<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *SpringSim '17 Spring Simulation Multi-Conference Virginia Beach, VA, USA – April 23 - 26, 2017*.

Citation for the original published paper:

Okhravi, C., McKeever, S., Kronlid, C., Baraldi, E., Lindahl, O. et al. (2017)
Simulating Market-Oriented Policy Interventions for Stimulating Antibiotics
Development

In: *ANSS '17 Proceedings of the 50th Annual Simulation Symposium*, 2 ACM Digital
Library

<https://doi.org/10.22360/springsim.2017.anss.008>

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-372624>

SIMULATING MARKET-ORIENTED POLICY INTERVENTIONS FOR STIMULATING ANTIBIOTICS DEVELOPMENT

Christopher Okhravi
Steve McKeever

Carl Kronlid
Enrico Baraldi

Department of Informatics and Media
Uppsala University
Kyrkogårdsgatan 10, Uppsala, Sweden
{christopher.okhravi, steve.mckeever}@im.uu.se

Department of Engineering Sciences
Uppsala University
Lägerhyddsvägen 1, Uppsala, Sweden
{carl.kronlid, enrico.baraldi}@angstrom.uu.se

Olof Lindahl
Francesco Ciabusch

Department of Business Studies
Uppsala University
Kyrkogårdsgatan 10, Uppsala, Sweden
{olof.lindahl, francesco.ciabusch}@fek.uu.se

ABSTRACT

The relative efficacy of intervention policies, aimed at stimulating development of antibiotics, can be estimated using Agent Based simulation. We propose that antibiotics development can be modeled as Markov Chains with time and cash loaded transitions, and that many intervention policies can be modeled as alterations to the stochastic distributions of said Markov Chains. Through the combination of these two models, Agent Based simulation can be used to estimate the relationship between interventions and the Expected Net Present Value of products. We apply this modeling to an intervention policy proposed by the EU-initiative DRIVE-AB, targeting the urgent need for antibiotics research and development due to increasing resistance. We focus on variants fully delinking profit from volume sales, and show that (1) implementation variations lead to differences in outcomes, and that (2) they exhibit diminishing returns.

Keywords: antibiotics, policy interventions, market entry rewards, decision support

1 INTRODUCTION

The markets for new antibiotics are not sufficiently attractive for many pharmaceutical developers. Monetary incentives are needed to secure a healthy pipeline of new antibiotics (Harbarth et al. 2015). This paper proposes that the relative efficacy of different designs of an intervention policy commonly referred to as a Market-Entry Reward (MER) (Rex and Outtersson 2016) aimed at stimulating pharmaceutical Research & Development (R&D), can be estimated using Agent Based Modeling (ABM) where boundedly rational agents make investment decisions based on a combination of public and private information which is altered by said policy interventions. While the simulation results could be reached using Monte Carlo methods we lay the ground work for incorporating more complex agent behaviors and policy interventions in the Agent

Based paradigm. Thus, shifting focus to the modeling of interactions between antibiotic developers and away from macro-level assumptions that eliminate heterogeneity.

The increasing prevalence of antibiotic resistance is eroding the efficacy of the currently available antibiotics (Laxminarayan et al. 2013). The problem is severe as antibiotics are the backbone of modern medicine and a necessary prerequisite for treating medical conditions ranging from cancer to broken bones to pneumonia (Harbarth et al. 2015). Unfortunately, there is a lack of R&D aimed at discovery and commercialization of new antibiotics to replace the old ones facing resistant bacteria (Rex and Outtersen 2016).

In attempts to mitigate this catastrophic combination of increasing resistance and decreasing antibiotics R&D, efforts (hereon referred to as interventions) are underway to incentivize pharmaceutical firms to either increase their efforts in, or re-commit to, antibiotics R&D (see e.g. Payne et al. 2015). However, in order to gauge how different designs of an intervention might result in different numbers and types of antibiotics, the impact of interventions on pharmaceutical firms' propensity to engage must be understood. The lack of antibiotics R&D stem from pharmaceutical firms either (1) not engaging in discovery resulting in no new molecules entering the pipeline, and/or (2) prematurely discontinuing molecules already in the pipeline. Thus, to fully model the efficacy of any given intervention we need to establish (a) what brings a new molecule into the pipeline, and (b) when and why a molecule in the pipeline is discontinued. This work is concerned with the second point.

The paper is structured as follows. We first describe what a MER is and why optimization suggests the need for simulation. In Section 3 we formally capture the decision process of pharmaceutical firms, and describe our implementation of this process. In Section 4 we outline the input used in preliminary experiments. In Section 5 how we model policy interventions, and in Section 6 how we argue validity in our model. Finally in Section 7 we report exemplary results that have been observed in preliminary runs. Subsequent sections conclude and suggest avenues for future research.

