

This is a repository copy of *Integrated testing strategies can be optimal for chemical risk classification*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/117767/>

Version: Accepted Version

Article:

Pitchford, Jonathan William orcid.org/0000-0002-8756-0902, Cussens, James orcid.org/0000-0002-1363-2336, Raseta, Marko et al. (1 more author) (2017) Integrated testing strategies can be optimal for chemical risk classification. *Mathematical Biosciences*. pp. 1-8. ISSN 0025-5564

<https://doi.org/10.1016/j.mbs.2017.05.005>

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Integrated testing strategies can be optimal for chemical risk classification[☆]

Marko Raseta*

*Department of Mathematics
University of York, UK*

Jon Pitchford

Department of Biology, University of York, UK

James Cussens

Department of Computer Science, University of York, UK

John Doe

Parker Doe Partnership LLP, UK

Abstract

There is an urgent need to refine strategies for testing the safety of chemical compounds. This need arises both from the financial and ethical costs of animal tests, but also from the opportunities presented by new in-vitro and in-silico alternatives. Here we explore the mathematical theory underpinning the formulation of optimal testing strategies in toxicology. We show how the costs and imprecisions of the various tests, and the variability in exposures and responses of individuals, can be assembled rationally to form a Markov Decision Problem. We compute the corresponding optimal policies using well developed theory based on Dynamic Programming, thereby identifying and overcoming some methodological and logical inconsistencies which may exist in the current toxicological testing. By illustrating our methods for two simple but readily gen-

[☆]This work was supported by the NC3Rs [grant number NC K001264 1]

*Corresponding author

Email addresses: marko.raseta@york.ac.uk (Marko Raseta),
jon.pitchford@york.ac.uk (Jon Pitchford), james.cussens@york.ac.uk (James Cussens),
john.doe@parkerdoe.com (John Doe)

eralisable examples we show how so-called integrated testing strategies, where information of different precisions from different sources is combined and where different initial test outcomes lead to different sets of future tests, can arise naturally as optimal policies.

Keywords: Optimal Integrated Testing Strategies, Dynamic Programming

1. Introduction

Notions of what society deems to be an acceptable testing regime for new chemicals are in a constant state of flux. Until 1999 it was acceptable in the EU to perform tests on guinea pigs in order to determine whether certain cosmetic
5 products were hazardous for human skin (Program (1999)). After 1999 this was replaced by the mouse LLNA (Local Lymph Node Assay), another animal-based method. More recently the EU ethical climate has changed again; by 2018 no new chemical to be used in the cosmetics industry can be tested on animals. Instead, chemicals need to be classified reliably using information from emerg-
10 ing in-vitro and in-silico assays, supplemented where possible by mathematical models. These new methods are likely to be less accurate than in-vivo tests, but are generally cheaper and less ethically problematic to implement. This presents a problem common across toxicology in general: can we make good
15 predictions about the risks associated with new chemicals without using animals at all? In other words, how best can we assemble uncertain information based on non-animal assays, so as to arrive at optimal ethical testing regimes?

Many important papers have emerged on this topic (Gabbert and Weikard (2010), Gabbert and van Ierland (2010), Gabbert and Weikard (2013), Jaworska J
20 (2010), Jaworska et al. (2013), Jaworska et al. (2015), Norlen H (2014)).

Indeed (Gabbert and Weikard (2010)) develops a theory that determines the optimal exposure level of any particular member of the population to the chemical and uses this theory to solve a decision problem of how to pick which chemical to test for hazard first from some finite set of possible chemicals. (Gabbert and

25 van Ierland (2010)) develops a framework allowing one to compute the optimal
battery of tests to assess a generic toxicological endpoint by means of a cost
effectiveness analysis (CEA). By contrast, (Gabbert and Weikard (2013)) devel-
ops a framework in which adaptive cost sensitive Integrated Testing Strategies
can be derived by means of a Value of Information technique (VOI). The au-
30 thors there distinguish between decision problems for competitive businesses
and regulators. Furthermore, (Jaworska et al. (2015)) begins by improving and
generalising previous work (Jaworska J (2010),Jaworska et al. (2013)) by devel-
oping more accurate potency class predictions of skin sensitisation potential of
chemicals via theory of Bayesian Networks and then uses these results together
35 with VOI framework to derive Optimal Integrated Testing Strategies for the
assessment of chemical hazard of chemicals. Finally, similar to (Gabbert and
van Ierland (2010)) and (Gabbert and Weikard (2013)), (Norlen H (2014)) uses
CEA in the context of performing a cost effectiveness analysis in the special
case of acute oral toxicity.

