Long Term Implications of Drug Policy Shifts: Anticipating and Non Anticipating Consumers

J.P. Caulkins, G. Feichtinger, R.F. Hartl,
P.M. Kort, A.J. Novak, A. Seidl,

Long Term Implications of Drug Policy Shifts: Anticipating and Non Anticipating Consumers

Jonathan P. Caulkins   Gustav Feichtinger   Richard F. Hartl

Peter M. Kort   Andreas J. Novak   Andrea Seidl

November 29, 2011

Abstract

We consider a semi-rational addiction model in which the user has perfect foresight over all things within the user's control, but not necessarily with respect to exogenous parameter shocks, e.g., those stemming from changes in national policy. We show that addictive substances are more likely to have state-dependent solution trajectories, and that in turn can create path dependence at the macro policy level; in particular, legalization may be an irreversible experiment. Also, in this model, shifting from a nuanced policy that differentiates between high and low intensity users, to a tougher one where the government makes life hard for every user reduces initiation considerably. Or may have perverse effects. In particular, we show that making the policy tougher in this way could drive some people from a "happy" stable saddle point equilibrium with moderate consumption into increasing rather than reducing their consumption and addiction stock. So implementing zero tolerance policies may increase rather than reduce aggregate drug use, depending on the population's distribution of parameter values and initial consumption stocks. Further, we consider the impact of announcing a policy change.

1 Introduction

There is a growing literature that addresses how rational agents deal with an addictive consumer good. Important early contributions to this theory of rational addiction are for example Ryder and Heal (1973); Stigler and Becker (1977); Iannaccone (1986) and Becker and Murphy (1988). Further important works in this area include Becker (1992); Dockner and Feichtinger (1993); Orphanides and Zervos (1994, 1995, 1998); see also Braun and Vanini (2003) and Gavrila et al. (2005).

To date the underlying system has always been assumed to evolve continuously, and typically with the external context remaining constant and only the user's consumption stock changing. In reality, however, sudden discontinuous changes in the environment may occur, for example those caused by a change in drug policy. The purpose of this paper is to provide insights about how consumers of addictive goods may react to policy changes and consider the resulting implications for a policy maker.

We consider a minor variant of the standard Becker and Murphy (1988) rational addiction theory (RAT) model but extended to allow for the possibility of an unanticipated

change in certain parameters describing the surrounding environment. By unanticipated change we mean that a consumer cannot foresee the implementation of a new policy ahead of time, or at least does not believe in its occurrence. This leads to a two-stage formulation of the user's dynamic optimization problem. We primarily consider what happens if parameters describing the attractiveness of consumption suddenly decrease or increase at an exogenously determined time.

The parameter change could be the result of macro forces or political processes that an individual consumer cannot forecast. An example would be the sudden implementation of a rule by a government, employer or building owner that one cannot smoke inside, making smoking much less attractive. A parallel example in the context of illegal drugs would be a change in sentencing policy. However, one can also imagine unanticipated shocks of other sorts – e.g., a sudden supply shortage as occurred with the Australian heroin drought (see Degenhardt et al., 2005), an overdose that makes the drug less pleasant to use, or an arrest leading to loss of a job and income, thereby increasing the opportunity cost of spending money on the drug because there is less income to cover rent and food. The user being arrested is a particularly interesting scenario because repeat offender laws mean that adding another arrest to the user's criminal history can change the expected sanction per unit of subsequent consumption. Even if the user correctly foresaw the expected number of future arrests, the timing is random. (Indeed, the classic modeling of the arrival of arrests within an individual's career is via a Poisson Process; compare Maltz, 1996.)

The basic finding is that allowing users of addictive goods to be surprised – as opposed to having Godlike omniscience – allows for a range of interesting and arguably more realistic behaviors.

We then contrast results when consumers do not anticipate the future policy switch with a scenario in which they do. These scenarios might be thought of as (1) perfect foresight regarding individual/local actions and outcomes but ignorance about future random shocks and policy actions vs. (2) perfect foresight in the classical sense of total omniscience. One of our aims is to find implications of consumer anticipation on consumer utility and on policy performance. It is interesting to ask whether a given policy performs "better" by some policy metric when consumers anticipate the change in policy or not. If so, then the policy maker might prefer to announce the policy switch beforehand.

Finally, we consider the possibility that the user forms expectations about future policy changes, but those expectations may not then be fulfilled. That is, the user attempts to anticipate a policy change, but guesses wrong. We particularly look at what can happen if a policy maker (credibly) announces a policy change for some point in the future and then to implement the change. It turns out that such an unfulfilled announcement can, in certain circumstances, lead to permanent changes in the long-run equilibrium use and consumption stock.

2   The Model

The one-stage variant of the model we consider was presented in Caulkins et al. (2011) and is based on Becker and Murphy (1988). In this model the consumer of an addictive good wants to maximize his or her utility which depends on the accumulated consumption stock

$S(t)$, as well as the current consumption rate $c(t)$, which we normalize to $0 \leq c(t) \leq 1$. This upper bound is motivated by the fact that a sufficiently large dose of a substance such as heroin or cocaine can be fatal; see Gable (2004). So putting an upper bound on c can be construed as solving the optimal control problem for someone who does not want to die.

In the sequel we omit time argument t. Like Becker and Murphy (1988, p. 679, eq. (7)) we consider a quadratic utility function

$$u_i(c, S) = \alpha_{ci}c + \frac{\alpha_{cci}}{2}c^2 + \alpha_{Si}S + \frac{\alpha_{SSi}}{2}S^2 + \alpha_{cSi}cS,$$

where the index $i = 1, 2$ is used to distinguish between the two stages. We assume that the two stages only differ in the values of the parameters.

Like Caulkins et al. (2011) we assume that the signs of the coefficients are as follows: $\alpha_{ci} > 0$, $\alpha_{cci} < 0$, $\alpha_{Si} < 0$, $\alpha_{SSi} < 0$, $\alpha_{cSi} > 0$, meaning that a consumer only derives positive utility from current consumption if it is not too large (i.e. $\alpha_{ci} > -\alpha_{cci}c$). We restrict ourselves to considering harmful addiction; thus, it makes sense to assume a negative utility of the accumulated consumption S. By the RAT definition of addictive behavior (see, e.g., Iannaccone, 1986; Dockner and Feichtinger, 1993) the marginal utility of current consumption increases with past consumption.

As is customary in these models, the state equation is

$$\dot{S} = c - \delta S,$$

where the δ denotes the exogenously given, constant rate of disappearance of physical and mental effects of past consumption.