2 MARKET ENTRY REWARDS

Delinkage, i.e. the detaching of revenues attained by a pharmaceuticals firm from the volume of sales for any particular drug, can be implemented through a Market Entry Reward (Rex and Outtersen 2016). A MER in the pharmaceutical context is intended to increase the willingness of a firm to push a project forward by increasing the size and certainty of the positive cash flows expected at the time of market entry. The revenue of a pharmaceutical product is expected to cover the costs incurred during R&D. Depending on a developer's time preference and profit requirements and due to the low probability of success, for a project to receive funding by economically rational investors, the area under the cash flow curve above zero may have to be significantly larger than the area below. A hypothetical irrational investment is depicted in Figure 1. A MER may mitigate this by promising a single or series of payments at market entry (Payne et al. 2015).

Figure 1: Illustrative depictions of hypothetical Market Entry Rewards.

From a business perspective, a MER reduces commercial risk (i.e. that a product will not be attractive for buyers) and, if promised before results are delivered, the competitive risk (i.e. that someone else commercializes a similar product before). From a public perspective, a MER has other benefits such as reducing the

incentive to aggressively market a drug in order to recoup costs, which, in the case of antibiotics, may help dampen the spread of resistance (Outterson 2010).

Depending on the size of the MER and who the benefactor is it may not be feasible to deliver the reward as a “lump-sum” directly at market entry (Rex and Outterson 2016). Hypothetically, while an upfront payment of x may not be feasible, paying x/y annually for y years may very well be. This installment of a payment will hereon be referred to as staging.

A MER may either completely separate revenue from free market sales (hereon referred to as fully delinked) or allows sales to complement the MER (hereon referred to as partially delinked). In the former the MER is granted *instead of* free market sales while in the latter *in addition to*.

The above points suggest that a MER can be thought of as a point in a three dimensional Cartesian coordinate system. Any MER must have a (1) size (i.e. how much is paid), (2) staging (i.e. over what period of time the payments are staged), and (3) delinkage level (i.e. what fraction of the originally expected market sales the beneficiary is allowed to realize). These characteristics captures the key aspects of Rex and Outterson (2016) as well as some more recent proposals stemming from the EU-initiative DRIVE-AB. The high number of permutations motivates the need for computational analysis to optimize the design of a MER.

3 MODELING ANTIBIOTICS R&D

The life cycle of a pharmaceutical product consists of the three major stages Discovery, Development and Launch (Blau et al. 2004). We refer to these as episodes. In Discovery, molecules that have effects on the target are identified. Permutations are experimented with and tested for toxicity. If no “worrisome toxic effects are observed” and some permutation suggests a successful drug, then the molecule is promoted to a “lead”. In Development the molecule is tested on healthy volunteers, then on patients with the disease, and finally in large-scale clinical studies conducted in concert with the relevant regulatory authority, e.g. FDA in the U.S. and EMA in Europe. The intent is to determine whether the molecule exhibits unacceptable side-effects. Finally a commercial plant is designed and constructed. If the drug is approved, it may proceed to Launch and be commercially sold. Pharmacovigilance monitors the product after launch, in search of unforeseen adverse effects (WHO and Collaborating Centre for International Drug Monitoring 2002).

The development of the new drug may be canceled due to technical failure (e.g. unwanted side effects, marginal efficacy, toxicity, etc.), and we will refer to this as failure. Furthermore, the development may be canceled due to poor financial prospects (e.g. in-house competition, low expected sales, etc.), and we will refer to this as termination. See Figure 2.

Figure 2: Stages and episodes of pharmaceutical R&D.

Each of the three stages can be divided into several detailed phases which we will refer to as stages (to avoid confusion in numbering as the first clinical phase often is referred to as phase 1). Innovation processes are often more iterative than commonly depicted (Van de Ven 1999), so how to splice the pharmaceutical R&D process into a series of discrete episodes is a matter of interpretation. Following Abbott and Vernon (2007), we divide the life cycle into six stages, but also introduce an explicit stage to represent sales. In Figure 2, stages are depicted (as solid boxes) inside their corresponding episodes (dashed boxes).

3.1 Modeling Go/No-Go Decisions

Due to the uncertainty of future financial flows inherent to pharmaceutical development, ENPV is commonly employed to evaluate investment opportunities (Svennebring and Wikberg 2013). ENPV (a.k.a. rNPV), is a risk-adjusted calculation of Net Present Value (NPV), which in turn is a way of discounting future money to its present value when evaluating an investment spanning over a long period of time.

