TECHNISCHE UNIVERSITÄT WIEN
Vienna University of Technology

Operations Research and Control Systems

When Two Become One: Optimal Control of Interacting Drug Epidemics

I. Zeiler, J.P. Caulkins, G.Tragler

Research Report 2010-07

March, 2010

Operations Research and Control Systems
Institute of Mathematical Methods in Economics
Vienna University of Technology

Research Unit ORCOS
Argentinierstraße 8/E105-4,
1040 Vienna, Austria
E-mail: orcos@eos.tuwien.ac.at

When Two Become One: Optimal Control of Interacting Drug Epidemics

I. Zeiler[*,a,1], J.P. Caulkins[b], G.Tragler[a]

[a] *Department of Mathematical Methods in Economics, Vienna University of Technology, Argentinierstrasse 8/105-4, A-1040 Vienna, Austria*

[b] *Carnegie Mellon University Qatar Campus, Adelaide Campus, Heinz College's School of Public Policy & Management School of Information Systems & Management, 5000 Forbes Ave. Pittsburgh, PA 15213*

Abstract

In this paper we set up a model of interacting drug epidemics by taking the "direct-product" of variations on a one-dimensional optimal control model that exhibits a so-called DNSS or Skiba threshold. Such a threshold reflects in drug policy the well-established paradigm of "eradication vs. accommodation"; that is, depending on the initial state, stabilizing a drug market at either a low or a high level of use is optimal.

We investigate whether and how this "eradication vs. accommodation" paradigm extends to higher dimensional models and whether and how the presence of a second drug can alter the optimal policy description for the first drug.

The main results are, first, the presence of a second drug market can dramatically alter or even reverse the optimal policy prescription for the first drug market. Second, with interacting drug markets "eradication vs. accommodation" is no longer a binary choice. Rather, we observe a fascinating series of situations with multiple optimal steady states and complex structures of optimal solutions.

Key words: optimal control, indifference points, thresholds, multiple equilibria, DNSS points, illicit drug use

[*] Corresponding author

Email addresses: irmgard.zeiler@ait.ac.at (I. Zeiler), caulkins@cmu.edu (J.P. Caulkins), tragler@eos.tuwien.ac.at (G.Tragler)

[1] Present address: Austrian Institute of Technology, Giefinggasse 2, A-1210 Vienna, Austria, irmgard.zeiler@ait.ac.at, Ph. +43(0) 664 815 78 17, Fax +43(0) 50550-6439

1. Introduction

Illegal drugs impose enormous costs on societies throughout the world and present challenging problems for policy makers. Over the last decade optimal control theory has been harnessed to provide crucial insights concerning how policy ought to adapt over the course of a drug epidemic (see for example Tragler et al., 2001; Behrens et al., 2000, 2002; Bultmann et al., 2008a,b).

These models employ mechanisms from epidemiology to describe the spread of drug use in a population. Drug users tend to start using by having contact to already active users, thus getting "infected" by using individuals. This leads to a dynamic feedback loop familiar from conventional epidemics. Recognizing these effects the literature concludes that policy should react dynamically to an evolving drug market. This paper advances that literature by exploring the optimal control of multiple interacting drug epidemics.

A common feature in these mathematical models are "tipping points" that indicate a certain size of a drug problem, above which "eradication", in the sense of the stabilizing the drug problem at a low level of use is not optimal anymore. Once the epidemic has grown beyond this certain size, an "accommodation" strategy that delays the growth of the problem rather than reversing it is the policy of choice. In this paper, each drug by itself would exhibit such a tipping point, and the primary question of interest is how if at all this may be affected by the presence of another drug epidemic. Inasmuch as the general framework of multiple contagious interacting "bads" (as opposed to goods) occurs in diverse contexts (multiple diseases, multiple biological pests, etc.) the overall framework and approach may be of general interest. However, the functional forms we use are specific to drug policy modeling.

While the "eradicate or accommodate" paradigm is reasonably well established, to date it has been examined primarily within one dimensional models. However, a more realistic model would recognize that there are often multiple important dimensions of market size or activity.

For instance, there is not just one drug. A common situation in the U.S. and parts of Europe is that there are robust markets for both heroin and for a stimulant such as cocaine (often in the form of crack). Many dependent users are polydrug users, willing and able to switch from one substance to another in response to changes in availability and price.

Likewise, there may be two distinct regions that can have separately measurable levels of activity but which are nonetheless linked inasmuch as drug sellers and/or users can migrate from one market to another. One can also think of multiple adjacent countries. E.g., at various times there has been considerable friction between the Netherlands and its neighbors over

"drug tourism".

So it is natural to ask whether and how the "accommodate vs. eradicate" choice generalizes to multi-dimensional markets. The strategy we employ to approach this question is to take the "direct-product" of variations on one of the most fundamental and widely studied optimal control models of drug markets, that of Tragler et al. (2001). Zeiler et al. (to appear) present some mathematical results for the very special case of two identically parameterized drug epidemics. Here we consider the more realistic case of asymmetric epidemics, by starting with symmetry and then varying one epidemic in various ways.

Substantively the most important insight is simply that the presence of a second drug market can dramatically alter the policy prescriptions for the first drug market, and what is optimal to do can depend crucially on the nature of the linkage between the two markets. For example, if a drug acts as a gateway that facilitates escalation to a second drug, the optimal policy may change from "accommodation" in the isolated case to "eradication" when the combined system is considered. In other cases the presence of the second drug "tips" the solution toward accommodation.

Methodologically we observe a fascinating series of results that reflect some important developments in optimal control theory. In optimal control theory "tipping point behavior" is indicated by the existence of so-called DNSS or Skiba points (henceforth DNSS, named after Dechert and Nishimura, 1983; Sethi, 1977, 1979; Skiba, 1978). A DNSS point indicates a point at which a decision maker is indifferent between choosing one of two or more optimal strategies, because each strategy would generate the same (minimal) cost. So this phenomenon is related to the familiar tipping points in uncontrolled models, but differs by pertaining to models with an objective function. Due to the higher dimensional nature of our model, DNSS thresholds are curves instead of points, which allows a variety of additional insights to optimal policy prescriptions.

