Interaction and modelling of entries to drug use

Trying to model two mechanisms: fear and social stigma

Comments welcome!

Hans O. Melberg
hom@sirus.no

SIRUS
Oslo, November, 2001

Social interaction, treatment and modelling changes in the number of drug users

General introduction

In this paper I want to focus on two mechanisms that may help to explain levels of drug use in a society: observational learning and self-reinforcing social stigma effects. In short, before deciding whether to experiment with drugs or not potential users learn about the dangers (and attractions) of using drugs by observing what happened to previous users and they consider the amount of stigma associated with taking drugs.

The general approach is very much inspired by Schelling's (1978) book on "Micromotives and Macrobehavior." He shows how social science may produce counterintuitive (but true) results by focusing on social interaction. One well known example is how we may end up with completely ethnically segregated neighbourhoods even if everyone wants to live in a mixed neighbourhood with, say, 60% of their own ethnic group and 40% of the minority. To me this showed the importance of focusing on social interaction between heterogeneous agents instead of assuming a "representative agent" as is often done in the economics literature (see Kirman).

A second inspiration for the current paper, is theories of bounded rationality and learning in evolutionary game theory (see Conlisk (1996) and Michihiro(1996)). This is partly because I encountered some formal problems with my original model (equilibrium selection in a situation with multiple equilibria), but also because they seemed to offer a more plausible interpretation of what was going on. In the original framework I had to make quite strong assumptions about rationality and learning. Hence, evolutionary game theory offered at least a potential way of both becoming more precise (reduce the number of plausible equilibria) without adding more implausible assumptions. In fact, weaker assumptions could produce better predictions. That at least, was the idea at the outset of my research.

The structure of this paper follows my line of inspiration chronologically. First of all I will formalize the micromotives I want to explore. From the microfoundations I then create a macromodel of the equilibrium level of drug use in a society. I then explore the implications of the macromodel. One of these is the potential existence of multiple equilibria and I this raises the issue of using evolutionary game theory. But, as I then discovered, evolutionary game theory cannot be used just to eliminate some of my equilibria, I had to reinterpret the whole model (including the microfoundations) in an evolutionary game theory frame.

The microfoundation

Assume that each individual (i) who turns 15 in time period (t) decides whether to use drugs based on a comparison of the expected utility of using drugs [$EU_{it}(D)$] and the expected utility if you chose to abstain from drugs [$EU_{it}(A)$]. The underlying assumptions is that the choice of whether to use drugs is at least influenced by some kind of calculation, as opposed to purely emotional or norm based behaviour. Moreover, for the sake of simplicity the choice is assumed to be a one upon a lifetime choice i.e. only people who decided not to experiment with drugs cannot later choose to do so.

(1) Use drugs iff: $EU_{it}(D) > EU_{it}(A)$

What is $EU_{it}(D)$ and $EU_{it}(A)$? The last - $EU_{it}(A)$ - is assumed to be a constant. To work out $EU_{it}(D)$ we need to consider the possible consequences of experimenting with drugs. For the sake of simplicity I shall assume that there are only two possible outcomes for individuals who experiment with drugs. Either you have a “junkie career” (unhappy) or you have a “yuppie career” (less unhappy). This assumption is meant to capture the fact that not all individuals who experiment with drugs end up as stereotypical “junkies.” In fact, only a very small minority of drug experimenters end up as junkies.

The two possible outcomes of experimenting with drugs - becoming a yuppie or a junkie – result in certain payoffs. A very simple way of formalizing this would be to say that $U_{it}(J)$ is the total (discounted) sum of utility you receive if you end up as a “junkie” (for individual I at time t) while $U_{it}(Y)$ is the total sum of utility if you en up as a yuppie. Note that $U_{it}(J)$ does *not* represent annual utility as a junkie. It represents the *total* (discounted) sum of utilities from the rest of your life if it turns out that experimenting with drugs results in addiction. Hence, it may include some years as a happy users (a yuppie), then some years as a junkie and then, finally, some years as a non-users (treated or “matured out”). The same goes for *U(Y)*. It does not only include years as a happy drug users, but also years as a non-users after being a “happy user.”