NPV is calculated as in Equation 1 where t is a discrete time period, N is the number of time periods of the investment, C_t the (out-of-pocket) cash flow (revenue or cost) at period t , and i the discount rate of the evaluator. The discount rate in NPV reflects the firm's opportunity cost and the perceived risk of the investment (Atrill, McLaney, and Harvey 2014). ENPV multiplies each cash flow used in an NPV calculation by the probability that it occurs, which according to the multiplication rule of Bayesian probability equates to the expected value (in terms of NPV) of the investment. ENPV is calculated as in Equation 2 where $P_{0 \rightarrow \omega}$ is the probability of reaching the final cash flow from the point of investment, and $P_{t \rightarrow \omega}$ the probability of reaching the final cash flow from time t (Stewart 2002).

$$\text{NPV}_N^i = \sum_{t=0}^N \frac{C_t}{(1+i)^t} \quad (1) \quad \text{ENPV}_N^i = \sum_{t=0}^N \frac{C_t P_{0 \rightarrow \omega}}{(1+i)^t P_{t \rightarrow \omega}} \quad (2)$$

Abbott and Vernon (2007) argued that ENPV is firmly grounded in microeconomic theory and used it as a basis to simulate the effects on pharmaceutical R&D investments as a result of reduced prices. We build upon this work by (1) moving from Monte Carlo modeling to ABM, (2) introducing a framework that enables modeling of more interventions (e.g. MER) than reduced prices, and (3) furthering agent heterogeneity. We model boundedly rational, profit-maximizing firms (agents) that make initial project investment decisions based on ENPV, and then proceed by attempting to develop those projects. Interventions (further explained in Section 5) alter the basis upon which ENPV is calculated and may consequently affect the agents investment decisions, and in turn the resulting number of projects (antibiotics) completed. The ENPV of a rational investment increases as the project passes milestones (Sertkaya et al. 2014), unless the conditions under which the agent operates change. As we have yet to incorporate such changes, firms in our model rationally only make a single relevant investment decision (the first) and then stick to their choice.

Sertkaya et al. (2014) and Sertkaya et al. (2016) modeled agents investing in projects based on indication data, while Abbott and Vernon (2007) used industry averages. Consequently they are modeling hypothetical firms developing hypothetical drugs, and have thus modeled firms that calculate ENPV before entering Pre-clinical. As we are instead modeling hypothetical firms developing the antibiotics that are in the pipeline today, we let agents calculate ENPV of their projects regardless of where they currently are in the development life cycle. While our results suggest the effects of an intervention on antibiotics currently in the pipeline, it lays the foundation for simulating the effects of interventions on future hypothetical antibiotics.

A rational agent employing ENPV will only invest in a project if $\text{ENPV}_N^i \geq 0$. However, as some firms require ENPV to be greater than a threshold (Sertkaya et al. 2014 used M\$100) they require profit beyond that accounted for by their discount rate and thus above 0. For this reason we endow firms with a threshold parameter (call it A_ω) and let rational agents in our model only invest in a project if $\text{ENPV}_N^i \geq A_\omega$. As discount rates may vary between firms we also endow agents with a discount parameter (call it A_i) and thus reformulate the condition to $\text{ENPV}_N^{A_i} \geq A_\omega$.

3.2 Modeling Imperfect Information

When evaluating an investment a firm must estimate how many time periods the investment will entail (N), the cash flow of each time period (C_t), and the independent probability of each time period (P_t), which

enables the derivation of $P_{0 \rightarrow \omega}$ and $P_{t \rightarrow \omega}$. Such data, while disputed, exists (e.g. DiMasi, Grabowski, and Vernon 2004). A firm can thus use the available industry averages for their indication group when calculating ENPV. We will refer to a calculation based on industry averages as publicly Available ENPV (AENPV).

If the firm knew precisely which numbers had been drawn from the distribution then there would be no need to calculate AENPV. Information would be perfect and the firm could forecast against their known numbers. We will refer to a calculation based on this information as Perfect ENPV (PENPV). PENPV may significantly differ from AENPV.

As perfect information is an unrealistic simplification, we follow the approach of Abbott and Vernon (2007) and assume that the firm is able to discern some information from the sample we've drawn (PENPV), and therefore combine this information with the publicly available information (AENPV) to form a new estimate. We assume that this deliberation can be approximated by a linear combination, and refer to the result as Information-Adjusted ENPV (IENPV). Calculating it as $IENPV_N^i = \alpha PENPV_N^i + (1 - \alpha) AENPV_N^i$ where the constant α represents how much information the firm is able to observe. We use $\alpha = 10\%$ which is in line with Abbott and Vernon (2007). We can thus re express the condition for investment as $IENPV_N^i \geq A_\omega$.