40
However, none of these explicitly accounts for the individual differences between
humans both in the exposure (i.e. environmental variability) and in the toxicity
corresponding caused by that exposure (i.e. individual variability). To be more
specific, only Gabbert and Weikard (2010) introduces a concept of toxicity for-
45 mally but treats it as constant for all members of the population. Moreover,
none of these papers combines these with the financial costs of chemical risk
classification in a mathematically rigorous fashion.

Any new testing strategy must be able to deal rationally with contradictory evidence. For example, one in-vitro assay may predict that certain chemical is a skin sensitiser, while another in-silico assay may predict that the same chemical is actually safe. The classification part of the argument in van der Veen et al. (2014) deals with this problem using a combination of majority voting and Bayesian Statistics. Jaworska J (2010) proposes assembling a Bayesian Network and combines this with the Weight-of-Evidence approach to overcome this issue. Thomas et al. (2012) proposes a strategy of "averaging probabilities", using empirical estimates of precision of each assay and then averaging these out in one "meta-assay". Each of these solutions may be pragmatic and defensible within the authors' given problem, but an over-arching logical framework would be a helpful step in confirming the value and risk associated with the removal of animal tests.

In what follows we shall propose a mathematical framework which seeks to simultaneously overcome the shortcomings mentioned above. The issues of imprecision, and of environment- and individual-level variability, fall naturally within theories developed for evolutionary ecology (Currey et al. (2007)). The efficient assimilation of evidence can then be handled by well-developed theories of Markov Decision Processes (Bellman (1957)).

We finish this chapter by surveying previous work in Toxicology and Medicine that is based on this mathematical theory. To the best of our knowledge very little work in Toxicology uses Markov Decision Processes, to be more precise, these are the works of Chang (2010) and Korthikanti et al. (2010). Chang (2010) is a rich summary of techniques used in contemporary in house pharmacological research and decision making illustrated with numerous examples. Among a vast range of mathematical and computational techniques used are the Markov Decision Processes as applied to the optimal decision making of a pharmaceutical business on whether to proceed from earlier (Phase 1 clinical trial) to later (Phases 2 and 3 of clinical trials) stages. Our work generalises this work in a

number of directions. Firstly, the models in Chang (2010) are developed for
80 the sake of a making a commercial business more profitable and do not take
into account the regulatory aspect of Toxicology, i.e. the fact that the company
may actually incur fines from regulatory bodies and lawsuits from individual
consumers in case they exhibit adverse outcomes as a consequence of using the
drug. The Markov Decision Process model in this paper takes this into account
85 via the mechanisms of misclassification costs: in case the company declares an
unsafe chemical as safe there will be serious consequences. Equally, if the com-
pany actually declares a safe chemical as unsafe it will lose money by not selling
the safe product in the market for which it possibly had an advantage over its
competitors. Thus our work bridges the two worlds: it allows the company to
90 maximise its profits while simultaneously acts in the best interest of the gen-
eral public. Secondly, Chang (2010) does not take neither variable exposure
to chemical among different members of the target population nor precision of
measurements used to test safety in the account. Instead, it relies on toxic-
ity levels observed in recently tested chemicals to draw conclusions on the new
95 chemical of interest while the transition probabilities of the Markov Decision
Model model are estimated from historical data and power calculations which,
by nature, cannot guarantee precisions of estimates in advance in the case of
unknown moments. Another problem with this approach is that there is no
guarantee the new chemical will share toxicity thresholds with the previously
100 used chemicals. These issues motivate the truly novel part of the methodological
work presented in this paper which is applicable in a general setting not necessar-
ily restricted to Toxicology and Medicine. Namely, the transition probabilities
between states of the Markov Decision Process in this paper are based on the
idea of integrating evidence from different measurements of the same quantity
105 in a non-contradictory way by using the information on the precision from the
instruments/assays used in the process. Knowing the value obtained in a less
accurate measurement and its precision we can get a probability distribution on
the values more accurate measurement of the same quantity can possibly take.
This simple observation has far reaching consequences; namely since the states

110 of the Markov Decision Process in our model are the collection of measurements
the above allows for a logically sound way of defining transition probabilities of
the model by first performing cheapest possible tests for each parameter of in-
terest. This procedure is justified by the existence of devices of varying precision
in many fields of human work. As far as Toxicology is concerned, this corre-
115 sponds to in-vivo, in-vitro and in-silico tests. By choosing to start our analysis
with a cheap in-silico test we remedy the problem outlined above: indeed we
get the transition probabilities of the model without having to resort to further
unknown characteristics of the chemical therefore bypassing a potentially cir-
cular argument. Another interesting work involving Markov Decision Process
120 in Toxicology is a theoretical paper of Korthikanti et al. (2010). The authors
approach the problem of model checking for a Markov Decision Process from
the Computer Science point of view using the language of Mathematical Logic;
actually as it turns out, the authors do not work with a conventional proba-
bilistic definition of a Markov Decision Process but define their own in another
125 set-up. Although motivated by an example of Insulin compartment model the
paper soon drifts into proving results in Mathematical Logic and holds little
practical value.