Some individuals end up as junkies and some end up as yuppies, but nobody knows in advance what he or she will become. Hence, in order to work out the expected utility of taking drugs, we the individual must estimate the probability of becoming a junkie. One way of doing so, would be to observe the outcome of previous generation’s experimentation with drugs. The probability of becoming a junkie could then be estimated by the share of junkies (*j*) of all drug users ($d = j + y$) in the last time period for which information is available.

$$(2)\quad E(p_{it}) = \frac{j_{t-1}}{j_{t-1} + y_{t-1}}$$

I do not claim that this is the only, best or rational way to estimate your probability of becoming junkie. It does, however, seem like one plausible factor that could affect your estimate. If many of the current users are junkies then one would expect people to believe that the probability of becoming a junkie (if you become a user) is high.

Finally, in order to capture the effects of social stigma I introduce a “moral cost” of experimenting with drugs (m_{it}). One might think about this as the cost of doing something that many people dislike – i.e. the cost of being an “outcast.” I shall also assume that individuals have different moral costs i.e. they differ in the extent to which they are influenced by social disapproval. The exact nature of the distribution of moral cost is important and I will experiment with different types; uniform distributions and normal distributions.

Altogether then, the expected utility of experimenting with drugs for an individual at a point in time is the utility he will receive in the two possible outcomes (junkie career or yuppie career) multiplied by their respective probabilities and adjusted for social stigma:

$$(3)\quad EU_{it}(D) = p_{it}U_{it}(J) + (1 - p_{it})U_{it}(Y) - m_{it}$$

One might argue that the formulation so far ignores many issues that are central to addiction. For instance, I do not explicitly model discounting which many people argue is an important phenomena when trying to explain addiction. I do not deny the importance of discounting, but the focus in this paper is something else, namely the effects of interaction through observational learning and social stigma. I want to isolate this and to do so I do not want to bring in more complications than necessary. Hence, more explicit modelling of discounting is left to be explored at a later stage.

Aggregation

So far all I have is a very general formulation of the decision problem. What I want, however, is an expression of the aggregate result if people make their decisions based on the microfoundation just described. This requires several assumptions, both in terms of simplifying assumptions and in terms of more substantial assumptions about the mechanisms of aggregation.

In order to make it easier to get analytic results, I now make the following simplifying assumptions:

(4) $p_{it} = p_t \quad \forall i$ (every individual uses the same probability of becoming a junkie)

(5) $U_{it}(.) = U(.) \quad \forall it$ (the utility of ending up as a junkie, a yuppie or an abstainer is the same for every individual at all times)

This means that every individual uses the same probability of becoming a junkie and that the utility of ending up as a junkie or as a yuppie is the same for every individual at all times. The decision problem for the individual is then reduced to comparing $EU(A)$ to the following expression:

$$(6)\quad EU_{it}(D) = p_t U(J) + (1 - p_t)U(Y) - m_i$$

In this formulation there is only one variable that change over time: the probability of becoming a junkie. From a given starting point with an arbitrary share of drug users in the population, the process will then be as follows: Based on the current share of junkies the individuals in the new generation will calculate the expected utility of experimenting with drugs. They will then compare this to the expected utility of the alternative. What we need to find out more about the aggregate result is to work out how changes in *p* affect expected utility and how this, in turn, affect the size of the group than begins to experiment with drugs.