3.3 Defining Project State Transitions

We assume that cash flows (i.e. costs or revenues) and durations (i.e. time spent) are incurred regardless of whether a drug is successfully refined from one stage to the next. Thus, C_t (the cash flow associated with t) and D_t (the duration of t) must be incurred before reaching state t , while P_t (the probability of completing t) must be the independent probability of successfully transitioning from t to $t + 1$.

The state transitions of a drug can, if we put terminations aside (due to being deterministic rather than probabilistic), be viewed as a Markov Chain. Loading transitions with cash flow and duration costs, we can express the above offset as in Figure 3a. Figure 3b depicts the state space of a project with three stages. Consequently we calculate the probability of reaching the final state from the initial state as $P_{0 \rightarrow \omega} = \prod_{t=0}^{N-1} P_t$, and the probability of reaching the final state from the current state as $P_{t \rightarrow \omega} = \prod_{t'=t}^{N-1} P_{t'}$.

Figure 3: Markov Chain representations of state transitions.

3.4 From Monte Carlo to Agent Based Modeling

The information presented up to this point can be used to produce an ENPV based Monte Carlo simulation. Similar to the results of Abbott and Vernon (2007) this would make the resulting ENPV, and thus the investment go/no-go decision the output of the simulation. Our work however seeks to move this forward by ultimately considering the outcome resulting from these go/no-go decisions in terms of the number of products that are successfully brought to market. For this reason we suggest an Agent Based approach, where agents are firms that (1) make go/no-go decisions for their projects by calculating IENPV (as elucidated above), and (2) experience actual failures, costs and durations. The simulator is initiated with an input file containing basic parameters (e.g. number of years to simulate), compounds that are in the pipeline, the

developers owning these compounds, and a collection of interventions to be tested (such as a MER). Agent and project characteristics are, at simulation initiation, randomly drawn from that input state space.

A simulation run (Figure 4a) consists of a number of ticks starting at $\hat{t} = 0$ and runs all the way to $\hat{T}$. For every tick, and for every active project, the two procedures `Termination` and `Development` are executed. An active project is a project that has steps remaining and is neither terminated nor failed. `Termination` (Figure 4b) is a procedure that calculates the ENPV of a project from the perspective of the owner and marks the project as terminated if ENPV is too low. `Development` (Figure 4c) is a procedure that subjects the developer of the project to the cash flow of the tick and attempts to refine the project one step further. If refinement fails then the project is marked as failed. Any given simulation run can be thought of as a sequence of state transitions. The model state, essentially behaves as the class diagram depicted in Figure 5.

Figure 4: Simulator flow charts, where S denotes the start and E the end of the flow.

We normalize the stage-based parameters P_t , C_t , and D_t , so that agents can operate in the same dimension of time. I.e. so that any two consecutive ticks, for any two agents, are equidistant. We assume that the cash flow and probability of successfully completing any single stage is uniformly spread over that stage. Thus, the developer will incur some cash flow and run some risk of failing at each simulation tick (step), given that the stage is associated with a non-negative cash flow ($C_t \neq 0$) and runs some risk of failure ($P_t < 1$).

Let $\hat{t}$ denote a normalized time step in the period t , where $D_{\hat{t}} = 1$, for all $\hat{t}$ in t . We can then derive the cash flow of any given step $\hat{t}$ as $C_{\hat{t}} = C_t/D_t$, and the probability of completing any given step $\hat{t}$ as $P_{\hat{t}} = P_t^{1/D_t}$. To exemplify, a stage with a duration of 1.5 years (18 months), a cash flow of -\$18M, and a probability of success of 70%, is in the simulator recalculated to a series of 18 steps where each step entails a cash flow of -\$1M (i.e. 18/18) and has a probability of success of approximately 98% (i.e. $0.7^{1/18}$).

Figure 5: Class diagram of the model state data structure.

3.5 Modeling the Current Pipeline of Antibiotics

Having established that we initiate a simulation with a model of the developers and compounds that are currently in the pipeline, we must now explain how we structure such an input model. For every agent, we must know its profit threshold (A_ω) and its discount rate (A_i). For every antibiotic (D), we must know all the stages it entails (D_S) and for every stage in its sequence of stages ($t \in D_S$) we must know its probability of success (P_t), cash flow (C_t), and duration (D_t).

