descriptive statistics for the external data are shown in Table 3. The state variable *risk-aversion* is assigned stochastically in simulation, drawn from a normal distribution defined by *ra-mean* and *ra-stdev*, variables assigned in meta-simulation (see Process Overview below). Environmental state variables are shown in Table 4. This table does not include variables used to display entities and the underlying GIS shapes

Table 1: Agent-entity variable descriptions.

Variable	Source	Description
Center-ID	D	A unique clinic identifier from the transplant data
latitude	D	The latitude of the transplant clinic (used for display but not in simulation in this model)
longitude	D	The longitude of the transplant clinic (used for display but not in simulation in this model)
total-count	D	The total number of transplants performed by the clinic
tacro-count	D	The number of transplants after which tacrolimus was prescribed (not used in this model)
start-t	D	The time step on which the clinic began transplanting kidneys
end-t	D	The time step on which the clinic stopped transplanting kidneys
ra-mean	M	Mean of the distribution from which risk aversion is Monte Carlo sampled (initially set to the mean for all agents)
ra-stdev	M	Standard deviation of the distribution from which risk aversion is Monte Carlo sampled (initially set to the mean for all agents)
adopted?	S	Boolean flag set to true when the clinic begins prescribing tacrolimus
apply-count	S	The number of times the innovation was applied
p-patient	S	The probability per time step of performing a transplant (computed from total-count, start-t, and end-t)
risk-aversion	S	The current Monte Carlo sample of risk aversion

Table 2: Link-entity variable descriptions.

Variable	Source	Description
adoption-strength	S	The strength of a link inferred from the relative rate of adoption of the agent-entities sharing the link
ownership?	D	True if the linked agent-entities share a common owner (not used in this model)
staff?	D	True if the linked agent-entities share any professional staff (not used in this model)
medschool?	D	True if medical staff have a medical school in common (not used in this model)

Source

D from transplant data

S set by simulation events

Table 3: Descriptive statistics of state variables.

Process overview and scheduling

There is a setup phase, a simulation phase, and a meta-simulation phase. The setup phase is executed once per simulation. The simulation phase is executed for a set number of time steps, representing days from October 1987 through 30 September 2013. The meta-simulation phase is run between very large ensembles of simulations.

Setup phase

In the setup phase, the the agents are instantiated by reading external data from a comma-delimited file, one agent per transplant clinic in the data file. The input data includes the total number of transplants performed by each clinic and the time steps over which that clinic performed transplants. Each agent computes a probability of performing a transplant, *p-patient*, by dividing the number of transpanlants by the number of time steps it operated. Input parameters for each agent also include *ra-mean* and *ra-stdev*, from which each agent sets state variable *risk-aversion* to a Monte Carlo sample from normal distribution $N(ra - mean, ra - stdev^2)$. The *risk-aversion* draw for each

Table 4: Environmental variable descriptions.

Variable	Description
fraction-adopting	The number of clinics adopting the tacrolimus innovation. This, as a time-series, is used to calibrate the model (see Process Overview)

agent is output for use in the meta-simulation phase. A fully-connected network is constructed by randomly linking pairs of agents.

Simulation phase

At each time-step, the start and end time steps determine whether a clinic-agent is active. If an agent is active, a random value is sampled from $U(0,1)$. If the random sample is less than the agent's *p-patient*, the clinic-agent performs a kidney transplant. Then that clinic-agent computes the network-game strategic complement payoff by polling all neighbors on the network (meaning all other clinics). The payoff is the fraction of neighbors that have adopted the innovation. The clinic-agent's risk aversion is the cost of adopting. If the payoff exceeds the cost, the clinic-agent adopts the innovation starting with this transplant. The output at each time step is the time step and the total fraction of clinics that have adopted the innovation at that time step.

Meta-simulation phase

After a large number of simulations (between 5,000 and 10,000), the R^2 for each simulation is calculated by comparing the simulation time series to the actual time series data. The R^2 for each simulation is regressed against the *risk-aversion* for each agent in each simulation. The marginal R^2 for each agent is used to update *ra_mean* for each agent, while *ra_stdev* is reduced for the next simulation ensemble. This is a form of simulated annealing to maximize R^2 by fitting each agent's *ra_mean*.

Design concepts

The overall design concept is that, in a tight, largely self-regulated community like kidney transplant clinics, each clinic is influenced by the practices of all the others. Some clinics may be more risk-tolerant where innovation is concerned: large research hospitals at universities, for example. While others, such as small private clinics, will wait until an innovation is established and well-documented before adopting it. A all-to-all network game with strategic complement payoffs is a flat-prior initial estimate of the interactions of kidney transplant clinics in the absence of data about specific relationships. There are no emergent properties in this model. The agents in this model have no adaptive or learning capabilities. These agents do not optimize and therefore have no objectives. The predictive goal of the model is to identify the relative thresholds at which clinics adopt an innovation. The networks on which these agents live are entirely ad hoc. The only interaction between these agents is the influence that neighbors have on the decision to adopt the innovation. The state variables are initialized from external data. There is no specific aggregation or collective behavior of these agents. The principle data collected are the time-series of innovation as a function of relative risk aversion of the agents, and the R-squared fit of the actual time-series of adoption. These results are presently being used to identify correlations with levels of risk aversion, such as clinic size, public or private funding, regulatory status with regard to patient outcome standards (in compliance, out of compliance, and on probation), location, local demographics, etc. A subsequent model may incorporate behaviors associated with these attributes to better model specific clinic outcomes.

Initialization

Agents are initialized from external data in a comma-delimited file.

Input Data

Center-identified data on transplanted patients were obtained under a confidentiality agreement from the United Network for Organ Sharing. These data cover the period from 1 October 1987 through 30 September 2013. The data provide partial records for 387,021 transplants and 2,636,830 follow-up examinations. The innovation of interest is the introduction of tacrolimus as an immunosuppressant therapy.

Submodels

There are no submodels.