This paper is organized as follows. First, we introduce the mathematical model of two interacting drug epidemics and its formulation as an optimal control problem, and solve it by applying Pontryagin's Minimum Principle. Second, we present numerical results on the optimal control of drug epidemics of equal importance and of different size, of early stage drug scenarios, and of epidemics in distinct jurisdictions. We close with a discussion of the results and important conclusions for controlling such interacting systems.

2. A Model of Interacting Drug Epidemics

2.1. Underlying Dynamics

Drug use can be thought of as an epidemic process, because there is a positive feedback from the current number of users to the number of new initiates. The one-dimensional model we use is based on Tragler et al. (2001), who analyzed the the optimal dynamic allocation of treatment and enforcement in illicit drug control. Tragler et al. (2001) tracked the total number of users $A(t)$ over time, not differentiating between lower- and higher-frequency users as in Behrens et al. (1999, 2000, 2002). This leads to a relatively simple model that is easy to adapt. The dynamics of $A(t)$ is governed by three flows: initiation, $I(t)$, which was a power function of the current number of users $A(t)$ and of the retail price p, natural desistance, $\Omega(t)$, which is users who quit without being assisted, and desistance induced via drug control, $U(t)$.

We modified this model in two ways. First, we replaced Tragler et al.'s initiation function by the more familiar logistic function, which better reflects the reality that there exists a maximum number of users even when no controls are applied. Second, the control "law enforcement" is omitted for simplicity, which leads to a constant retail price. Bultmann et al. (2008a,b) and Grass et al. (2008) discussed these modifications at length.

The impact of control spending for "treatment", $u(t)$, is modeled as increasing the flow out of drug use. We choose the same functional form as in the model of drug treatment of Rydell et al. (1996). In fact, the control can be seen as representing any control driven outflow, but for ease of exposition we refer to it as "treatment". The impact of control spending is assumed to have diminishing returns ("cream skimming"). This can be thought of as reflecting some ability to target spending so that when funds are limited, they are allocated first to the most cost-effective programs. Natural desistance, $\Omega(t)$, is modeled as occurring at a constant per-capita rate.

2.2. Interaction Among Models

Several forms of interaction between the epidemics are possible. We call the interaction "migration"; it can be understood as either physical migration, e.g. from one city to another, or in a wider sense as a change of primary substance of abuse. The system is "closed" with respect to this interaction, i.e. when an individual migrates from one population to another the *total* number of drug users in the system is not (directly) affected by that flow.

Here, migration is modeled as a "push" mechanism. That is, the number of users migrating is independent of the number of users in the target population. The connecting flows are constant per capita rates, denoted by γ_A and γ_B. γ_A denotes the rate at which users leave population A by migrating into

B, and γ_B denotes the rate at which users leave population B by migrating into A. We exclude balanced polydrug use, so at one moment of time either A or B is the primary substance or location of abuse for every individual.

2.3. Mathematical Formulation

In optimization there has to be a metric which measures what is "best". We follow other optimal control models of illicit drug use and define the objective to be minimizing the weighted sum J of the costs of drug use and of drug control spending over a long (infinite) planning horizon.

Use is measured in quantity consumed given by the term $A(t)p_A(t)^{-\omega_A(t)}$, where p_A is the retail price per gram and ω_A is the absolute value of the short run price elasticity of demand. In the objective functional quantity consumed is weighted by the average social cost per gram consumed, κ_A and κ_B, respectively. Spending on drug control, $u_A(t)$ and $u_B(t)$, enter the functional linearly.

The cost functional is *separable* in use and control of each substance. That is, the model can handle drugs of different social costs, e.g. heroin and cannabis. With the separable control variables the decision maker is able to direct financial resources dynamically to each of the epidemics.

The costs are discounted at an annual rate r for the usual reasons.

Letting $A(t)$ and $B(t)$ denote the current number of users and $u_A(t)$ and $u_B(t)$ denote the control effort exerted on the indexed epidemic at time t, the model is given by

$$\min_{u_A(t),\, u_B(t)\geq 0} J = \int_0^\infty e^{-rt}(\kappa_A p_A^{-\omega_A} A(t) + \kappa_B p_B^{-\omega_B} B(t) + u_A(t) + u_B(t))\, dt, \tag{1}$$

subject to

$$\begin{aligned}
\dot{A}(t) &= k_A p_A^{-a_A} A(t)(m_A - A(t)) - c_A \left(\frac{u_A(t)}{A(t)}\right)^{z_A} A(t) \\
&\quad - \mu_A p_A^{b_A} A(t) - \gamma_A A(t) + \gamma_B B(t), \\
\dot{B}(t) &= k_B p_B^{-a_B} B(t)(m_B - B(t)) - c_B \left(\frac{u_B(t)}{B(t)}\right)^{z_B} B(t) \\
&\quad - \mu_B p_B^{b_B} B(t) + \gamma_A A(t) - \gamma_B B(t),
\end{aligned}$$

where k is a constant governing the rate of initiation, a is the elasticity of initiation with respect to price, m is the maximum number of users, c is the control efficiency proportionality constant, z is a parameter reflecting diminishing returns of treatment, μ is the baseline rate at which users quit without treatment, and b is the elasticity of desistance with respect to price.

The five terms in each state equation model initiation, control driven outflow, natural desistance and migration, respectively.

2.4. Model Solution

The optimal control problem (1) is solved by applying Pontryagin's Minimum Principle (cf. Feichtinger and Hartl, 1986; Grass et al., 2008; Leonard and Long, 1992), which provides necessary optimality conditions. These are used to transform problem (1) into a so-called *canonical system*, which here is a four-dimensional system of ordinary non-linear differential equations.

To derive the canonical system we consider the current value Hamiltonian $\mathcal{H} = \mathcal{H}(A, B, \lambda_A, \lambda_B, u_A, u_B)$ defined by

$$\mathcal{H} = \lambda_0(g_A(A, u_A, \rho_A) + g_B(B, u_B, \rho_B)) + \lambda_A \dot{A} + \lambda_B \dot{B}, \tag{2}$$

where λ_A and λ_B denote the costate variables in current value terms and $\lambda_0 \geq 0$ is a nonnegative constant multiplier associated with the integrand of the objective function

$$g_A(A, u_A, \rho_A) + g_B(B, u_B, \rho_B) = \kappa_A p_A^{-\omega_A} A + \kappa_B p_B^{-\omega_B} B + u_A + u_B. \tag{3}$$

The time argument t is omitted for brevity. If $\lambda_0 = 0$, then the Hamiltonian minimizing condition would yield at least one $u_i = \infty$, $i \in \{A, B\}$, because $\mathcal{H}$ is monotonically decreasing in u_i and there is no upper bound on the controls. That would maximize rather than minimize the objective function (1). Hence, we can conclude that $\lambda_0 \neq 0$ and may set $\lambda_0 = 1$ without loss of generality.