A rise in the probability of becoming a junkie leads to a reduction in the expected utility of experimenting with drugs; the size of the reduction is determined by:

$$(7)\quad \frac{\partial EU(D)}{\partial p} = U(J) - U(Y)$$

I assume that this is negative i.e. that the utility of a career as a junkie is lower than the utility of a career as a yuppie [$U(J) < U(Y)$]. Note that the difference is a constant i.e. the relationship between changes in p and the expected utility of experimenting with drugs is linear. Note also that the degree to which expected utility change depends on the difference between utility of junkie career and utility of yuppie career; the more “unhappy” the junkies are relative to the “yuppies” the larger the reduction in expected utility when there is a rise in p. One would expect that this translate into a “large” reduction in the number of individuals who wants to use drugs. The question is then what kind of relationship we have between changes in expected utility and people starting to use drugs.

How does changes in expected utility relate to changes in the number (or share) of people starting to use drugs. Recall that the decision-rule used by individuals is.

Use drugs iff $EU_{it}(D) > EU(A)$

Since the partial derivative of EU(D) with respect to *p* is a constant, we have the following relationship:

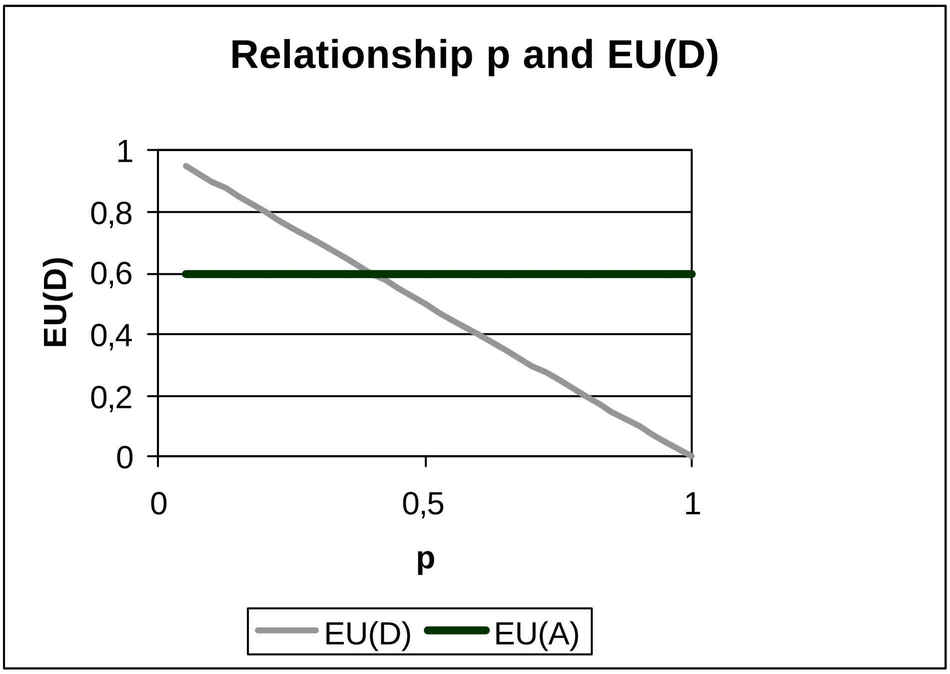

In words: An individual (in the incoming generation) decides to use drugs is his expected utility is above some constant (here EU(A)). The expected utility depends linearly on *p* – the probability of becoming a junkie.

How do we translate this micro-mechanism into an aggregate equation? The easiest way to do this is that say that moral costs are uniformly distributed. We would then have that the share of the new generation that begins to experiment with drugs is a linear function of *p*. If moral costs are not uniform, then there is not a one-to-one relationship between changes in expected utility and changes in the number of people entering drug use. This is an extension I will study later.

Assuming a uniform distribution is necessary, but not sufficient. I also have to make sure that the cut-off points "appropriate." For instance, imagine that the "limit" for using drugs is very high, then it might be the case that a small reduction in fear is not enough to make anyone begin to experiment with drugs. Hence, when the "cut-off points" are "outside" the relevant interval the relationship between changes in *p* and changes in the number of drug users is not a continuous linear function. In order to make thinks work I then either have to make the model more complicated to account for the mentioned possibility (that changes in *p* do not affect the number – either because everybody wants to use it even if *p* changes slightly or because nobody wants to use it despite the change in *p*). The easiest way out is simply to assume that the "limit" is at *p*=0. That is, I assume values [for EU(A), U(J), U(Y)] which implies that at *p*=0 everybody use drugs and at *p*=1 nobody use drugs. By doing so I assume away the possibility that changes in *p* do not affect the number of users.