When it comes to applications in Medicine the literature is much larger. We
130 discuss in detail a variety of different applications (Fakih (2006), Kurt et al.
(2011), Shechter et al. (2008), Alagoz et al. (2007), Alterovitz et al. (2008),
Nunes et al. (2017) and Sloan (2007)) . As mentioned in the above, the main
novelty of this paper, logically consistent aggregation of evidence from differ-
ent measurements in a non-contradictory way and without a need to resort to
135 empirical estimates still stands. Fakih (2006) develops a general framework for
learning efficient approaches to medical diagnosis. It resembles this manuscript
and Bayer-Zubek (2004) in the sense that the states are cumulative history of
observations, which in turn, guarantees the Markovian nature of the process but
it resorts to the empirical estimates of transition probabilities between states
140 based on observed frequencies. Kurt et al. (2011) develops a discrete time,

but unlike this paper, an infinite-horizon Markov Decision Process to maximise the patient's quality-adjusted life years prior to them having either a stroke or developing a Coronary Heart Disease. The infinite horizon is justified by the large number of visits to the doctor by Type 2 Diabetes patients. The model
145 resembles the one of ours in that it has a terminal state which in turn insures convergence. Finally, transition probabilities are computed as a combination of equations based on medical knowledge and empirical observations. Shechter et al. (2008) develops a Markov Decision Process that aims to maximise the expected lifetime or quality-adjusted life years. Similarly to Kurt et al. (2011) this
150 model contains an absorbing state which can be reached from any other state and is an infinite-horizon problem for the same reason as Kurt et al. (2011). Furthermore, exactly as Kurt et al. (2011), the rewards in the problem are not measured only in monetary units as in our case but instead in the units of Health Economics, namely quality-adjusted life years, while the transition
155 probabilities of the model are estimated empirically from data. Alagoz et al. (2007) develops a Markov Decision Process for optimal choice of when to go for a liver transplant and then, furthermore, should one accept the part of a liver of a living-donor or the entire cadaveric liver. This is again a discrete-time infinite-horizon problem for the same reasons as in the previous papers. Again,
160 the model contains a naturally emerging absorbing state while the cost is measured in Health Economics units. Transition probabilities are estimated from clinical data. Furthermore, the authors derive in closed form the sufficient conditions for the existence of simple threshold-based optimal policies. Alterovitz et al. (2008) develops a Markov Decision Process to optimise the probability of
165 the steering needle reaching a target through the noisy environment generated by the uncertainties emerging from the needle- soft tissue interactions and other obstacles that are too small to detect with great but still limited resolution of medical images. The problem does not have costs in the sense of the papers above. They formulate the problem as a discrete time acyclic Markov Decision
170 Process and solve it by value iteration. Finally, the transition probabilities do not emerge from the imprecision of measurements of needle's position, turning

angle and velocity but instead are a consequence of the stochastic environment in which the needle operates. Nunes et al. (2017) develops a Markov Decision Problem for the optimal control of hospital admissions for non-emergency patients. Transition probabilities are computed by means of an interplay between
175 a complicated model based on deep understanding of the specific process and empirical considerations. This is a discrete state model considered where the limit of expected average reward is the object to maximise while the costs are expressed in monetary units. Sloan (2007) develops a Markov Decision Process
180 to determine the optimal replacement of broken medical equipment with either new equipment or already used (but sterilised and refurbished) equipment. As in this paper, cost and penalties are expressed in monetary terms. This is a discrete time Markov Decision Process. Transition probabilities between states are made on the authors' guesses and serve to demonstrate how the optimal
185 policies change subject to variation of these. This is similar to our paper with the difference that we study the sensitivity of optimal policies to monetary costs instead.

In summary, we develop a Markov Decision Process for the assessment of hazard
190 of new chemicals while taking into account the exposure, inter-person variability in toxicity and misclassification costs. The result is optimal cost-effective sequential testing strategies for assessment of risk of chemicals. These so-called "optimal policies" rationally integrate outputs from tests with different costs and precisions. The integrated testing frameworks which emerge achieve the
195 keys aims outlined in Pastoor et al. (2014), Embry et al. (2014).

2. Theory

2.1. Individual- and Population-level Models

Much existing literature in quantitative toxicology seems to ignore, or implicitly to average out, the large differences that may exist between the exposure
200 of individuals to potentially toxic chemicals, and the variable effects of a fixed

exposure to these chemicals on these individuals. When the purpose of testing is to identify the risk of rare toxicologically harmful outcomes, it seems mathematically expedient to account for this variability explicitly.