In each stage the consumer acts as if this stage endures forever (i.e. chooses the solution path with the highest objective value leading to one of the steady states of the current stage) and only adapts this strategy when the policy change takes place at time τ. For this kind of problem, two-stage methods actually are not needed. Solving one-stage models for each of the sets of parameter values and then jumping from stage 1's optimal solution trajectory to stage 2's trajectory at the exogenously given time is sufficient to characterize the solution. But to state the problem more formally, in the first stage the consumer is confronted with the optimization problem

$$\max_c \int_0^\infty e^{-rt} u_1(c, S) dt,$$
$$\text{s.t.} \quad \dot{S} = c - \delta S,$$
$$0 \leq c \leq 1, \qquad S(0) = S_0,$$

where $u_1(c, S)$ is the utility of the consumer in stage 1 and $r > 0$ denotes the discount rate. Then in stage 2, the consumer faces following problem

$$\max_c \left[\int_0^\tau e^{-rt} u_1(c_1^*(t), S_1^*(t)) dt + \int_\tau^\infty e^{-rt} u_2(c, S) dt \right],$$
$$\text{s.t.} \quad \dot{S} = c - \delta S,$$
$$0 \leq c \leq 1,$$

where $(c_1^*(t), S_1^*(t))$ denotes the optimal solution path determined by the stage 1 problem. Note that one can easily extend this approach in a similar manner to three and more stages.

The current value Hamiltonian in stage i is

$$\mathcal{H}_i = \alpha_{ci}c + \frac{\alpha_{cci}}{2}c^2 + \alpha_{Si}S + \frac{\alpha_{SSi}}{2}S^2 + \alpha_{cSi}cS + \lambda_i(c - \delta S),$$

where λ_i denotes the costate variable. By applying Pontryagin's maximum principle, see e.g. Grass et al. (2008), we find that

$$c^* = -\frac{\alpha_{ci} + \alpha_{cSi}S + \lambda_i}{\alpha_{cci}},$$

meaning that the optimal control c^* is within the interior of the admissible region ($0 < c^* < 1$) iff $-\alpha_{cSi}S - \alpha_{ci} < \lambda_i < -\alpha_{cSi}S - \alpha_{ci} - \alpha_{cci}$. Note that in general the optimal control as well as the costate will be discontinuous at the switching point if the consumer cannot anticipate the change.

The costate equation is

$$\dot{\lambda}_i = (r + \delta)\lambda_i - \alpha_{Si} - \alpha_{SSi}S - \alpha_{cSi}c.$$

The Legendre-Clebsch condition is fulfilled as $\alpha_{cci} < 0$. The sufficiency conditions are fulfilled iff $\alpha_{cci} < 0$, $\alpha_{SSi} < 0$ and $\alpha_{cSi}^2 < \alpha_{cci}\alpha_{SSi}$ as the Hamiltonian is jointly concave in state and control then.

It is possible to find one interior steady state with

$$\hat{S}_i = \frac{-\alpha_{Si} - \alpha_{ci}(r + \delta)}{(r + 2\delta)\alpha_{cSi} + \delta\alpha_{cci}(r + \delta) + \alpha_{SSi}}$$

and two boundary steady states with

$$\left(\hat{S}_i^0 = 0, \hat{\lambda}_i^0 = \frac{\alpha_{Si}}{r + \delta}, \hat{c}_i^0 = 0\right) \qquad \text{and} \qquad \left(\hat{S}_i^1 = \frac{1}{\delta}, \hat{\lambda}_i^1 = \frac{\delta(\alpha_{cSi} + \alpha_{Si}) + \alpha_{SSi}}{\delta(r + \delta)}, \hat{c}_i^1 = 1\right),$$

See Caulkins et al. (2011) for the derivation of these expressions and for a closer consideration of the feasibility as well as the stability properties of these steady states.

Since in our model the utility maximizing consumer cannot anticipate the change of the system, he or she acts in each of the two stages as if the stage would last forever. In this paper we focus on policy measures that make drug consumption more or less attractive. Thus, similarly to Caulkins et al. (2011), we can derive a 2 X 2 table (Table 1) that distinguishes between cases in which current consumption is or is not highly rewarding (the columns) and the substance is or is not highly addictive (the rows) in the sense that past consumption strongly enhances the reward from current consumption. The different regions, which can be distinguished by the number and admissibility of the steady states, have implications for how a consumer would act in stage $i = 1, 2$.

	consumption is	
	more attractive $\alpha_{ci} > -\alpha_{Si}/(r+\delta)$	less attractive $\alpha_{ci} < -\alpha_{Si}/(r+\delta)$
low addiction $\alpha_{cSi} < \frac{1}{r+2\delta}[\delta(-\alpha_{cci}-\alpha_{ci})(r+\delta) -\delta\alpha_{Si}-\alpha_{SSi}]$	Region 1	Region 3
high addiction $\alpha_{cSi} > \frac{1}{r+2\delta}[\delta(-\alpha_{cci}-\alpha_{ci})(r+\delta) -\delta\alpha_{Si}-\alpha_{SSi}]$	Region 2	Region 4

Table 1: Parameter regions depending on different parameters

Region 1 If consumption is attractive and not very addictive then only the interior saddle point $\hat{S}_i$ is admissible; i.e., the drug consumer would always target some intermediate level of use.

Region 2 If consumption is attractive and addictive then only the steady state at the boundary with maximum consumption $\hat{S}_i^1$ is feasible; i.e., in the long-run the user's habit would always grow to reach the maximum possible non-lethal rate of consumption.

Region 3 If consumption is neither attractive nor addictive, then only the boundary steady state with no consumption $\hat{S}_i^0$ is admissible; i.e., the consumer would always reduce consumption to zero over time.

Region 4 The most interesting case occurs when consumption is addictive but not very attractive. Then all three of the steady states are admissible. This means the optimal behavior might be history-dependent[1]; i.e., it depends on the state value (consumption stock) at time zero and at the time of the policy change. Then so-called Skiba points (also known as indifference-threshold or DNSS points; see Grass et al., 2008) can be found, where a user is indifferent between approaching $\hat{S}_i^1$ or $\hat{S}_i^0$.

These cases help explain what otherwise might look like a paradox concerning the location of the interior steady state. If consumption is less attractive (right hand column of Table 1) then the numerator of the expression for the interior steady state is positive, and the steady state is only relevant if the denominator is positive as well. In that case, increasing the consequences of the addiction stock (α_S larger in absolute value) or reducing the pleasures of current use (α_c smaller) both move the interior steady state to the right – which tends to increase the range of initial values that are to the left of the Skiba point, meaning there is a larger range of initial addition stocks for which eventual abstinence is optimal. On the other hand, if the good is more attractive (left hand column of Table 1), the interior steady state is only relevant if both the numerator and the denominator are negative. Then the same parameter changes (consumption is less pleasurable and/or addiction stock is more consequential) move the steady state to

---

[1]Note that even though all three steady states are admissible, for certain parameters the solution path leading to one of the boundary steady states might be better than the solution path leading to the other boundary steady state for any initial state value. The interior steady state is always unstable, thus it can be excluded as long term outcome.