There must exist two non-positive multipliers η_A and η_B for the non-negativity constraints on the controls such that the partial derivatives of the Lagrangian function $\mathcal{L} = \mathcal{L}(A, B, \lambda_A, \lambda_B, u_A, u_B, \eta_A, \eta_B)$ with respect to the controls are equal to zero when evaluated along an optimal path. The Lagrangian function is defined by $\mathcal{L} = \mathcal{H} + \eta_A u_A + \eta_B u_B$, and the complementary slackness conditions, $\eta_i \leq 0 \wedge \eta_i \, u_i = 0$, $\forall \, i \in \{A, B\}$, imply that as long as the optimal controls are in the interior of the feasible region, i.e. are strictly positive, the Lagrangian function $\mathcal{L}$ reduces to the Hamiltonian $\mathcal{H}$ and the solutions of the constrained problem equals the unconstrained one.

So, if $\forall i \in \{A, B\} : u_i > 0$, the necessary optimality condition on the controls is to minimize the Hamiltonian subject to the constraints at any instant of time, that is $u_i^\star = \arg\min_{u_i} \mathcal{H}$. Using the convexity of the Hamiltonian with respect to the controls, which is proven by the positive definiteness of

the Hessian matrix,

$$\frac{\partial^2 \mathcal{H}}{\partial u_i \partial u_j} = \begin{pmatrix} -\frac{1}{A} c_A \left(\frac{u_A}{A}\right)^{z_A - 2} (z_A - 1) z_A \lambda_A & 0 \\ 0 & -\frac{1}{B} c_B \left(\frac{u_B}{B}\right)^{z_B - 2} (z_B - 1) z_B \lambda_B \end{pmatrix},$$

$i, j \in \{A, B\}$, we can derive $u_A^\star = A(c_A z_A \lambda_A)^{\frac{1}{1-z_A}}$ and $u_B^\star = B(c_B z_B \lambda_B)^{\frac{1}{1-z_B}}$ for the controls by setting $\mathcal{H}_{u_i} = 0$. If any $u_i = 0, i \in \{A, B\}$, the complementary slackness condition yields $\eta_i \leq 0$, where an explicit form for η_i can be derived from $\mathcal{L}_{u_i} = 0$.

Furthermore, the costate equations are given by $\dot{\lambda}_A = r\lambda_A - \mathcal{H}_A$ and $\dot{\lambda}_B = r\lambda_B - \mathcal{H}_B$. The limiting transversality conditions for the costates λ_A and λ_B are $\lim_{t\to\infty} \mathbf{e}^{-rt}\lambda_A(t) = 0$ and $\lim_{t\to\infty} \mathbf{e}^{-rt}\lambda_B(t) = 0$ and hold if states and costates approach a stable state. The state and costate equations form, after substituting for the controls, the canonical system.

Problem (1) yields long-run steady state solutions that are saddle-node equilibria. The candidate trajectories for optimal solutions are paths that converge into such a saddle point. These paths are calculated numerically using the method presented in Grass et al. (2008).

We use the parametrization of Bultmann et al. (2008a,b), which roughly describes the current US cocaine epidemic. With these values the underlying one-state model manifests a DNSS point. All scenarios analyzed are built by using round percentage changes to some of the baseline parameter values that preserve the DNSS structure in the unconnected isolated systems.

3. Results

3.1. Drugs of Equal Importance

We look first at the case of two drugs of roughly equal importance and for which users can flow in either direction, from drug A to drug B or vice versa. A typical application example would be one city or country dealing with two "hard" drugs, such as heroin and cocaine/crack. The connecting flows are assumed to be symmetrical, i.e. the rate at which users switch to the other drug is identical for drugs A and B. Furthermore, the state equations and cost functionals are identically parameterized.

As a foil, consider first what happens if there is no interaction between the two populations. Then we have merely the "direct product" of two one-state optimal control problems. The optimal control strategy depends then on the initial state of each epidemic only. In the phase portrait, Figure 1, this is represented by the two straight DNSS *lines* that separate the locally optimal steady states. Since each epidemic converges either to its high (H) or

low level of use (L), the two-dimensional model has four optimal outcomes, (L, L), (H, L), (L, H), and (H, H). If the system starts at a point on such a DNSS line the decision maker can pursue eradication or accommodation at equal costs for at least one epidemic. On the central intersection point of the DNSS lines there are four options: either eradication or accommodation of both drugs or one of two discordant strategies, where one epidemic is driven to its low level of use, while the other converges to the high-level of use.

Once even a modest amount of interaction is introduced this situation changes. Figure 2 shows the phase portrait for $\gamma_A = \gamma_B = 0.005$. That is, every year just one half of one percent of drug users switch their primary substance of abuse to the other drug. With connected systems every action taken on one epidemic (indirectly) influences the other, because the connecting flows change. The asymmetric steady states move from the boundary to the interior of the state space, and are now separated by five different DNSS curves. In the figure these DNSS curves are shown together with the DNSS lines of the zero interaction case (grey, dashed). The area between the solid DNSS curves and the dashed DNSS lines are a set of initial values for which the optimal outcome has changed because of the introduction of this very modest interaction. When this set is of significant size we know that the system reacts sensitively to changes of the respective parameters.

In Figure 2 it is obvious that interaction favors the "pure" eradication or accommodation strategies; the area in which it is optimal to maintain just one drug at a small level of use gets smaller. However, the effect is modest, because the systems are only very loosely connected.

A more relevant case is depicted in Figure 3, where $\gamma_A = \gamma_B = 0.01$, so the migration rate has about the same magnitude as the natural desistance. Here, the quality of the solution has fundamentally changed. The two asymmetric steady states are no longer optimal, i.e. if the system starts in such a state the decision maker's better choice is to use control to leave this steady state for the low-level of use. Consequently we have two instead of four optimal steady states and these are separated by a single DNSS threshold. Zeiler (2007) has shown that this structure holds for any $\gamma > 0.0051987$.