The aggregate results

The share of yuppies in the whole population at any point in time (y_t) can be found by taking last years share, subtracting the share that leaves the group every year and adding the share of the new generation that starts to use drugs and. I shall assume that the exit process is exogenous i.e. there is a certain percentage of the yuppies (b_1) that leave the group every year (they might die, quit, or become junkies). The “input” process, however, is endogenous. Based on the microfoundation described above the size of the group of new users depends linearily and negatively on *p*. It also depends on the size of the new generation. We may use g_t to symbolize the size of the new generation as a share of the total population (assume this is a constant for now)[1]. Finally, we know that the degree to which a change in *p* affects expected utility (and hence also the size of the new group of users) depends on U(J)-U(Y).

$$(8)\quad y_t = y_{t-1} - b_1 y_{t-1} + b_2 (1 - E(p_t))$$

where

$$(9)\quad b_2 = f\left(U(J) - U(Y), g\right)$$

and

$$(10)\quad E(p_t) = \frac{j_{t-1}}{j_{t-1} + y_{t-1}}$$

As for the size of the group of junkies, we may write this as the share of junkies last year, subtract those leave (exogenously determined; $_1$% leave the group every year) and those who enter (every year $_2$% of the yuppies become junkies). The aggregate dynamics for the size of the junkie group is then:

$$(11)\quad j_t = j_{t-1} - d_1 j_{t-1} + d_2 y_{t-1}$$

Some results

If one tries to solve () and () as a system of difference equations one will soon discover that this is rather difficult since there are no standard solution techniques for non-linear difference equations. It is however, possible to find potential stationary equilibria by setting inflow equal to outflow:

$$(12)\quad b_1 y = b_2 \left(\frac{y}{j + y}\right)$$

$$(13)\quad d_1 j = d_2 y$$

It is then easy to find that one stationary state is given by:

$$(14)\quad y^* = \frac{b_2 d_1}{b_1 (d_1 + d_2)}$$

[1] Assuming a constant g (the new generation as a share of the old) is not as innocent as it sounds. Since there is a new generation every year the outflow must be equal to g every year in order for g to be constant. If there were no deaths then each new generation would represent a smaller and smaller share of the existing population (i.e. g would not be constant). Hence, when assuming g is constant I also have to include deaths in the model. The total number of deaths every year should be equal to g, but I also believe I have to make some assumption about the distribution of death across different groups. Maybe that the same percentage dies in every group (although implausible this is perhaps what I have to assume or?). I have already assumed a certain percentage that leaves both groups (junkies and juppies) every year, so deaths cannot be larger than this. I have not specified the percentage deaths among abstainers or the proportion of exiting junkies who die relative to those who are treated. It is possible I have to be more explicit about this, but I was hoping that it was possible to sidestep the issue by arguing first that this is a minor problem (the effects on the relative shares of junkies/abstainers/yuppies may be small) and/or that it can be assumed using the “free” variables that I have not specified (mentioned above). I am not entirely happy with this and will try to take a closer look at it later.

$$(15)\quad j^* = \frac{\boldsymbol{b}_2 \boldsymbol{d}_2}{\boldsymbol{b}_1(\boldsymbol{d}_1 + \boldsymbol{d}_2)}$$

There is also another stationary state, namely y*=0 and j*=0.