Type	r	δ	α_{c1}	α_{cc1}	α_{S1}	α_{cS1}	α_{SS1}
Person 1a	0.05	0.1	5	-10	-0.6	1.15	-0.15625
Person 1b	0.05	0.1	4.5	-10	-0.6	1.15	-0.15625
Person 4	0.05	0.1	3.25	-10	-0.6	1.3	-0.15625

Table 2: Parameter values used for the numerical calculations

the left. But in that case the interior steady state, if it exists, is a saddle point and long-run equilibrium. So in either circumstance, those parameter shifts do have the expected effect of tending to reduce consumption, even though that is not so readily apparent from the equation for the interior steady state.

In the previous model formulation it was assumed that a drug consumer does not know that a policy change will take place. However, we also want to deal with the question of how a drug consumer would adapt if a policy maker credibly announced a coming change. In order to deal with this problem, we have to formulate our problem as a multi-stage model (see, e.g., Grass et al., 2008):

$$\max_c \int_0^\tau e^{-rt}u_1(c,S)dt + \int_\tau^\infty e^{-rt}u_2(c,S)dt$$
$$\text{s.t.} \quad \dot{S} = c - \delta S,$$
$$0 \leq c \leq 1,$$

where at the exogenously determined switching time τ the matching condition (see Tomiyama and Rossana, 1989; Makris, 2001)

$$\lambda_1(\tau) = \lambda_2(\tau)$$

has to hold. λ_i denotes the costate in stage $i = 1, 2$. The analytical expressions for the optimal control and the costate equations for the two stages are the same as before.

3 Numerical Results for Unforeseeable Policy Changes

This section derives implications of policy switches which are not being foreseen by consumers. Hence, before the policy switch occurs, consumers act as if the original policy will last forever.

We consider how different policy changes affect different types of consumers that can be distinguished by their utility of consumption, their degree of addiction as well as their disutility of addiction. The particular sets of parameters can be seen in Table 2, the type number refers to the corresponding case in Tab. 1.

3.1 Making Drug Consumption Less Attractive & Long Term Impact of Crackdowns

Fig. 1 shows what can happen if α_c decreases in the second stage, e.g., because of an increase in the probability of arrest per gram consumed or an increase in the expected
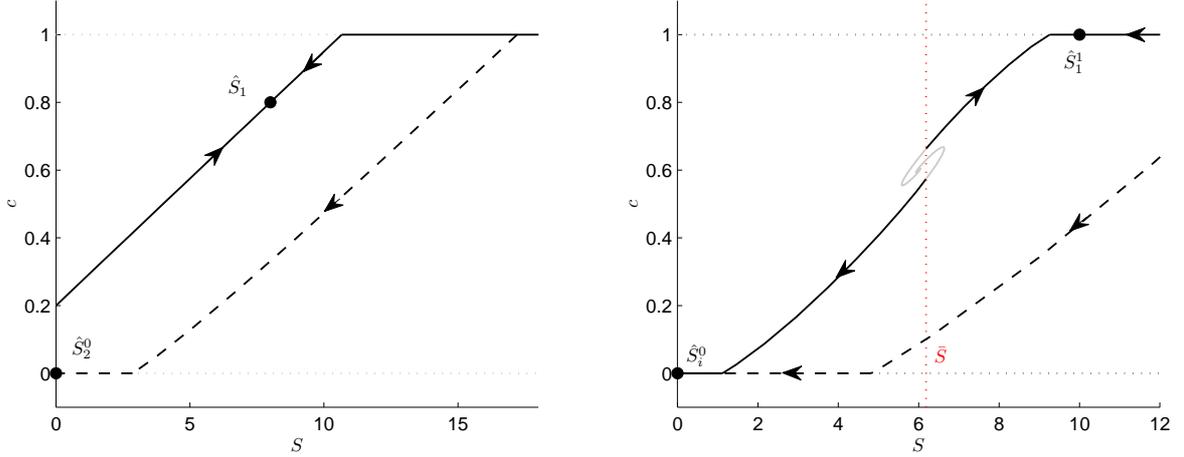
6

Figure 1: The solution path for Stage 1 is depicted by a black solid line, the parameters correspond to Person (type) 1a (left panel) of Table 2 and Person 4 (right panel). As soon as the policy changes reduces α_c by 2.5 (i.e. to $\alpha_{c2} = 2.5$ in the left panel and to $\alpha_{c2} = 0.75$ in the right panel), the consumer chooses the dashed solution path.

punishment given arrest, assuming arrest risk is proportional to the quantity consumed. Proportionality there makes sense inasmuch as many arrests occur around the purchase of the drug. Supply control interventions that drive up price would also reduce α_c if money makes a linear contribution to utility. The higher the price, the less money one has left over at any given current consumption level c.

In the left panel of Fig. 1 we can see following scenario: In the first stage, a consumer would always approach an interior saddle point. However, in the second stage, consumption becomes less attractive, so the consumer would decrease consumption and finally end up not consuming anything at all. The right panel depicts a scenario where the consumer would, depending on the initial accumulated consumption, either increase or decrease consumption, to approach a steady state with no or maximum non-lethal consumption. At the Skiba point $\bar{S}$ he or she has the choice between the two options. In the second stage, consumption gets so unattractive, that, no matter how large the state value at the time of the switch between stages, the user always reduces consumption and approaches the steady state with no consumption.

This result corresponds to the conventional case. Tougher enforcement makes use less attractive, so use declines, eventually to zero. Indeed, the policy change can even induce a 'cold turkey' response, which means that a drug consumer drastically decreases or even stops consumption even though that might not be pleasant in the short-run. Cold turkey here is triggered by the change of the parameter α_c making consumption much less attractive.