Since this critical value of γ is so small, one might say that interaction reduces the policy question to the familiar one of eradication versus accommodation, i.e. the two epidemics may be treated as one large drug problem. But this is not quite true, since the DNSS curve is *convex* in the states. So it is not just the total number of users in the system of drug use that determines the outcome, but the relation of the initial states of A and B. At the central DNSS point (which is a DNSS point for all identically parameterized systems with symmetric migration) the total number of users in the system is roughly 10.4 million. Because of the convexity, the total number of users

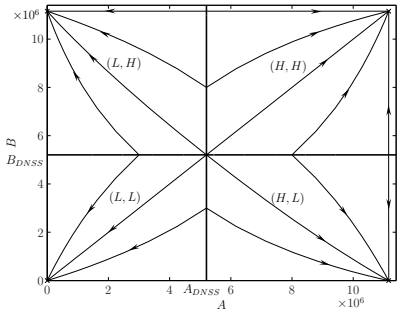

Figure 1: Identically parameterized epidemics without migration, $\gamma = 0$.

such that the decision maker is indifferent between eradicating both or accommodating both epidemics increases when moving away from this central DNSS point. In policy terms this means that eradication of both epidemics is indicated even though the total number of users exceeds the sum of the separate, single drug DNSS thresholds.

It may seem counterintuitive that eradication is easier when the initial conditions are asymmetric, but optimal control exploits the asymmetry by attacking first the market where progress is easiest.

To conclude when there is only a moderate level of interaction between two epidemics of equal importance choosing between "eradication" or "accommodation" is still the question to ask, but the answer depends now in part on the relative numbers of users of the two drugs, not just on the total number of users.

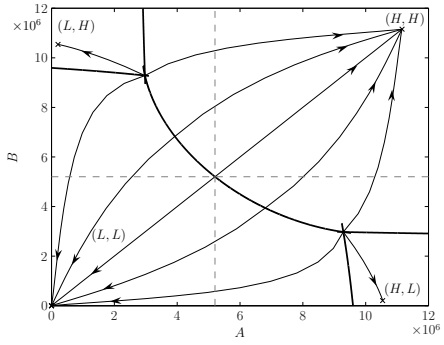

Figure 2: Identical epidemics with little migration, $\gamma = 0.005$.

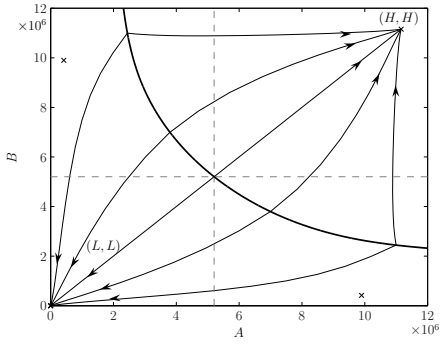

Figure 3: Identical epidemics with modest migration, $\gamma = 0.01$.

3.2. Interacting Epidemics of Different Size

The case discussed in the last section, where both drugs are of equal size, is important but in many countries use of one of the substances is more widespread. In the Americas stimulants are more widely abused, whereas in Europe the opiates market is much larger.

We model such a situation by doubling the maximum number of users of drug B, m_B, as compared to the base case. This lets B grow to a high steady state that is about twice the size of the base case, and it lets B grow *faster*. That is, drug B is more attractive than A in the sense that users are more easily initiated into use. Furthermore, we double the social costs incurred by B, so the large drug is also of greater relative importance.

Figure 4 shows a phase portrait of this case with symmetric migration at rate $\gamma_A = \gamma_B = 0.005$. The dashed lines depict the DNSS thresholds of the respective epidemics without interaction. Please note, that the DNSS threshold of the isolated large epidemic B is now about twice its size in the base case. So increasing the maximum number of users and increasing the social cost parameter leads to an exact scale up of the drug problem, where the point of indifference is measured relative to the base case and not in absolute numbers.

Figure 4 shows that the rate of migration necessary to reduce the problem to a binary choice between either eradication or accommodation of both drugs is even smaller than in the previous section. The value of $\gamma = 0.005$ used here is sufficient to reduce the number of optimal steady states to two. Furthermore, the single DNSS curve differs only slightly from the DNSS threshold of an isolated, large drug B. That is, the large drug B dominates the system to such an extent that the long run behavior of both epidemics depends almost only on the initial state of B.

So, in the more realistic setting of interacting epidemics that are of different size, what is optimal for the larger drug is generally optimal for the whole system. And this is true even when the rate of migration between the drugs is small. If a decision maker has the problem of a well established, large drug B, then stabilization of a newly emergent but less important drug at a low level of use can never be optimal, even if the upcoming drug grows more slowly. On the other hand if a drug is already well-established when a potentially more serious drug emerges, then it may be optimal to quickly eradicate the established minor drug to prevent it from fueling growth of the more problematic drug. "Minor" in this context is relative - half as much use at the high steady state and half as costly per user for a total problem size one quarter as great. Conceivably this could be interpreted to question whether the US perhaps should have worked harder to eradicate heroin before

cocaine/crack came along, although it is not clear that the policy instruments available were actually powerful enough to make that possible.

3.3. One-Directional Flows

Drug use "careers" often progress through a predictable sequence of stages of use. These stages can be different levels of use of one substance (e.g. casual use of heroin is followed by regular use that leads to a stage where treatment is needed, Caulkins et al. (2007)) or a "progression" from use of one substance to use of another. Here, we model such a progression in the use of different substances by letting users only migrate from drug A to drug B.

There are two important subcases depending on whether the social costs associated with use of the first drug are relatively minor or are of the same magnitude as the social costs of the second drug. The former case (when use of the first drug creates relatively few problems besides the risk of escalation) matches some people's understanding of the relationship between so-called soft drugs (notably cannabis) and hard drugs. It is well established that drug users often follow established patterns of escalating through a sequence of substances (e.g., first alcohol, then tobacco, marijuana, and on to whatever hard drug is common in that cultural context). There is enormous debate about whether the escalation sequence is causal, so we should be clear that our model of the relationship is causal. Halving use of drug A will halve escalation from A to B in this model. The second case of similar social costs might describe the interaction between stimulants and opiates (dependent stimulant users sometimes use opiates to take the edge off the long-term stimulant use and become frequent opiate users).

A further distinction regards initiation. In a stimulants vs. opiates case direct initiation to drug B is not uncommon, whereas in a cannabis to hard drugs escalation scenario initiation to B might occur primarily via earlier use of drug A.