The main interesting feature of the solutions so far emerges if we examine the total number of users (*d=j+y*). In equilibrium this is a very user friendly expression:

$$(16)\quad d^* = \frac{\boldsymbol{b}_2}{\boldsymbol{b}_1}$$

What is interesting about this, however, is that the parameter $_1$ does not appear in the solution to *d**. This is interesting because $_1$ might be interpreted as a policy controlled parameter. That is, the number of people that "exit" from the junkie group every year depends on how much money the authorities choose to spend on treatment. More money spent on treatment means a higher exit rate (a higher $_1$). One might expect that more treatment for junkies would reduce the overall number of drug users, but in my model this is not true as one can see from the fact that $_1$ does not appear in the solution for *d**. Given the non-intuitive nature of this conclusion it requires some elaboration.

Removing one user (a junkie) immediately reduced the number of drug users by one. However, because the number of junkies is reduced, the fear of using drugs among potential users is also reduced (they no longer see as many junkies as before). When the fear of using drugs is reduced, more people will start to use drugs and this increases the number of users. And, although not intentionally designed to do so, in my model the two effects balance each other exactly. That is: Eliminating one junkie will result in one more yuppie and the net effect on the number of drug users is zero.

Having noticed that $_1$ is not a part of the solution to d*, we might look closer at exactly the variables that enter the solution. Previously $_2$ was defined as: $\boldsymbol{b}_2 = f\left(U(J) - U(Y), g\right)$, so:

$$(17)\quad d^* = \frac{f(U(J) - U(Y), g)}{\boldsymbol{b}_1}$$

Or, as one would expect, the share of drug users depends on the utility difference between a yappie career and a junkie career.

Extensions, criticism and moving towards evolutionary game theory

Extension

One of the extensions that I have tried to work with is to experiment with different ways of including social stigma. So far I have simply assumed that it is uniformly distributed. One extension could be to assume a normal distribution. Another (not mutually exclusive) extension could be to make the stigma dependent on the number of people engaged in the activity. This could be done linearly as in:

(18) $m_{it} = k_i\,(1-d_t)$

In which case the social stigma cost of using drugs is reduced the more people who do so. This is probably plausible for some interval, but the effect could also be modelled like this:

(19) $m_{it} = k_i\,(0.5-d_t)$

In which case we would have a kind of "conformity" effect. People would tend to do what the majority does. (The expression would be negative when more than 50% use drugs and this means that in the expression for expected utility of drugs the moral cost becomes a moral benefit since subtracting a negative number is the same as adding). At least locally speaking it does not seem implausible to argue that in a group where almost everybody uses drugs there is a "social" pressure towards using drugs and not the other way around. When experimenting with these formulations I soon faced a new problem: multiple equilibria. In itself this is not a problem: if the context is one in which multiple equilibria really can (or do) exist, then it is only an advantage if our model reflects this. If however, some of the equilibria are less likely to appear, then we might try to find solution concepts that eliminate these.

Criticism

After presenting the model described above, I am often asked whether I really believe that potential addicts "calculate" the costs and benefits and based their decision on this calculation. The underlying criticism being that the decision to use drugs cannot be modelled as a "rational choice." There are several possible reactions to this line of criticism. One would be to argue that the model is an "ideal type" that simply exaggerates a tendency. That is, all I need is that people at least to some extent consider the incentives they are facing before using drugs. One might also try to defend the assumption using Friedman "as-if" justification – arguing that even if people do not calculate the way described above it often seems like their behviour can be well described as if they did. One might also, of course, stand hard and argue that people really do calculate. For instance, I recently read a report about how the number of births in Norway was highly correlated with "the profitability" of giving birth on special dates. In short, in the past ten years the Norwegian laws about maternity benefits and so on has been changed four times. On all four occasions there was a significant difference in the number of births in the "profitable" direction AND the law was not announced nine months in advance so there was little possibility of "planning" to be profitable in that sense. All in all, it seems like women, consciously or unconsciously was responsive to incentives even when it comes to such an "involuntary" matter as the date of giving birth. For better or worse, however, people who dislike rational choice models have probably heard all of these arguments and are unlikely to be moved much. My best reply, then, would be to create a model based on weaker assumptions about rationality and examine if my results still hold.