Let us now ask the question what happens if the reduced attractiveness of the drug is only temporary, e.g., is caused by some transitory crackdown and the user anticipates neither the onset nor the termination of the program. If a drug consumer assumes that this new, lower utility of drug consumption will remain forever, he or she would decrease consumption. However, as soon as the attractiveness of the drug rebounded to its original level, the user will again adopt the strategy indicated by the solid line. It is sometimes
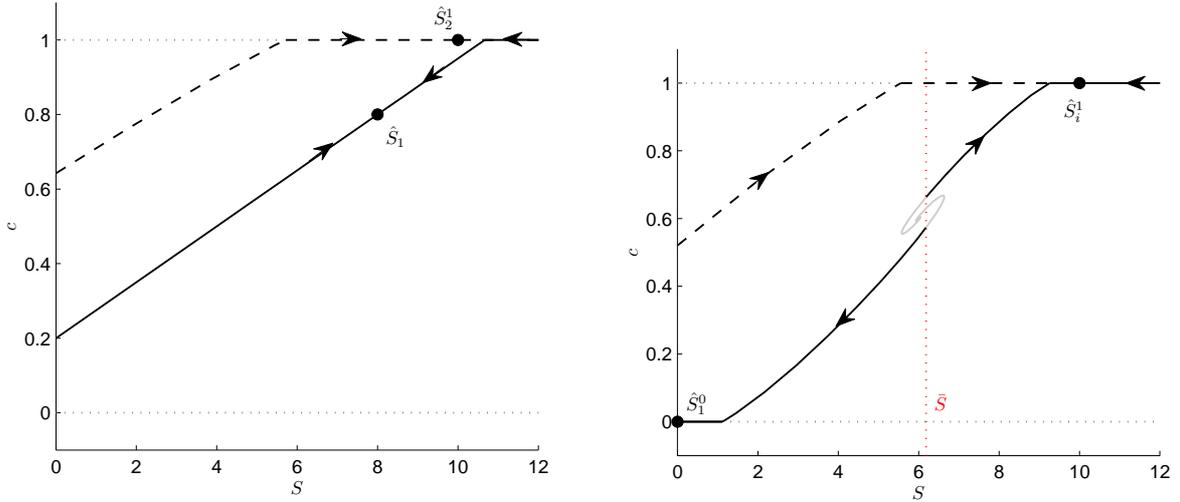
Figure 2: The solution path for Stage 1 is depicted by a black solid line, the parameters correspond to Person (type) 1a (left panel) of Table 2 and Person 4 (right panel). As soon as the policy changes increases α_c by 2.5 (i.e. to $\alpha_{c2} = 7.5$ in the left panel and to $\alpha_{c2} = 5.75$ in the right panel), the consumer chooses the dashed solution path.

asserted that temporary crackdowns can bring no lasting benefits. Indeed, returning to the solid line may mean returning to a higher level of use, but not necessarily. For the type of person depicted in the right panel of Fig. 1, this temporary measure can affect the consumption behavior so much, that if during this second stage the consumption stock S falls below level $\bar{S}$, the degree of addiction would be so weak in the third stage, that, after an upward jump, the user would further decrease consumption and become abstinent in the long run.

3.2 The Irreversibility of Legalization

Here we consider what happens if consumption becomes more not less attractive, e.g., due to legalization reducing prices and eliminating sanction risks so α_c increases. Figure 2 reveals that in this case a consumer would increase consumption and approach the steady state with the maximum consumption. Of particular interest is the right panel, where this policy change induces a consumer who might otherwise become abstinent in the long-term to instead start to use more. Thus, if a subsequent policy maker who opposed drug consumption wanted to undo this policy change and reestablish the stage 1 conditions, he or she would have to be careful not to wait too long. Otherwise the consumption stock would increase so much, that it would not be optimal for the consumer anymore to approach the steady state with no consumption, even if the legalization were repealed.

This scenario might also have implications concerning the impact of marketing measures related to legal addictive products like alcohol, tobacco, chocolate, coffee or junk food. Let us assume a firm increases the attractiveness of its product, e.g. by lowering its price, adding some free toy if the product is aimed at children, or promoting the status of using the product. If a firm can keep up this increased attractiveness forever, then even

people, who might otherwise have consumed this product only on a moderate level or not at all, might instead develop a larger consumption stock. If a marketing measure only increases the attractiveness for a limited amount of time, consumption only increases for a certain time and will get back to its original steady state level[2] – except for the type of consumer considered in the right panel of Fig. 2. In this case[3], people who would in the long run stop using the product, might be induced through this marketing measure to increase their consumption level so much that their addiction reaches a level that they would continue to consume more in the long-run when the attractiveness of the product goes back to its original level.

This means for example that a certain type of consumers - but not all - might be permanently influenced by temporary marketing measures. If one thinks that from the perspective of a child, the change in objectives that accompany becoming an adult is to some extent an unexpected change, then a politician trying to fight obesity by prohibiting fast food chains from adding toys to their unhealthy junk food meals for kids (as happened in San Francisco some time ago[4]) may have an intellectual basis for that market intervention.

3.3 Omitting the Distinction between Light and Heavy Consumers

This scenario is inspired by US drug policy history, specifically for cocaine. Early in a drug epidemic, there may be low dependence, lots of people initiating, and official policy that is relatively permissive. That was certainly true in the US in the 1970s when cocaine was publicly labeled as a "soft drug" that was not addictive. US policy got much tougher during the 1980s. This led to a massive increase in arrests, change of laws to make sentences much longer, and other laws were implemented to deny various social benefits to people who had been convicted for a drug offense. In our model the tougher policy translates into a smaller α_c. Section 3.1 shows how this affects consumption. The present section analyzes the implications when this policy is accompanied by not distinguishing so much between light and heavy users. In other words, before the policy change only the people that were heavy (frequent) users were punished, implying that α_c was rather high but α_{cc} was severely negative. However, after the policy change, "everybody" consuming drugs is punished, implying a large decrease in α_c, but also α_{cc} moving closer to zero because less distinction is made between light and heavy users. (It does not go all the way to zero because it still costs more money to buy a lot than a little of the drug, even if the enforcement risk were the same across users.)

Figure 3 shows such scenarios. We see that initiation is high during Stage 1, as was the case for cocaine in the U.S. in the late US 1970s. On the other hand, initiation is low during Stage 2, as when cocaine initiation plummeted after US drug policy got tougher in the 1980s.

---

[2]It does not play a big role in this context whether people expect the policy to change again back to the original one, here it would only affect the amount consumed; this can be easily shown by calculations similar to those in Section 4.

[3]Note that in this case the expectations of the consumer can be crucial; compare Section 4.2.