In all discussed cases we set $\gamma_B = 0$, i.e. there is no flow from epidemic B to A. To stay in an area where the qualitative results are robust to changes of the interaction parameter in the identical case, we set $\gamma_A = 0.01$.

3.3.1. One-Way Interaction Can Reverse The Optimal Policy

Here, two equally costly drugs attract new users equally and the epidemics are equally parameterized, except for the one-way flow from A to B.

Figure 5 depicts this solution. The DNSS lines of the uncoupled systems are shown by dashed grey lines. There are three optimal long-run solutions: stabilization of both drugs at a low or high level of use, (L, L) and (H, H), and an eradication of the early drug while accommodating the late drug.

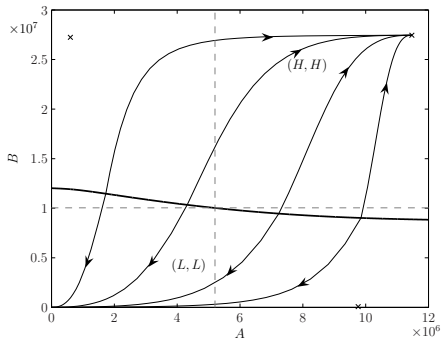

Figure 4: Epidemics of different size and equal importance with little, symmetric migration; $m_B = 2m_A$ and $\gamma_A = \gamma_B = 0.005$.

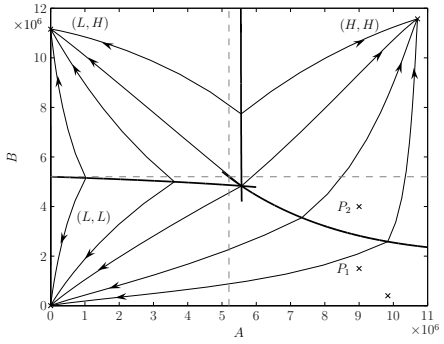

Figure 5: Identical epidemics with a one-directional flow from A to B, $\gamma_A = 0.01$ and $\gamma_B = 0$.

So a first finding here is that pursuing an eradication strategy on the late drug without controlling the first drug is not advisable. This can be explained easily; the one-directional flow acts as an additional outflow of drug A and an additional inflow to drug B. So in an asymmetric steady state with a larger A there is a constant stream of users that makes the control of B at a low level of use very difficult and hence too costly to be optimal.

Compared to the case of symmetrical flows, the DNSS curves that separate (L, H) from its adjacent equilibria are only slightly affected. The interesting DNSS curve is again the (L, L)-(H, H) DNSS curve that runs now only through the lower right "quadrant" (defined by the DNSS points of the uncoupled systems). If the systems were isolated, for all initial values in this quadrant the early drug A would converge to a high steady state and the late drug B to a low-level of use.

In contrast, suppose the asymmetric system starts at a point such as P_1 that lies below the (L, L)-(H, H) DNSS curve. If there were no progression from drug A to B the high costs of driving epidemic A down would outweigh the corresponding benefits and use of A would be allowed to remain high. But if drug A acts as a gateway that facilitates escalation to a second drug B there is a large additional area in which an eradication strategy for A should be applied to prevent the whole drug system from growing to a high rate of use for both drugs. So if there is a second drug B, the optimal policy for A changes from accommodation to eradication.

Furthermore, the DNSS curves do not intersect the A axes in the relevant area of the state space. So, if the second drug B is sufficiently small, the low level steady state of A should be approached for all initial values of A, even if escalation would have been optimal in the isolated system. Thinking

chronologically, if drug B simply did not exist before A stabilized, but then drug B suddenly became available, that change might imply a radical revision in what the optimal policy is for the established drug A.

If the system starts at a point such as P_2 that lies above the (L, L)-(H, H) DNSS curve, the optimal policy changes for the late drug from eradication to accommodation. The additional inflow from A leads in this area of the state space to a situation where the decision maker's better choice is to let both epidemics converge to (H, H) instead of (L, L). That is, the prior existence of an established "feeder drug" makes eradication less appealing for a newly emergent drug that would come later in the natural progression of use through a sequence of different substances.

So here we show a case where the familiar "accommodation" versus "eradication" choice gets turned on its head. In a significant area of the state space, a discordant strategy should be applied that leads to an asymmetric steady state, (L, H). For all other initial values the crucial factor that determines the final outcome is once more the combination of the initial states and the quality of the linkage of the two systems, and not so much the total number of users. So once the system is two-dimensional, the quality of the solution gains complexity.

3.3.2. A Lower Social Cost Favors Eradication

Next we consider the same set up, (identical drugs with a one-way flow from A to B) but now social cost per user of drug A is half that of drug B.

This modest variation changes the optimal policy thoroughly. Figure 6 shows the optimal solution of this case together with the DNSS curves of the previous case. As before there are three optimal steady states, the symmetric low and high levels of use and the asymmetric (L, H) state. However, the DNSS curves that separate the basins of attraction have fundamentally changed, greatly increasing the area in the state space for which a pure accommodation strategy is optimal.

An interesting phenomenon occurs at the sharp bend of the (L, L)-(H, H) DNSS curve. Here the DNSS curve is intersected by an indifference curve that separates two different ways of getting to the same long-run steady state. Such a phenomenon can only occur in optimal control models of higher dimension and is hence of great interest.

Here the occurrence of the indifference curve is directly related to the differences in the social costs parameter. When starting on this indifference curve the decision maker has two options to drive the system to the high level of use. S/he can opt for pt_1, a path where the number of costly users of drug B is larger or for pt_2, where more control is exerted to keep B small as long as possible. However, the indifference curve is very small. In a later

example we show such an indifference curve that is of significant size (see Section 3.4.1).

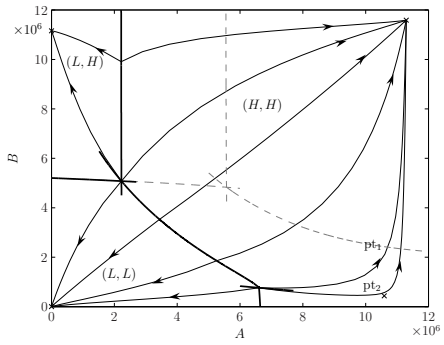

Figure 6: Early drug A is less costly than late drug B, $\kappa_A = 0.5\kappa_B$, $\gamma_A = 0.01$ and $\gamma_B = 0$.