Possible answer: Use Evolutionary Game Theory

I am recently new to the field of evolutionary game theory, but it seems at least to be concerned about the same topics mentioned above – both the problem of multiple equilibria and the problem of excessively strong rationality and information assumptions. For instance, Ritzberger and Weibull (1995, 1371) note that "the rationalistic foundation of this [Nash equilibrium] approach is quite demanding" and among the alternatives they argue that "particularly promising seems the approach taken in evolutionary game theory." In the rest of this paper I will then try to reformulate my model in an evolutionary frame. Based on this I will end by comparing the two approaches and trying to be slightly critical of the claims of evolutionary game theory.

Stability and equilibria selection: Using EGT as an add-on!

In the model as it is there are already two stationary equilibria. First a situation in which there are no drug users ($j^*=0$, $d^*=0$), and – second - the more interesting solution developed above. It is rather obvious that (0,0) is not a very stable equilibria. It is not stable in the sense that as soon as somebody somehow (trembeling hand!) start to use drugs, the dynamics is such that we at once move farther and farther away from (0,0). In other words; a small deviation from that equilibria leads us away from it. The opposite is the case with the other equilibria. If – somehow, too many people use drugs (more than the equilibrium level), the level be reduced since many in the new generation will be very "scared" given the high level of junkies. Hence the second equilibrium is stable, the first is not. And, conversely, when few people use drugs the level of "fright" is low and more people will start. (see Appendix 1 for the results of a computer simulation using some arbitrary starting values).

Hence, my initial hope that the concept of evolutionary stability could help me in solving problems of multiple equilibria has not really been tested yet. One might say that my dismissal of the zero-solution was based more on evolutionary arguments than on Nash demands. This may be true, but I have some small queries. For something to be evolutionary stable it has to be a Nash equilibrium (the strategy is the best response given the other strategies) AND it has to be "mutant resistant" in the sense that the strategy has to give the highest payoff even if the population is invaded by some small fraction who also plays your ("mutant") strategy. I am unsire whether this is enough and I am also unsure about whether it is necessary. It is not necessary in the sense that I could just use Selten's "trembeling hand" refinement of Nash in order to eliminate the zero-option. And the ESS concept may not be enough in the sense that in a zero equilibrium there is no drug users to imitate (if the underlying logic in the evolutionary setup is that agents imitate the ones with higher payoff). If the underlying logic is that people follow a "trial and error" approach it is easier to use ESS to dismiss the zero-solution. But, this all points to the need for specifying the whole model in en evolutionary frame instead of just trying to add it after working out the equilibria (at which point you would probably say: of course!).

An evolutionary setup

Given my limited knowledge and experience with evolutionary game theory (EGT), I will simplify the model as much as possible. Hence, while the original model contained some individual heterogeneity, I will avoid this in the evolutionary frame. In fact, the only reson I introduced heterogeneity (different sensitivity to social stigma) in the original model, was the desire to avoid that I would only get corner solution (either everybody or nobody start to use drugs). In fact, one might consider it a strength of EGT that it is possible to end up in equilibrium with mixed populations (both users and non-users) even if people's underlying preferences are identical. In any case, there is at least no need to introduce heterogeneity of the kind I had in order to get a mixed equilibrium.

How do we get a mixed population equilibrium? EGT simply interprets the probabilities placed on the various strategies in a mixed strategy as expression of the share of people in the population that will play that strategy. This is the so called "mass-action" interpretation of Nash equilibrium points.

What are the strategies and their payoffs? There are only two possible strategies: Either you use drugs or you do not. Previously I assumed that choosing not to experiment with drugs yielded a payoff of EU(A), which later was assumed to be a constant (in which case I should write U(A) instead of EU(A)) And the expected yield from experimenting with drugs was:

(20) $EU_{it}(D) = p_t U(J) + (1 - p_t)U(Y) - m_i$

I have to reformulate this for several reasons. First of all I want to include the social stigma effect and this means that the payoff from "not using drug" can no longer be assumed constant. Instead it varies with the share of the population that use drugs. Indeed, this seems to be one of the strong points of EGT – its ability to model behaviour where the payoff depends on the share of the population that use the other strategies.