[4]See http://articles.latimes.com/2010/nov/02/business/la-fi-happy-meals-20101103 (last accessed Nov. 29th, 2011)
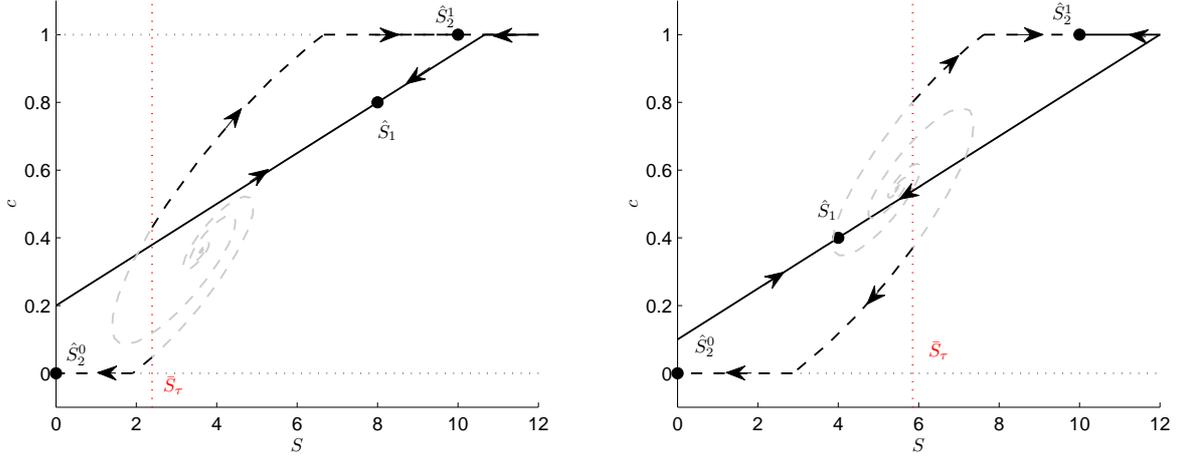
Figure 3: The solution path for Stage 1 is depicted by a black solid line, the parameters correspond to Person 1a (left panel) of Table 2 and Person 1b (right panel). As soon as the policy change leads to a decrease of α_c by 2 and an increase of α_{cc} by 4 (i.e., $\alpha_{c2} = 3$ and $\alpha_{cc2} = -6$ (left panel) or $\alpha_{c2} = 2.5$ and $\alpha_{cc2} = -6$ (right panel)), the consumer chooses the dashed solution path.

In the first stage the consumer would approach a saddle point. After the change of parameters, however, the consumer's strategy crucially depends on the level of accumulated consumption at the time τ of the change. If $S(\tau)$ is below some level $\bar{S}_\tau$ it is optimal to reduce and eventually quit consumption entirely. However, if $S(\tau)$ exceeds this level, then the consumer should increase consumption even beyond the stage 1 saddle, so this kind of policy change might be a mistake if the policy maker's objective is to reduce consumption: in the second stage the consumer increases consumption and ends up in a steady state with a higher accumulated level of consumption than in the first stage.

But what is probably the most interesting aspect of this scenario is that one can see that the same policy change can reduce consumption for one person but increase it for another. The left and the right panel of Fig. 3 depicts two consumers with only a slightly different utilities of consumption. Since for consumer type 1a in the left panel, $\alpha_{c1} = 5$ is higher than for consumer type 1b ($\alpha_{c1} = 4.5$), person 1a would consume more and, consequently, the steady state level of consumption stock is also higher in the first stage. Now let us assume that both consumers have reached their steady states and the policy is changed for both consumers at the same time and in the same way (i.e. α_c is decreased by 2 and α_{cc} increased by 4). We can see that consumer 1a finds himself now on the right side of the Skiba point, so he would increase his consumption, while consumer 1b finds herself on the left side of the Skiba point and would reduce her consumption. So, even though two persons differ only by 10% on one parameter, the same policy action produces diametrically opposed effects on the two different people. This means that the overall effect of the policy change on a population of people who have different parameters is ambiguous. This implies that the ultimate outcome depends on the distribution of parameter values across that population.

Fig. 4 reveals that for parameters such as those used for the phase portrait in Fig. 3 the switching time is crucial for the optimal long-run solution. We saw in Fig. 3 that when
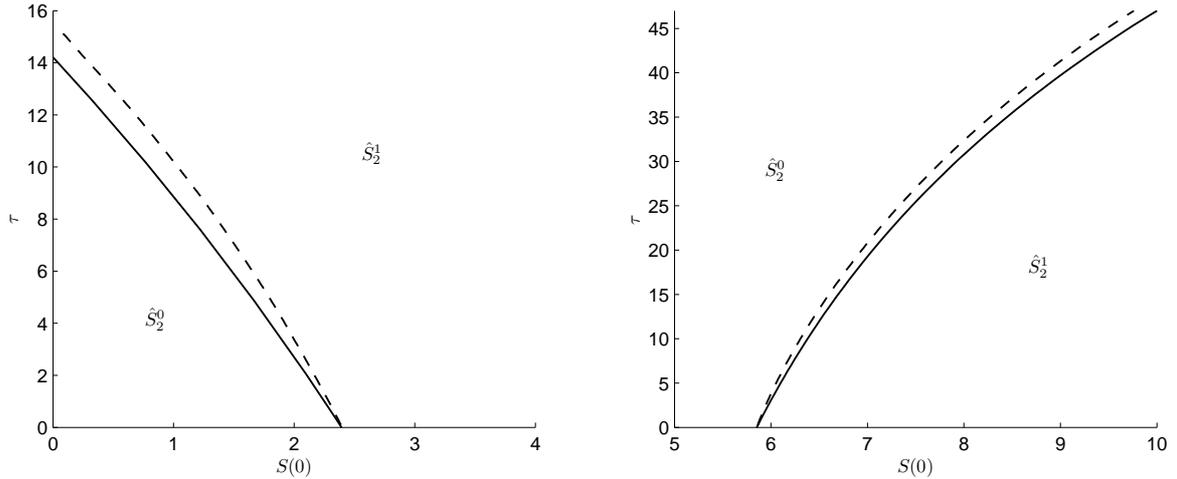
Figure 4: Skiba curve depending on initial state value and switching time corresponding to Fig. 3, i.e. the left panel corresponds to Person 1a and the right panel to Person 1b. $\hat{S}_2^j$, $j = 0, 1$ denotes the steady state the consumer would end up at for the given initial state value and switching time. At the red line the consumer is indifferent between the two options for the non-anticipation case, while the black, dashed line depicts the anticipation case.