While the unambiguous specification (JSON Schema) of the structure of valid input data is too verbose to include in this paper, we provide a few illustrative examples. The root object of a simulation input must define the key firms as an array of firms with antibiotics in the pipeline. All firms must define the parameters `discount_rate` (A_i) and `threshold` (A_ω), as well as an array of projects with the antibiotics currently owned by the firm. Each project must define its name and an array of episodes. Each episode must define its type, and an array of stages. Each stage must define the name of its type and the parameters `prob` (P_i), `cash` (C_i), and `time` (D_i), where `cash` can be substituted for `cost` which is treated as the additive inverse of `cash`. An illustrative example of a single firm with an antibiotic expecting significantly fewer stages than a real world antibiotic is depicted in Figure 6a.

```
- name: firm one
  discount_rate: { min: 0.09, mid: 0.11, max: 0.24 }
  threshold: { min: 0, mid: 100, max: 100 }
  projects:
    - name: project one
      episodes:
        - type: development
          stages:
            - type: pre-clinical
              prob: { min: 0.3, mid: 0.6, max: 0.7 }
              cost: { min: 40, mid: 50, max: 60 }
              time: { min: 1, mid: 8, max: 24 }
            - type: phase-1
              prob: { min: 0.3, mid: 0.6, max: 0.7 }
              cost: { min: 40, mid: 50, max: 60 }
              time: { min: 1, mid: 8, max: 24 }
            - type: market
              stages:
                - type: during-patent
                  cash: { min: 40, mid: 50, max: 60 }
                  time: { min: 1, mid: 8, max: 24 }
                  prob: 1
                  name: MER
                  effects:
                    episodes:
                      additions:
                        - target: market
                          cash: { min: 0, max: 10000 }
                          time: { min: 0, max: 120 }
                          prob: 1
                      alterations:
                        - target: market
                          cash:
                            operator: "*"
                            operand: { min: 0, max: 1 }
```

(a) A firm with a project.

(b) A distribution of MERs.

Figure 6: Fictive and partial input data examples.

As our simulator is stochastic, the parameters of any particular firm and antibiotic may be defined using numeric distributions and not merely fixed numbers. Any numeric parameter (except the randomizer seed, the number of years to simulate, and the number of simulation trials to run) of a simulation input may be replaced by a triangular or uniform distribution. When a parameter is defined as a plain number (e.g. the probability of the market stage in Figure 6a) it is treated as a point estimate, when defined as an object with `min` and `max` then it is treated as a uniform distribution, and when defined as an object with `min`, `mid`, and `max` (e.g. the probabilities of the development stages in Figure 6a) it is treated as a triangular distribution. In a simulation run, all firms will be instantiated as agents endowed with instances of projects as defined by their project definitions. All distributions will be sampled.

To emphasize the flexibility of our approach, consider how the probability (`prob`) of the market stages of the project defined in Figure 6a, is set to 1. The project will thus not be able to fail during market years, which can be thought of as not taking pharmacovigilance (i.e. surveillance of drugs after launch) into account.

4 DATA ACQUISITION

To supply the simulator with a model of the antibiotics that are currently in the pipeline, we map existing antibiotic projects to publicly available data based on indication category. Data on what antibiotics are in the

pipeline, who develops them, and what potential indications they may target is sourced from Pew Charitable Trusts (2016), a nonprofit organization with reports on the pipeline of antibiotics in clinical development (i.e. post pre-clinical) for the US market.

Data on expected costs (C_i), development durations (D_i), probability of success (P_i), and cash flow during market years (C_i), for all indications are triangularly distributed and sourced from Sertkaya et al. (2014). The indications are (1) acute bacterial otitis media, (2) acute bacterial skin and skin structure infections, (3) community acquired bacterial pneumonia, (4) complicated intra-abdominal infections, (5) complicated urinary tract infections and (6) hospital acquired/ventilator associated bacterial pneumonia. Similar to Sertkaya et al. (2014) the market cash flows distribution is produced by multiplying each point of the, for the indication in question, distribution of the expected market share with the total market size. While Sertkaya et al. (2014) normalized cash flows to 2012 levels to account for inflation we have not further normalized.

Three problems arise from using the PEW and aggregated indication data when attempting to gauge the efficacy of interventions. First, the data does not indicate how far a given project has progressed in its current phase. This is mitigated by reducing the time remaining for any given project by the delta between the date of project initiation and August 2016. Initiation dates are based on press releases and have been vetted by an expert associated with DRIVE-AB. Second, some projects have several potential indications, but the data for cost, time, attrition rate, market size and market share from Sertkaya et al. (2014) are grouped by single indications. In such cases the indication with the lowest market size has been chosen. Third, data for some indications of drugs reported by PEW are not available in Sertkaya et al. (2014). For these cases, a miscellaneous data set was constructed. All properties of the miscellaneous data set are triangular distributions. For cost, time and attrition rate, the lower limit was set to the lowest lower limit of all the six indications, while the upper limit to the highest upper limit, and the mode to the mean. Market size was set to the mean of all indications.