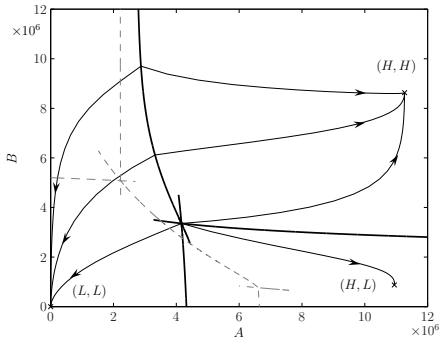

Figure 7: Most initiation happens via the early and less costly drug A, $k_B = 0.7k_A$, $\kappa_A = 0.5\kappa_B$, $\gamma_A = 0.01$ and $\gamma_B = 0$.

3.3.3. Lower Social Costs and Sequential Initiation Bring Back Eradication

The previous scenario might loosely describe escalation from cannabis to heroin use except for the fact that initiation directly into heroin was just as common as cannabis initiation, whereas in reality few people start using heroin without first using cannabis. We take this into account by reducing the initiation constant of the late drug B by 30%. This can be interpreted as assuming that 70% of heroin initiation is independent of cannabis use, but 30% is causally connected in the sense that halving cannabis use would halve that 30% of heroin initiation.

This has two consequences in the system dynamics. First, the total number of users of drug B is less "dangerous" in the sense that the number of new users per current user is smaller than in the base case. Second, the inflow of drug A has a relatively larger weight in the dynamics of drug B, so a reduction of this inflow has a greater impact on the system dynamics of B.

Figure 7 shows the new phase portrait. Here, the steady state (H, L) is an optimal steady state. This contrasts with the previous results, where eradication of the early drug was only optimal in combination with eradication of the late drug. But with the extra weight on the inflow from A to B, stabilization of the late drug at a low level is optimal for some initial values, even though the early drug grows to its high level of use.

The second striking result is the loss of optimality of the (L, H) steady state. So far stabilizing the early drug at a low level of use while allowing escalation of the late drug was a candidate for an optimal policy. But changing the initiation constant led to a situation where the early drug gained so

much weight in the dynamics that eradication of the late drug always pays off whenever the early drug is small enough. That is a striking structural change to arise from changing one parameter by only 30%.

With a further decrease of the initiation constant the picture changes once more. When the late drug attracts new users directly to only a very small extent, then the initial state of the early drug dominates the situation even though the early drug has substantially smaller social costs. Hence, if the early drug A is small enough, eradication is the policy of choice; if it has grown beyond a certain threshold, both epidemics should converge to a high level of use, regardless of the initial value of the late drug B. So, if the progression to a late drug happens almost exclusively via another drug, the early drug should determine the policy choice between eradication and accommodation, Inasmuch cannabis use is already well-established in many regions around the world, this could be interpreted as an argument against trying to achieve a low drug use equilibrium.

Concluding we can say, that in a system where users "progress" from one drug to the next, the nature of this progression is crucial. When the more costly substance attracts migrating users mainly from an "earlier" substance (i.e. marijuana, ecstasy or LSD), this earlier substance governs optimal policy (Figure 7). However, if the diffusion process of initiation of new users through existing users is of relatively high importance for the later drug (Figure 6), then the presence of the feeder drug may alter the optimal strategy for the late, costly drug from "accommodation" to eradication.

3.4. Distinct Jurisdictions

The extreme version of asymmetric costs is when the two states represent drug using populations in adjacent but distinct jurisdictions. That is, one drug, here B, does not directly impose any societal costs on the decision maker because those users live in another jurisdiction.

We distinguish between two important cases. In case one, the policy maker cannot or does not want to influence policy making in the other country. The other country's epidemic is assumed to grow uncontrolled from the decision maker's perspective.

In case two the decision maker in country A is willing to fund interventions in the second country, if this intervention would bring benefits to his or her own country, because the reduced number of users in the adjacent country leads to smaller migration of users from B to A. This could occur with a decision maker from a large and affluent country that funds treatment in a smaller, poorer neighbor.

3.4.1. Case One: An Uncontrollable Neighboring Drug Epidemic Changes the Question from 'Where To Go' To 'How To Get There'

In this scenario drug B does not impose costs in the objective functional, so $\kappa_B = 0$. Also drug B cannot be controlled by the decision maker, which is modeled by setting the efficiency of control to zero, so $c_B = 0$.

As it is assumed that B grows uncontrolled to a high level of use, population B provides a constant inflow to A. Hence, an overall low level of use cannot be reached from any initial value where either A or B is positive.

In the phase portrait, Figure 8, this manifests in a single, stable steady state, where both drugs are at a high level of use, although the number of users in A is slightly smaller than in B because control is exerted on A in its high steady state.

This is not surprising. What is more striking is the existence of an indifference curve, and one of significant length. If the system starts on this indifference curve, for example on the initial value of the trajectories pt_1 and pt_2, the decision maker has two options. Moving to the left of the indifference curve, on pt_1 control is used to keep the number of costly users A low for as long as possible, i.e. until the second epidemic has grown so close to its steady state value that the constant stream of new users makes further control inefficient. Moving to the right of the indifference curve, on pt_2, the system converges more or less directly to the high level of use.

In a slightly different interpretation, we can say that when control begins before epidemic A has grown beyond the level indicated by the indifference curve, the optimal policy is to fight hard in order to maintain the drug problem at a small level for as long as possible. If drug A has grown beyond this curve, an 'accommodation' strategy is optimal, while the epidemics converge directly to the overall high level of use.

Figure 9 shows the remarkable differences in the optimal policies. It shows how the optimal controls along pt_1 and pt_2, $u_A^{pt_1}$ and $u_A^{pt_2}$, depend on the state of A. Both start at $A(0)$, the origin of the discussed paths. The control curve to the left of the vertical line corresponds to the part of trajectory pt_1 that lies to the left of its initial value in Figure 8. Here control is used to such an extent that it causes the number of users in A to drop quickly to a minimum level. However, eventually drug B's population grows to the point of being a significant source of inflow to drug A. The effects on A's population are briefly resisted, but by the time the number of users grows back to its initial level, the tide of immigration is overwhelming and even less control is used than would initially have been used when following trajectory pt_2. So point $A(0)$ indicates a point of indifference between two substantially different optimal policies.