Second, whereas I previously assumed that individuals estimated p_t (the probability of becoming a junkie) by the share of current junkies out of the total number of current users ($j+y$, or simply d), this seems wrong in an EGT frame. The relevant parameter is the *actual* payoff received, not what individuals mistakenly expect based on incorrect estimation of p. It is the actual payoff that will make other people imitate (or avoid) a strategy. And, if we assume standard "replicator dynamics" the growth rate of the population using or not using drugs depends on the difference between the (actual) payoff of the current strategy and the average (actual) payoff of the alternative strategy. In short, given that the payoff should reflect actual values, it seems wrong to use the same definition in the EGT frame as I did in the original model.

The problem is probably more deeply rooted. Whereas I originally put quite a lot of effort into creation good microfoundations for the original model (and justifying the linear relationship), EGT tend to assume that there is some kind of process in the background (imitation, conscious or unconscious learning or trial and error) that generate a replicator dynamics. On the one hand this is strength in the sense that the aggregate results are compatible with many different micro-stories. That is, it has been shown that the replicator dynamics applies for several types of underlying processes (such as the ones mentioned above) given some additional assumption (about how often people evaluate their decisions and so on; for instance R+W assume it is a poisson process). Relying on the replicator dynamics is also a strength in the sense that within EGT there is a set of tools ready to be used so you do not have to invent the wheel every time you create a model. On the other hand, it might not be equally well suited for all kinds of models and purposes. In the original model the main force "driving" the result was exactly the "fear" and how it changed over time (depending on the levels of yuppies and junkies). Whether the fear really was well grounded or true was not the question. Moreover, by deriving the microfoundation myself I made the connections more precise, explicit and I did not have to use the assumptions needed to produce the replicator dynamics. I am not arguing that one is better than the other, just noting that the various approaches has pros and cons and that these varies depending on what you want with the model.

Finally, the last blow to the reformulation of the model in an EGT frame occurred when I tried to pursue the "generational" interpretation. Previously I assume that every new generation made up their minds about drugs and that it was a once upon a lifetime decision. It might be possible to do this in EGT.

The lesson then, is to use EGT only to examine the stigma mechanism without pretending to be able to translate the whole model. I then have the following strategies and payoffs (assuming the social stigma mechanism mentioned previously:

Strategy	**Payoff**
Abstain (A) (i.e. Not experiment with drugs)	U(A)+(0.5-d)
Drugs (D) (i.e. Experimenting with drugs)	pU(J)+(1-p)U(Y)-(0.5-d)

In equilibrium the payoffs must be equal (since only there is there no incentive to change strategy), which implies that:

$$(21)\ \ d^* = \frac{1}{2}(1 + U(A) - pU(J) - (1-p)U(Y))$$

I cannot really make up my mind whether this is plain silly or whether there really is a lesson in that conclusion. First, some restrictions clearly have to be put on the constants so that the share of drug users does not fall below zero or above 1. Second, the sign of U(A) seems wrong (but this depends on d again). It also seems difficult to compare this to the previous model (since b1 b2, d1 or d2 does not enter into this model), but this could be fixed with more complicated expressions for utility. As mentioned a junkie career consists of some years as a happy user, some times as an unhappy user and finally (on average) some time as an abstainer (after being a user). That is:

$$(22)\ \ U(JC) = \lambda^{J}U(J) + \lambda^{Y}U(Y) + \lambda^{JA}U(JA)$$

$$(23)\ \ U(YC) = \lambda^{Y}U(Y) + \lambda^{JA}U(YA)$$

In short: The total expected utility of a junkie career - U(JC) – is the sum of years in the various states (first as yappie, then junkie, then abstainer after being a junkie) multiplied by the annual utility of being in that state. The total expected utility of a yuppie career is simply the number of years as a yuppie multiplied by the annual payoff in that state added to the annual payoff of being an abstainer after first having been a yuppie multiplied by the number of years in that state.[2] Note also that I have takes some liberty with the notation here: U(J) and U(Y) now represents annual utility, while total utility - what I previously labelled U(J) and U(Y) - is now labelled (JC) and U(YC).