Data on agent thresholds (A_ω) and discount rates (A_i) stem from Sertkaya et al. (2014). They used a M\$100 threshold, but we explore a larger spectrum by sampling a triangular distribution between \$0 and M\$100 where M\$100 is the mode. Regarding discount rate, it has been suggested that early stage investments are discounted more heavily than late stage investments (Villiger and Hoejer Nilsen 2011). While this is not yet accounted for in our work we, in line with Sertkaya et al. (2014), triangularly distribute firm discount rates between 5% and 24% with a mode of 11%.

5 MODELING INTERVENTIONS

Policy interventions alter properties of a system and may in exchange require some agent(s) to adhere to some condition(s) (e.g. restraining from advertising or obligations to conduct additional trials), see e.g. Rex and Outtersen (2016). In this work we explore effects, but not conditions, which can be thought of as the no-condition condition, where any intervention is granted regardless of circumstances.

In natural language, an intervention effect may be described as: decreasing phase 1 costs by 10%, increasing all market years' expected return by 1M\$, turning the probability of success of discovery into 5%, etc. Simulating a broad set of interventions poses two immediate challenges: (1) constructing a model for unambiguously specifying interventions, and (2) specifying how an instance of the intervention model affects the R&D process described in Section 3.

As established, agents in the simulator use stage-based input data to calculate IENPV of all projects by transforming stages into steps and combining perfect (PENPV) with public information (AENPV). An intervention can be thought of as altering the stage-based model state before sampling, allowing both AENPV and PENPV to be altered. If $\mathbb{S}$ is the set of all possible simulation input states before projects are sampled, then an intervention i that have a discrete effect, can be viewed as the function $i : \mathbb{S} \mapsto \mathbb{S}$ that maps from

one state to another. This describes a large family of interventions, but disregards interventions that change over simulation time (e.g. correcting for inflation) and interventions that appear after the passing of some simulation time (i.e. future interventions). We assume that developers only accept an intervention if it makes them financially better off, i.e. IENPV is higher in the second state than in the first.

In defining the function of an intervention we consider three fundamentally different types of intervention effects: replacements, alterations, and additions. They act upon the stage configuration of some non-proper subset (hereon referred to as target) of the stages of a project. A replacement replaces the target with a new sequence of stages. An alteration varies the values of some properties of the target stages. Finally, an addition adds a sequence of stages to the target that will be executed in parallel (hereon referred to as a track). To specify an intervention target we use string matching. Consider e.g. how the targets in Figure 6b match the type of the market episode in Figure 6a.

The introduction of additions requires reformulation of the ENPV formula to include the cash flows and probabilities of all tracks. Let a be a track belonging to the set of all tracks $\mathbb{A}$ of a project, and let C_t^a , P_t^a , and D_t^a refer to C_t , P_t and D_t within the track. The formula can then be reformulated to:

$$\text{ENPV}_N^i = \sum_{a \in \mathbb{A}} \sum_{t=0}^N \frac{C_t^a P_{0 \rightarrow \omega}^a}{(1+i)^t P_{t \rightarrow \omega}^a}.$$

While our approach lays the foundation for modeling a broad family of interventions, this work focuses on MERs. As elucidated in Section 2, a MER can be thought of as a point in a three dimensional Cartesian coordinate system, where x defines staging, y size, and z fraction of delinkage. By treating the three dimensions as uniformly distributed random variables we can, in Monte Carlo fashion, explore the space of variations. To exemplify, if we're interested in MERs of sizes \$0 to \$B10 ($0 < Y < 10,000$), staged between 0 and 10 years ($0 < X < 120$) with anything between full delinkage and no delinkage ($0 < Z < 1$), we can provide the simulator with the intervention specified in Figure 6b. Note that we're making use of two effects: one addition that introduces the MER cash flows and one alteration that reduces the size of the free market.

6 VERIFICATION & VALIDATION

As opposed to simulating physical or biological systems, the effects of policy on R&D does not readily lend itself to empirical testing. One cannot demonstrate behavioural similarities or perform controlled trials. However, as the high level of stochasticity means that our simulator does not suggest a single future, but rather a space of reasonable outcomes, it is intended to aid policy makers when refining interventions. While we cannot empirically validate these outcomes, we can (1) argue validity in the assumptions, (2) attempt to retrodict and (3) publish predictions that will have to stand the test of time. Regarding the first point, we argue that as ENPV is firmly grounded in microeconomic theory, it is a rational approach to decision making, and that our micro-level assumptions thus reflect the way decisions are de facto being made.

7 EXEMPLARY RESULTS

The following results are examples of what can trivially be derived using our approach. While the simulator supports a wide range of interventions, the here reported results are based on the intervention experiments outlined in Table 1. All size and stage ranges reported in the table are uniformly distributed. All MERs are fully delinked ($z = 0$) and therefore completely replace the expected market of a project.