both types of consumers have already reached an interior saddle point in the first stage, due to the same policy change one consumer might end up abstinent, with the other one consuming as much as possible. Now let us assume the drug consumers have not reached their steady state level at the moment of the policy switch. Fig. 4 shows how the long run outcome depends on the switching time for different initial values. For example in the left panel a decision maker has to be careful not to wait too long until implementing a new policy if the initial addiction level is low. The reason is that the incentive to consume drugs is larger in the first stage, so if the policy cannot be implemented fast enough, the consumption stock would become so big that in the second stage it would still be advantageous to further increase consumption. On the other hand when the policy is changed soon enough, a consumer would become abstinent in the long-run if the initial addiction level is low. However, if the initial degree of consumption stock is high, it does not matter when the policy is changed, the consumer would always consume as much as possible in the long-run. At the depicted curve itself, a drug consumer has the choice at the switching point whether to further increase or decrease the consumption of the drug. If we assume a policy maker wants drug consumers to become abstinent, it is interesting to see that for drug consumer type 1a the policy maker would have to be careful not to wait too long implementing the new policy if the initial state value is low. However, for consumer type 1b, when the initial addiction is high, a policy maker has to wait long enough before changing the policy. The reason is that consumption is not so advantageous in the first stage for for someone with a high consumption stock, and so a consumer would decrease consumption. Thus, we see even more clearly than before that while a certain type of drug policy change might lead to the desired result for one person, it might lead to the opposite outcome for another one.
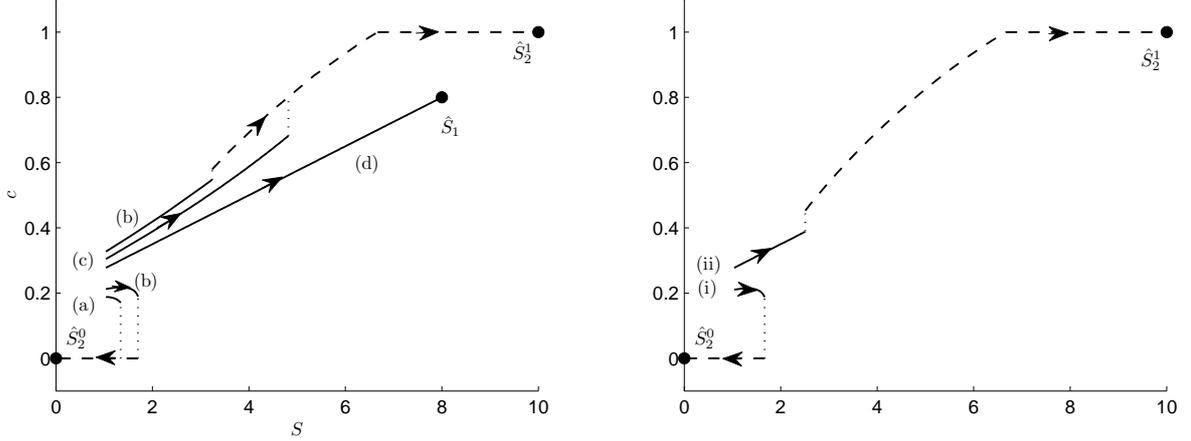
Figure 5: The solution paths for person type 1a for the initial state value $S_0 = 1.035$ if the consumer can anticipate the policy change for the exogenously determined switching time (a) $\tau = 5$, (b) $\tau = 10$ and (c) $\tau = 20$ (d) $\tau = \infty$ (left panel). The right panel depicts solution paths for $\tau = 9.5$ if the consumer does (i) and does not (ii) anticipate the policy change. The solid line denotes the solution of the first stage, while the dashed line shows the one in the second stage.

4   Effects of Announcing a Policy Change

In the following section we will consider how a consumer would adapt his behavior if the policy change is announced. This means that a consumer knows when a policy change is going to occur and how it will look like. We assume that the policy maker is credible and reliable (unless stated otherwise), i.e. the policy change occurs the way it is announced.

4.1   Adapting the Consumption Behavior

If one announces a future policy in a way that people believe it will occur, this can influence behavior today. The left panel of Fig. 5 shows this in a simple way. Again we consider Person 1a and a policy change where α_c decreases by 2 and α_{cc} increases by 4, meaning that in the second stage a tougher policy is implemented, where less distinction is made between light and heavy users. For the initial state value $S_0 = 1.035$ we can see the following: If the announced time of the policy change is early, the consumer would consume less the sooner the announced policy change will occur, while if the announced time is long, he or she would consume more. Effectively, this means that users prepare for the coming change by choosing consumption levels that are "closer" to those that they will head toward in the long run after the policy shock, and do so in a more concerted way the shorter the announced switching time is. Thus, if the policy shock is producing an outcome that is desirable to the policy maker, then the policy maker can purchase a little more of that benefit "for free" just by warning the users of the coming change. However, the longer the announced time is, the closer is the behavior to a solution where the consumer expects the policy never to change.

The parameters in this case are chosen in a manner that in the second stage both boundary steady states are admissible and are candidates for the optimal long run so-
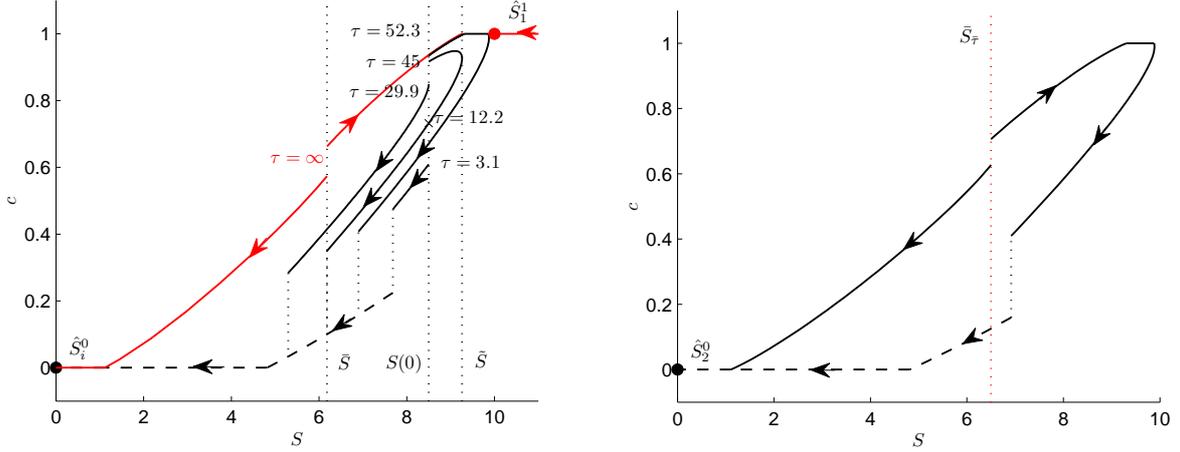
Figure 6: The solution paths for person type 4 for the initial state value $S_0 = 8.5$ if the consumer can anticipate the policy change for different exogenously determined switching times. For $\tau = 88.7$ (right panel) and $S_0 = 6.5$ the consumer is indifferent between two solution paths leading to the same steady state. The solid line denotes the solution of the first stage, while the dashed line shows the one in the second stage.

lution. For the switching time $\tau = 10$ we have a Skiba point starting at $S_0 = 1.035$: a consumer is indifferent between consuming little and ending up in the no addiction steady state and consuming more and becoming severely addicted. (This explains why solution (b) occurs twice.) In this sense this figure relates to Fig. 4. Like in the non-anticipation case, we can see how the final outcome depends on the announced switching time. Fig. 4 reveals that, comparing the case where a drug consumer can and cannot anticipate a policy change, the Skiba curve, where a consumer is indifferent between ending up in the low and high addiction steady state, shifts upwards if the consumer can anticipate the policy change. Thus, for certain initial state values and switching times, like for the example shown in the right panel of Fig. 5, if a policy maker acts without announcing a policy change, the user will converge to the high steady state, but if the policy maker makes the same policy change at the same τ, but does announce it ahead of time, then the user will converge to the lower steady state. This means that a (relatively inexpensive) communication strategy "buys" a very big change in the long run outcome "for free".