Using the same set of assumptions as before about inflow and outflow rates, the average time as a yuppie for a person who enters the yuppie group is:

$$(24)\ \ \lambda^{Y} = \frac{1}{\beta_1}$$

Similarly, the average time (years) as a junkie (given that you become a junkie) is:

$$(25)\ \ \lambda^{J} = \frac{1}{\delta_1}$$

We the

Strategy	**Payoff**
Abstain (A) (i.e. Not experiment with drugs)	*U(A)+(0.5-d)*
Drugs (D) (i.e. Experimenting with drugs)	$pU(J)\frac{1}{\delta_1} + (1-p)U(Y)\frac{1}{\beta_1} - (0.5-d)$

And the solutions becomes:

[2] Note that it would be easy to include discounting at this point, but so far I do not see that it adds much to the analysis of social interaction. Although on second thought it could have a great impact in the sense that the state of the "scary" addicts may be believed by potential addicts to be far ahead in time and hence discounting reduces the importance of the fear factor.

$$(26)\quad d^* = \frac{1}{2}\left(1 + U(A) - pU(J)\frac{1}{\boldsymbol{d}_1} - (1-p)U(Y)\frac{1}{\boldsymbol{b}_1}\right)$$

One could also eliminate p by examining the inflow parameters. We know the average number of years spent as a yuppie and we know the percentage of yuppies that every year become junkies. This allows us to work out the probability that the individual will become a junkie. I have not done so, however, because I already have a solution in which treatment effort (increasing $\boldsymbol{d}_1$) reduces the number of drug users. However, it is a very different model since the main driving force is no longer the "fear" but "social stigma." It does, maybe, point to a need to use more complex utility functions in the original frame. That is, I should formulate it in a way that allows a person to ask questions like "If you increase the average time period by some measure how will this affect the equilibrium size of the groups?"

Conclusion

My attempt to reformulate the model in evolutionary game theory was unsuccessful. While the old frame seemed good at handling the "fear" mechanism, the evolutionary frame (within my limited abilities) was not. On the other hand, the evolutionary frame could much easier handle the "social stigma" mechanism. I will continue to work on both in order to create a better unified treatment of the two.

References

Conlisk, J. (1996): "Why bounded rationality?", *Journal of Economic Literature*, 34, 669-700.

Kirman, A.P. (1992): "Whom or what does the representative agent represent?", *Journal of Economic Perspectives*, 6, 117-136.

Michihiro, K. (1996) "Evolutionary game theory in economics" Ch. 8 in (eds.) Kreps, David M & Wallis, Kenneth F. : *Advances in economics and econometrics* . Cambridge: Cambridge University Press.

Ritzberger, Klaus & Weibull, Jörgen W. (1995) "Evolutionary selection in normal-form games", *Econometrica* 63: 1371-1399.

Schelling, T.C. (1978) *Micromotives and macrobehavior*. New York : W. W. Norton.

Appendix 1:

I ran some simulations to examine the speed of convergence. The figure below shows the result from one such simulation in which the starting point was set as follows:

60000	Y0 (yuppies)
10000	J0 (junkies)
0,142857	Initial probability of becoming a junkie
70000	Total drug users
30000	(b1)
0,2	(b2)
0,05	(d1)
0,07	(d2)

I wanted to examine how fast the system converged upon equilibrium so I introduced a "shock" after about 50 periods. The number of junkies was then about 62 000 and I just assumed an exogenous shock that reduced their numbers to 50 000. It turned out that the system returned to (the same) (almost) equilibrium after about 25 years.