Data for each experiment stem from 200 runs each progressing a single tick (due to only being interested in initial termination decisions). The input data (available upon request) described in Section 4 was used in all experiments. Summary statistics are reported in Table 1. Comparing any experiment to Experiment 1

Table 1: Summary statistics of all experiments. N=200 per experiment.

	Description	Non-terminated			Description	Non-terminated	
		μ	sd			μ	sd
1	No intervention	1.961	23.915	5	1\$B MER lump-sum	0.889	30.210
2	0-4\$B MER lump-sum	3.322	30.665	6	2\$B MER lump-sum	0.606	32.565
3	0-4\$B MER 5 yrs staged	2.987	30.705	7	1\$B MER 5 yrs staged	1.441	28.230
4	0-4\$B MER 0-5 yrs staged	3.252	30.420	8	2\$B MER 5 yrs staged	0.886	31.410

enables gauging of its benefit. Experiment 2 and 3 allow for optimization of size when either lump sum or 5 years of staging, respectively, is assumed. Experiment 4 is a distributed MER and can be used to optimize the balance between size and staging. Experiments 5-8 are non-distributed MERs and thus test a MER against the stochastics inherent to project and firm data.

The size of an intervention stimulating a finite set of projects expectedly suffers from diminishing returns. This is exhibited in Experiments 2 and 3, and illustrated in Figure 7. When the most expensive project in the set of all projects has been sufficiently stimulated then all less expensive projects must too be sufficiently stimulated and thus the global maxima of non-terminated projects is reached.

Figure 7: Lump sum (Experiment 2) vs 5-year staged MER (Experiment 3) of varying sizes.

Developers with a high discount rate will prefer an early cash flow over a later one. Staging an intervention therefore has a cost in terms of the amount and type of projects the intervention is able to stimulate to continuation (non-termination). This is illustrated in Figure 7 where the intervention size (x) for staged MERs is higher at most numbers of non-terminated projects (y). Optimizing the balance between intervention staging and size can be aided by statistically evaluating experiments such as Experiment 4. Since, as previously argued, large lump sums may be unpalatable, the relevance of such optimization is apparent.

A MER paying a developer substantially more than the developer needs (surplus profits) to pursue development may be a suboptimal resource allocation scheme as the surplus could have been recycled into a new MER. We refer to this phenomena as overcompensation and calculate it as the difference between IENPV and the agent's threshold (A_ω), but set it to 0 if the resulting number is negative (i.e. if the firm was subjectively undercompensated). In Experiment 4, the mean overcompensation for a project was ~14,648\$M, and the highest ~29,001\$M. Importantly, overcompensation regards profits beyond required profits, and not merely profit. Minimizing overcompensation is thus not to deny developers profit, but to optimize the allocation of money into MERs.

8 CONCLUSION

This paper describes a simple Agent Based extension of a Monte Carlo simulation model designed to evaluate the efficacy of variations of a Market Entry Reward intervention intended to incentivize pharmaceutical firms to engage in antibiotics R&D. Previous work (Sertkaya et al. 2016, Sertkaya et al. 2014, Abbott and Vernon 2007) have looked at hypothetical developers developing hypothetical drugs. We refine the mode of investigation by looking at hypothetical developers developing drugs that are currently in the pipeline. Sertkaya et al. (2016) used a decision-tree based on ENPV to investigate a lump sum MER versus a five year delay in competition from generics. We build upon this work by (1) moving from a decision-tree approach to stochastic simulation which enables exploration of a large spectra of characteristics for any given intervention, which in turn enables fine-tuned optimization. Furthermore, by (2) using an information-adjusted ENPV calculation inspired by the work of Abbott and Vernon (2007) which, arguably, more closely resembles the decision-making process of pharmaceutical firms. Abbott and Vernon (2007) used Monte Carlo modeling to evaluate the effects of decreased prices on expected R&D. We build upon this work by (3) introducing a framework that enables modeling of more interventions (e.g. MER) than increased prices, and (4) furthering agent heterogeneity.

9 FUTURE WORK

As part of DRIVE-AB, we have a long list of mandatory enhancements that are mostly concerned with ensuring mimicking of known developer behavior. Allowing projects to change ownership, i.e. sales and/or in-licencing, is crucial to capturing projects that are initially innovated by biotech SMEs and then brought to market through either venture capital or a change of ownership to a large pharmaceutical company. Further, the sensitivity of the information parameter α must be explored, the reduction in time on market based on time to market considered, and experiments of partial delinkage run. Finally, the trade-offs inherent to smaller markets must be explored to understand the effects of sustainable consumption on interventions.

ACKNOWLEDGMENTS

This work has received support from the EU/EFPIA Innovative Medicines Initiative Joint Undertaking DRIVE-AB grant n°115583.

This work does not necessarily represent the view of all DRIVE-AB partners.