The Skiba point is also interesting for another reason. Due to the exogenous switching time the consumer is indifferent between two solution paths starting in the same direction, which is possible in this setup due to the exogenously given switching time. In a standard one-state optimal control model this cannot happen; see Caulkins et al. (2010). In addition, an exogenously given switching time can generate non-monotonic solution paths, which will be shown in the next section. This is especially remarkable, because Hartl (1987) proves that optimal trajectories are monotonic in autonomous one-state optimal control models.
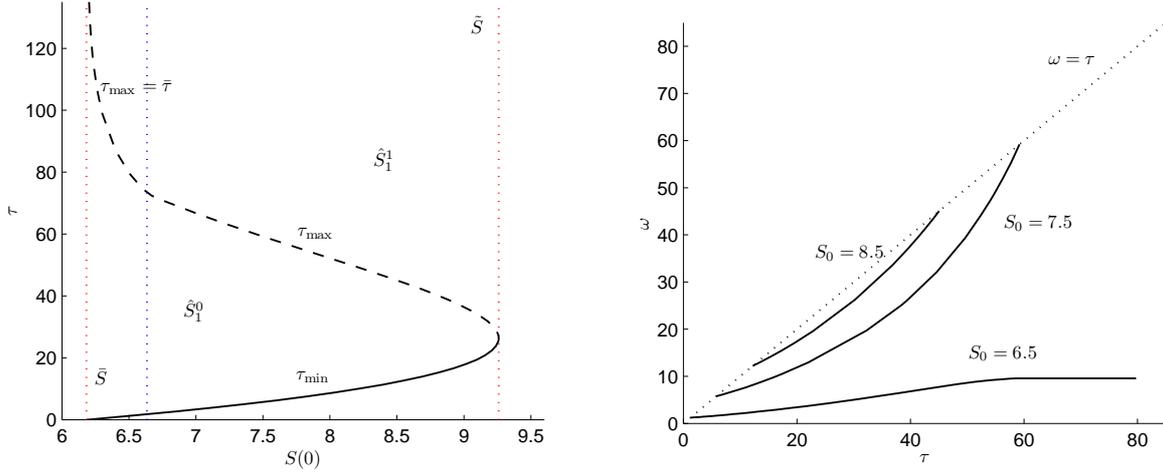
Figure 7: Minimum and maximum time for which a policy shift can be announced so that an anticipating consumer would change his behavior in such a manner that an actual policy change might not be necessary at all. The time, at which a consumer can realize that a new policy is not going to happen, ω, so that he would still move toward the abstinent steady state after realization, depends on the announced switching time (right panel).

4.2 Exploiting Anticipation Effects

The previous section showed that announcing a policy change can have a huge impact on which steady state a consumer ends up in. We now push this further and show that even just announcing a policy that to the user's surprise is not then executed can also have such a dramatic effect. In particular, consider a scenario where the long run outcome only depends on the initial state value, if one never implements a new policy, i.e. when $\tau = \infty$. The consumer expects the utility of consumption to decrease in the second stage, and the parameters correspond to those in Section 3.1. Fig. 6 shows solution paths for different switching times and the initial state value $S_0 = 8.5$. A user would adapt consumption according to the time when the policy change is expected to occur; e.g., if the user expects the policy change to happen soon, drug consumption is reduced much more drastically than if the policy change is not expected to happen for a relatively long time. This relates of course to solution paths getting closer to the solution with $\tau = \infty$ as τ increases. For $\tau = \infty$ we find a Skiba point at $\bar{S} = 6.813$.

What is particularly interesting about the solution paths in Fig. 6 is the following: For certain switching times starting at an initial state value $\tilde{S} > S_0 > \bar{S}$, i.e. the consumer would approach the high consumption steady state if Stage 1 would last forever. However, one can merely by announcing a new policy achieve that a consumer reduces consumption so much that by a time $\omega \leq \tau$, i.e. still in the first stage, the consumption stock is so small that $S(t) \leq \bar{S}$ for all $t \geq \omega$. This means that if in fact the policy change never actually occurred, the consumer only realized this at some time $t > \omega$, the consumption stock would have already fallen so much because of this announcement that the optimal strategy would continue to decrease consumption and become abstinent on the long run. Effectively this also means that a policy maker does not have to implement a new policy

14

if the consumption stock has already fallen this much. $\tilde{S}$ denotes the degree of addiction above which just announcing a policy change and not implementing it would only lead to a short term change of the consumers behavior.

Figs. 6 and 7 show that if one wants to push a consumer to abstinence just by announcements (rumors, promises, threats), a policy maker has to be careful about the timing. While a user would decrease consumption the most if the policy change is announced to happen soon (i.e. $\tau < \tau_{\min}$), the level of addiction would not decrease fast enough, so when the user realizes that the new policy is not going to be implemented, he or she would start again to consume more. On the other hand if the announced time is too high (i.e. $\tau > \tau_{\max}$), then the optimal strategy of the drug user would be too close to what would be chosen if the policy change were never expected to occur. For certain initial state values and switching times, however, there are indifference points (i.e., the consumer has the choice between two different solution paths leading to the same steady state; see Grass et al., 2008) as depicted in the right panel of Fig. 6, and also in the left panel of Fig. 7, i.e. the points where $\tau_{\max} = \bar{\tau}$. Starting at such a point a consumer is indifferent between either decreasing consumption and directly approaching the second stage steady state or first increasing consumption and only decreasing it after some time. As already mentioned, the non-monotonicity of the solution path relates to the exogenous switching time.

The right panel of Fig. 7 shows that ω, i.e. the time when a consumer can notice that a policy change is not going to occur and still approach the no-consumption steady state, depends on the announced switching time τ and the initial state value. It has to hold that $\omega \leq \tau$ and, consequently, ω can neither be smaller than $\tau_{\min}$ nor larger than $\tau_{\max}$. If the announced switching time equals the minimum possible time to reach the Skiba point, it makes of course sense that ω must be equal τ. Any lower value for ω means that a consumer would notice that the new policy is not going to be implemented when his addiction level is still so high that when adapting his optimal strategy expecting the current policy last forever, he or she would still approach the steady state with a high level of addiction. Also, for certain initial state values it holds that, if the announced switching time equals the maximum time which makes the consumer sufficiently change his behavior that he ends up in the abstinence steady state, ω must also be equal to this time. This reason for this is that for these initial state values and switching times, the optimal path is the one where one would switch to the second stage exactly on the Skiba point. However, if the consumer is indifferent between two solution paths for $\tau_{\max}$, this is not the case. Then the consumer would only end up on the preferential side of the Skiba point, if he or she chooses the monotonous solution path (for any $\tau < \tau_{\max}$ he or she would choose a monotonous solution path), thus ω will be smaller than $\tau = \tau_{\max}$.