This sheds new light on the form in which tipping points and thresholds for different optimal policies can occur. Here we have an example of a set initial values that qualify as 'tipping points' even though the final outcome is uniquely given. Depending on the side of the indifference curve the system starts in, the optimal policy 'tips' either towards an aggressive control strategy (on the left side) or a moderate strategy (the right side). So, the indifference curve separates two optimal policies of very different quality, but both trajectories lead to the same place. This shows that the question to ask in higher dimensional systems is not only about the best final outcome, but also about the best way to get there.

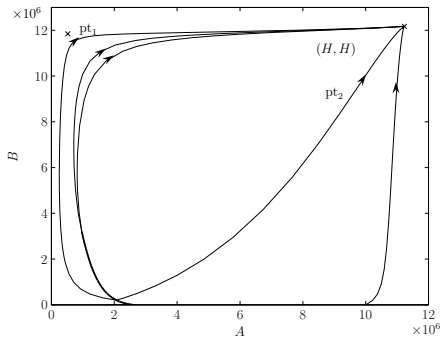

Figure 8: Similar drugs in different regions where only drug A is controllable and imposes costs to the decision maker, $\kappa_2 = 0$, $c_2 = 0$ and $\gamma = 0.01$.

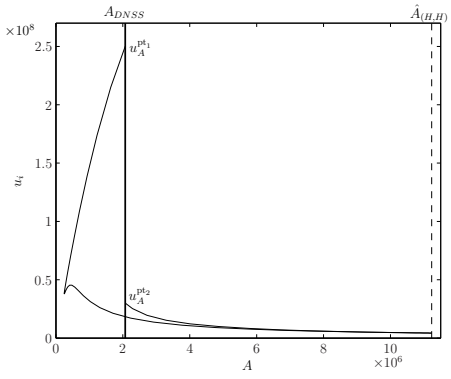

Figure 9: Optimal controls along two paths starting on the indifference curve. Only drug A is controllable and only it imposes costs $\kappa_2 = 0$, $c_2 = 0$ and $\gamma = 0.01$.

3.4.2. Case Two: Charity towards a Drug Using Neighbor May Be Self-Serving

A decision maker responsible for region A might be willing to fund drug control in an adjacent country, B, even if drug use in region B does not (directly) impose costs on the decision maker, since controlling population B reduces the number of users that migrate from B to A.

This scenario is reflected by having both control efficiencies, c_A and c_B, at baseline levels but by setting the social cost parameter of B, κ_B, to zero.

Figure 10 shows a phase portrait of this situation together with a dashed line marking the DNSS curve from Figure 3 which was identically parameterized except that $\kappa_2 = \kappa_1$ rather than $\kappa_2 = 0$. There are again two different optimal strategies, eradication of both epidemics and accommodation of both epidemics. However, the area in which eradication is optimal is very small when compared to Figure 3, because the direct benefits of lower use of drug

B are externalities not realized by the decision maker's objective functional. From a transcendent social welfare perspective, not enough is invested in control when the beneficial effects on one subpopulation are not valued by the decision maker.

More concretely for decision maker A eradicating drug B, while A is large, does not bring any benefits for the decision maker. Conversely, if B was large, it would provide a continuous stream of users migrating into A, the control of which is not optimal. However, for smaller levels of interaction a discordant strategy is optimal.

The (L, L) and (H, H) states are separated by a single DNSS curve that is intersected by an (H, H)-(H, H) indifference curve. As in the last section, this indifference curve separates two different strategies that both lead to the same steady state, but on distinct paths. However, this curve is relatively short, so its impact on the structure of the optimal solution is limited.

We concentrate here instead on solutions with a "classical" DNSS behavior. Figure 11 shows how the controls depend on state A, when the system starts at such a DNSS point, P. When eradication is pursued (to the left of the vertical line) the controls start at a very high level, before decreasing as eradication process. To the right side, the system converges to the overall high level of use (H, H). Consequently, both controls are at a lower level. Note that the dashed line representing control u_B is positive throughout the course of the epidemic; it is always in A's self-interest to fund at least a little drug control in Region B if cream skimming lets the first dollars be directed toward very high value targets.

So, the important policy implication from this case is that investing in the control of a drug problem in a distinct jurisdiction can be beneficial, even if this drug problem does not directly impose costs on the decision maker.

4. Discussion and Conclusions

For more than a decade optimal control theory has been applied to epidemiological models of illicit drug use. One of the most important resulting ideas is the paradigm of multiple optimal steady states with low and high levels of use leading to a strategic choice between "eradication" versus "accommodation", which means driving drug use down toward some low level steady state or allowing it to grow to a high level steady state. However, so far only single populations have been considered.

This paper extends the analysis from a single drug market to two interacting drug markets, and in the process generates at least three important insights. First, modeling single drug markets may be a potentially hazardous simplification, since the presence of a second drug market can dramatically

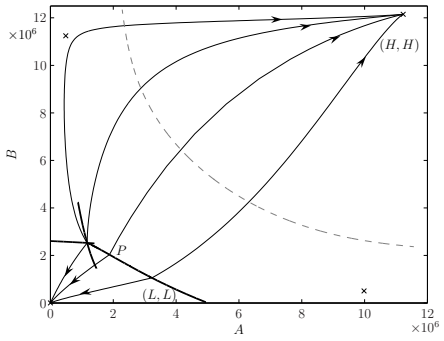

Figure 10: Similar drugs in different regions. Both drugs can be controlled, but only drug A imposes costs, so $\kappa_2 = 0$, $c_2 = c_1$ and $\gamma = 0.01$.