REFERENCES

- Abbott, T. A., and J. A. Vernon. 2007. "The Cost of US Pharmaceutical Price Regulation: A Financial Simulation Model of R&D Decisions". *Managerial and Decision Economics* vol. 28 (4/5), pp. 293–306.
- Atrill, P., E. McLaney, and D. Harvey. 2014, August. *Accounting: An Introduction*, 6/E. Pearson Higher Education AU.
- Blau, G. E., J. F. Pekny, V. A. Varma, and P. R. Bunch. 2004. "Managing a portfolio of interdependent new product candidates in the pharmaceutical industry". *Journal of Product Innovation Management* vol. 21 (4), pp. 227–245.
- DiMasi, J. A., H. G. Grabowski, and J. Vernon. 2004. "R & D Costs and Returns by Therapeutic Category". *Drug information Journal* vol. 38, pp. 211–223.
- Harbarth, S., U. Theuretzbacher, and J. Hackett, on behalf of the DRIVE-AB Consortium. 2015, February. "Antibiotic research and development: business as usual?". *Journal of Antimicrobial Chemotherapy*, pp. dkv020.

- Laxminarayan, R., A. Duse, C. Wattal, A. K. M. Zaidi, H. F. L. Wertheim, N. Sumpradit, E. Vlieghe, G. L. Hara, I. M. Gould, H. Goossens, C. Greko, A. D. So, M. Bigdeli, G. Tomson, W. Woodhouse, E. Ombaka, A. Q. Peralta, F. N. Qamar, F. Mir, S. Kariuki, Z. A. Bhutta, A. Coates, R. Bergstrom, G. D. Wright, E. D. Brown, and O. Cars. 2013. "Antibiotic resistance—the need for global solutions". *The Lancet Infectious Diseases* vol. 13 (12), pp. 1057–1098.
- Oттerson, K. 2010. "The legal ecology of resistance: the role of antibiotic resistance in pharmaceutical innovation". *Cardozo L. Rev.* vol. 31 (June 2008), pp. 613–2521.
- Payne, D. J., L. F. Miller, D. Findlay, J. Anderson, and L. Marks. 2015. "Time for a change: addressing R&D and commercialization challenges for antibacterials". *Phil. Trans. R. Soc. B* vol. 370 (1670), pp. 1–12.
- Pew Charitable Trusts 2016. "Antibiotics currently in clinical development". Technical report.
- Rex, J. H., and K. Otterson. 2016. "Antibiotic reimbursement in a model delinked from sales: a benchmark-based worldwide approach". *The Lancet Infectious Diseases* vol. 16 (4), pp. 500–505.
- Sertkaya, A., A. Jessup, and H.-H. Wong. 2016, September. "Promoting Antibacterial Drug Development: Select Policies and Challenges". *Applied Health Economics and Health Policy*, pp. 1–6.
- Sertkaya, S., J. Eyraud, A. Birkenbach, C. Franz, N. Ackerley, V. Overton, and K. Otterson. 2014. "Analytical Framework for Examining the Value of Antibacterial Products". Technical report, Eastern Research Group, Inc.
- Stewart, J. J. 2002, April. "Biotechnology Valuations for the 21st Century". Technical report, Milken Institute.
- Svennebring, A. M., and J. E. Wikberg. 2013. "Net present value approaches for drug discovery". *Springer-Plus* vol. 2 (1), pp. 1.
- Van de Ven, A. H. 1999. *The Innovation Journey*. Oxford University Press.
- Villiger, Ralph and Hoejer Nilsen, Nicolaj 2011. "Discount rates in drug development".
- WHO, and Collaborating Centre for International Drug Monitoring. (Eds.) 2002. *The importance of pharmacovigilance: safety monitoring of medicinal products*. Geneva, WHO. OCLC: 50740044.

AUTHOR BIOGRAPHIES

CHRISTOPHER OKHRAVI is a PhD student of Information Systems. His research concerns simulating the effects of policy intervention on research & development.

STEVE MCKEEVER is an Associate Professor of Information Systems. He holds a PhD in Mathematics from the University of Oxford. A recent focus has been supporting simulations at the life science interface.

CARL KRONLID is a PhD student of Engineering Sciences. His research is focused on innovation and organizing, with a focus on collaborative high-tech innovation projects.

ENRICO BARALDI is a Professor of Engineering Sciences. His research concerns strategies in business networks, innovation, product development and the commercialization of science.

OLOF LINDAHL is an Assistant Professor of International Business. His research interests lie in the management of multinational firms, particularly in the mechanical engineering and pharmaceutical industries.

FRANCESCO CIABUSCHI is a Professor of International Business. His research concerns management of multinationals, corporate strategy, state enterprises, innovation management and entrepreneurship.