Note, however, that whether a consumer would adapt just because of an announcement, depends on the assumption that the user really believes what a policy maker announces. Also the assumption that, when the consumer realizes that the new policy is not going to happen, his or her utility does not change, e.g., due to disappointment or a general reevaluation of the actual situation, plays a big role. So this particular finding may be more of an analytical oddity than a practical recommendation; there is, after all, a reputational cost to crying wolf.

5 Conclusion

We showed that sometimes a policy change concerning an addictive good can have irreversible effects for some types of persons so that changing and then restoring the policy will not always return the level of use to its original equilibrium. This can apply to both liberalizations (e.g., legalization) and also to policies that increase sanctions (crackdowns).

We also showed that a given policy change can have diametrically opposed effects on future consumption for different users, even though their preferences might not much differ by very much. Hence, the aggregate effect of a policy can depend on the relative proportions of different types of people in the population.

We also showed that a policy maker has to be very careful not only about what new policy to implement, but also about when to implement it. While waiting might produce an outcome more desirable for the policy maker with respect to one person, it can lead to the opposite effect for another. So again, relative proportions of different types of people in the population matter.

We also verified that within this model a policy maker can influence the behavior of drug consumers today just by announcing a future policy shift. That in and of itself is not surprising; after all, that is almost the definition of a drug user having foresight. However, an interesting implication of this couple with the existence of threshold behavior is that policy makers could in principle permanently alter the consumption behavior of users merely by announcing (credibly) a policy change without ever implementing it. If while the policy is still being anticipated, the user's response to the expected change pushes his or her consumption stock across a Skiba threshold, then even if the user subsequently realizes that the policy will not in fact be implemented, the rational user's long run behavior will still forever be influenced by that past "misunderstanding". So in theory, if the policy change leads to a situation the policy maker favors, the policy maker might be able to obtain this outcome "for free", although, it is important that the announced time until the policy change is supposed to occur is neither too short or too long, since different expectations lead to different adaptions of the consumption pattern.

The results are also interesting from a methodological point of view. We saw that if we have a Skiba point in the first stage (i.e. for $\tau = \infty$), for certain initial state values and exogenously given switching times, indifference points leading to the same steady state can be found in the multi-stage version of the model. Further, if we have Skiba in the second stage (i.e. for $\tau = 0$), we also have Skiba points in the multi-stage model for certain initial values and switching times. The exogenously given switching time acts more or less as a second state variable in terms of allowing non-monotonic solution paths, and also Skiba and indifference points from which more than one optimal trajectory can depart while moving "in the same direction" with respect to the state.

Further, the approach of using multi-stage techniques to deal with the question of how a decision maker would react to a policy change if he or she can or cannot anticipate it, might also be relevant for a wide range of other applications, e.g., environmental regulations, terror control, etc.

The model presented here can be considered as a precursor to some kind of dynamic game with two players, namely the drug consumer and the policy maker. The simplest extension then would be to let the decision maker optimally decide when to implement a

new policy. We saw that there are certain scenarios when it pays to announce a new policy in advance. However, since politicians too often make it hard for anyone to consider them to be truthful, a drug consumer might not necessarily believe any kind of announcement and act accordingly. An additional state could be added representing the "credibility" of the politician, and the politician would face a meta-problem of deciding how optimally to spend down that asset by strategically reneging on past statements at the best possible time.

Acknowledgments

References

Becker, G. S., 1992. Habits, addictions, and traditions. Kyklos 45 (3), 327–346.

Becker, G. S., Murphy, K. M., 1988. A theory of rational addiction. Journal of Political Economy 96 (4), 675–700.

Braun, N., Vanini, P., 2003. On habits and addictions. Journal of Institutional and Theoretical Economics 159, 603–626.

Caulkins, J. P., Feichtinger, G., Grass, D., Hartl, R. F., Kort, P. M., Seidl, A., 2010. Skiba points in free end time problems: the option to sell the firm. In Submission.

Caulkins, J. P., Feichtinger, G., Hartl, R. F., Kort, P. M., Novak, A. J., Seidl, A., 2011. Multiple equilibria and Skiba points in a rational addiction model. Research Report, available at http://orcos.tuwien.ac.at/research/research_reports/.

Degenhardt, L., Reuter, P., Collins, L., Hall, W., 2005. Evaluating explanations of the australian 'heroin shortage'. Addiction 100 (4), 459–469.

Dockner, E. J., Feichtinger, G., 1993. Cyclical consumption patterns and rational addiction. The American Economic Review 83 (1), 256–263.

Gable, R. S., 2004. Comparison of acute lethal toxicity of commonly abused psychoactive substances. Addiction 99 (6), 686–696.

Gavrila, C., Feichtinger, G., Tragler, G., Hartl, R. F., Kort, P. M., May 2005. History-dependence in a rational addiction model. Mathematical Social Sciences 49 (3), 273–293.

Grass, D., Caulkins, J. P., Feichtinger, G., Tragler, G., Behrens, D. A., 2008. Optimal Control of Nonlinear Processes: With Applications in Drugs, Corruption and Terror. Springer, Heidelberg.

Hartl, R. F., 1987. A simple proof of the monotonicity of the state trajectories in autonomous control problems. Journal of Economic Theory 40 (1), 211–215.

Iannaccone, L. R., 1986. Addiction and satiation. Economic Letters 21 (1), 95–99.

Makris, M., 2001. Necessary conditions for infinite-horizon discounted two-stage optimal control problems. Journal of Economic Dynamics and Control 25 (12), 1935–1950.

Maltz, M. D., 1996. From poisson to the present: Applying operations research to problems of crime and justice. Journal of Quantitative Criminology 12 (1), 3–61.

Orphanides, A., Zervos, D., 1994. Optimal consumption dynamics with non-concave habit-forming utility. Economic Letters 44 (1–2), 67–72.

Orphanides, A., Zervos, D., 1995. Rational addiction with learning and regret. The Journal of Political Economy 103 (4), 739–758.

Orphanides, A., Zervos, D., 1998. Myopia and addictive behaviour. Economic Journal 108 (446), 75–91.

Ryder, H. E., Heal, G. M., 1973. Optimal growth with intertemporally dependent preferences. Review of Economic Studies 40, 1–33.

Stigler, G. J., Becker, G. S., 1977. De gustibus non est disputandum. American Economic Review 67 (2), 76–90.

Tomiyama, K., Rossana, R. J., 1989. Two-stage optimal control problems with an explicit switch point dependence: Optimality criteria and an example of delivery lags and investment. Journal of Economic Dynamics and Control 13 (3), 319–337.