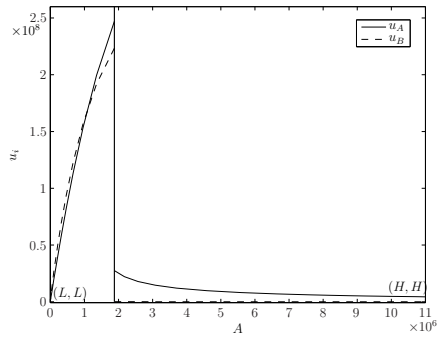

Figure 11: Optimal controls for similar drugs in different regions. Both drugs can be controlled, but only drug A imposes costs, so $\kappa_2 = 0$, $c_2 = c_1$ and $\gamma = 0.01$.

alter the policy prescriptions. In Section 3.1, for example, we saw that the introduction of a modest amount of interaction can dramatically change the quality of the optimal solution. So, the optimal strategy may depend on the relative sizes of the initial levels of drug use, not just on the individual initial states or their sum. Section 3.3.1 offered another example that underlines this point is presented in which the presence of a second drug market can altogether reverse the optimal policy prescription for the first drug market. However, Section 3.2 shows that a simplification to a one-dimensional system may be justifiable when drugs whose epidemics are of different importance and size are considered. Then the joint optimal strategy reduces essentially to what is optimal for the more important epidemic.

A second insight is that "eradication" versus "accommodation" is not necessarily a binary choice when there are multiple populations. The optimal policy for a system of drug markets may be a mix of strategies, the combination of which depends on the nature and the initial states of the drug markets involved. In all three early stage drug scenarios discussed in Section 3.3, we saw that there is a significant area in the state space where a discordant strategy is advisable meaning, one drug is driven to a low level steady state while the other grows toward its high level steady state. In the example presented in Section 3.4.1 we saw that in two-dimensional models the relevant question to ask is not only "Where are we going?", but also "How should we get there?". Indifference and, hence, threshold behavior does not only occur in the form of two optimal outcomes but as well in the form of two different optimal *paths* that lead to one certain long-run solution.

Third, we have shown that a parallel population, that is not costly itself

but that is connected to the population of concern, can have a tremendous impact on the optimal solution (Section 3.4). So decision makers should be advised to consider controlling not only drug use within their borders but also funding interventions for adjacent populations.

From a more general perspective, we presented a case study of taking the "direct product" of two one-dimensional optimal control models (with DNSS threshold behavior).

In the last years numerical methods and computational power have advanced to the point that two-dimensional optimal control problems with higher dimensional DNSS points can be solved in reasonable time and for a wide range of cases. This enables researchers to analyze and perform sensitivity analysis on more complex models. This opens a whole variety of possible, higher dimensional applications, built by generalizing one-dimensional systems to two-dimensional ones.

Such *parallel* systems are of relevance in many application domains. Further research could include other epidemiological models, such as the optimal control of multiple, interacting diseases (e.g. HIV and HCV) or infectious diseases in different populations (e.g. HIV/AIDS in adjacent countries). Furthermore, one might be interested in generalizing from one to two dimensions some famous examples in the economics literature, like the Ramsey Growth Model (Skiba, 1978) or a version of the contagion model for optimal advertising (Sethi, 1979).

Several insights gained here can be carried over to the analysis of the "direct product" of optimal control models in general (cf. Zeiler et al., to appear). Most notable is the realization that a little interaction can matter a lot. In most of the named applications, state dependence and tipping points have occurred in one-dimensional settings. These phenomena are common and important, so a deeper understanding of both the technical aspects and the policy application of them are of great interest and should be analyzed in models of different applications and different complexity.

Acknowledgements

This work was supported by the Austrian Science Fund (FWF) under contract No. P18527-G14, the Qatar National Research Fund, and by a Robert Wood Johnson Foundation Investigator Award in Health Policy Research. The views expressed imply no endorsement by the Robert Wood Johnson Foundation.

References

Behrens, D. A., Caulkins, J. P., Tragler, G., Feichtinger, G., 2000. Optimal control of drug epidemics: prevent and treat – but not at the same time? Management Science 46, 333–347.

Behrens, D. A., Caulkins, J. P., Tragler, G., Feichtinger, G., 2002. Why present-oriented societies undergo cycles of drug epidemics. Journal of Economic Dynamics and Control 26, 919–936.

Behrens, D. A., Caulkins, J. P., Tragler, G., Feichtinger, G., Haunschmied, J. L., 1999. A dynamic model of drug initiation: implications for treatment and drug control. Mathematical Biosciences 159, 1–20.

Bultmann, R., Caulkins, J. P., Feichtinger, G., Tragler, G., 2008a. How should drug policy respond to market disruptions? Contemporary Drug Problems 35 (2&3), 371–395.

Bultmann, R., Caulkins, J. P., Feichtinger, G., Tragler, G., 2008b. Modeling supply shocks in optimal control models of illicit drug consumption. In: Lecture Notes in Computer Sciences (LNCS). Vol. 4818. Springer, Heidelberg, pp. 285–292.

Caulkins, J. P., Dietze, P., Ritter, A., 2007. Dynamic compartmental model of trends in Australian drug use. Health Care Management Science 10 (2), 151–162.

Dechert, W. D., Nishimura, K., 1983. A complete characterization of optimal growth paths in an aggregated model with a non-concave production function. Journal of Economic Theory 31, 332–354.

Feichtinger, G., Hartl, R. F., 1986. Optimale Kontrolle ökonomischer Prozesse - Anwendungen des Maximumsprinzips in den Wirtschaftswissenschaften. Walter de Gruyter, Berlin.

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.

Leonard, D., Long, N. V., 1992. Optimal Control Theory and Static Optimization in Economics. Cambridge University Press, Cambridge.

Rydell, C. P., Caulkins, J. P., Everingham, S. S., 1996. Enforcement or Treatment: Modeling the Relative Efficacy of Alternatives for Controlling Cocaine. Operations Research 44 (6), 687–695.

Sethi, S. P., 1977. Nearest feasible paths in optimal control problems: Theory, examples, and counterexamples. Journal of Optimization Theory and Applications 23 (4), 563–579.

Sethi, S. P., 1979. Optimal advertising policy with the contagion model. Journal of Optimization Theory and Applications 29 (4), 615–627.

Skiba, A. K., 1978. Optimal growth with a convex-concave production function. Econometrica 46, 527 – 539.

Tragler, G., Caulkins, J. P., Feichtinger, G., 2001. Optimal dynamic allocation of treatment and enforcement in illicit drug control. Operations Research 49 (3), 352–362.

Zeiler, I., 2007. Optimal dynamic control with dnss curves: multiple equilibria in epidemic models of hiv/aids and illicit drug use. Ph.d. thesis, Vienna University of Technology, Faculty of Mathematics and Geoinformation, Vienna, Austria.

Zeiler, I., Caulkins, J. P., Tragler, G., to appear. Optimal control of interacting systems with DNSS property: The case of illicit drug use.