205 We start by providing an alternative view on exposure. Since we cannot know the exact exposure of every single human to the chemical of interest, the natural way to model it is probabilistically. We shall therefore denote exposure experienced by an individual by a random variable X having some distribution with density function f , say, where f would typically correspond to one of many
210 classical distributions (normal, log-normal, gamma). For illustrative purposes, here we will assume f is characterised by a single unknown parameter a^* , say. This could correspond, for example, to the variance of the exposure distribution.

We now turn to modelling the variability in toxicity caused by a given exposure. A typical dose-response curve is sigmoidal: Toxicity is typically small up to some threshold level of exposure, beyond which it increases rapidly with exposure before eventually saturating. This behaviour is efficiently captured by a relationship of the form

$$g(x) = Ax^p/(1 + x^p),$$

where A is the level at which toxicity saturates, x is the exposure and p is a
215 model parameter to be inferred from data. However, due to individual-level variability, it makes sense to model toxicity at a given level of exposure as a random variable, having mean specified by the function g together with some unknown variance to be inferred from data. Rigorously, we shall denote toxicity at a level of exposure $X = x$ by a random variable T^x . This random variable
220 will have mean

$$\mathbb{E}(T^x) := g(x)$$

Similarly to the case of exposure (above), we denote the probability density function h^x for the random variable T^x as being characterised by a single parameter b^* .

225 For the sake of transparency, in what follows assume A and p are known, and
that the unknown parameters are a^* (relating to exposure variability) and b^*
(relating to individual variability in response to exposure). The purpose of any
optimal testing regime will be, broadly, to find the most efficient way to estimate
ranges for a^* and b^* so as to determine whether or not a chemical is safe. This
230 purpose will be made more precise below, while the generalisation to wider sets
of unknown parameters is left to the Discussion.

2.2. Definition of safety

From the point of view of a pharmaceutical company it is impossible to predict
whether a particular individual who buys their product will be adversely affected
235 by it. Again, only probabilistic language makes sense: for example a criterion
that no more than 1 in 1000000 customers is affected at a level exceeding some
defined threshold. Suppose γ represents this threshold of toxicity and let T
be the random variable indicating the toxicity experienced by an individual
sampled randomly from the population. Rigorously, using the logic and notation
240 introduced for exposure and toxicity above, we will declare our chemical of
interest to be safe for human use if and only if

$$P(T \geq \gamma) \leq 10^{-6}. \quad (1)$$

By conditioning on the exposure level and using the generalised law of total
probability we obtain:

$$P(T \geq \gamma) = \int_0^\infty P(T^x \geq \gamma) f(x) dx \quad (2)$$

$$= \int_0^\infty \int_\gamma^\infty h^x(t) dt f(x) dx. \quad (3)$$

In other words, the chemical is safe for human use if and only if:

$$\int_0^\infty \int_\gamma^\infty h^x(t) f(x) dx dt \leq 10^{-6}. \quad (4)$$

245 The important mathematical step here it that calculating the probability a
chemical is safe is reduced to finding the value of a well defined (but unknown)

double integral i.e. to evaluating a (possibly complicated) function of a^* and b^* alone. Thus (4) can be, in general, written as

$$F(a^*, b^*) \leq 10^{-6} \tag{5}$$

for some well defined function F . In other words, all we need to deduce whether
250 our chemical is safe or not is knowledge of the parameters a^* and b^* . Some
values of a^* and b^* correspond to safe chemicals, and some to unsafe chemicals.
It is the job the any testing strategy to identify the real values of a^* and b^* with
sufficient accuracy to determine in which category to place the chemical.

2.3. Precision of Measurements

255 Biological practice shows that one hardly ever obtains the same numerical value
upon repeating the same test in the laboratory. Rather, the values observed are
random variables. Suppose we have several in-vitro and in-silico tests at our
disposal for the assessment of our chemical of interest. Some tests will relate to
exposure (i.e. tests attempting to estimate a^*) and some will relate to toxicity
260 (to estimate b^*). All tests will be inaccurate, and we assume that in-vivo tests
are the most accurate, followed by in-vitro and in-silico tests respectively.

For mathematical convenience we assume that test outcomes can only take a
finite number of values (i.e. that the data are discrete rather than continuous).
265 This simply allows integrals such as equation 4 to be expressed as finite sums, which
does not lose generality because the scale of discretisation can be arbitrarily fine.

Specifically, suppose that a series of n independent in-vitro tests for a^* has out-
comes (random variables) a_1^v, \dots, a_n^v , and that m independent in-silico tests for a^*
270 have outcomes a_1^s, \dots, a_m^s . Similarly let b_1^v, \dots, b_p^v and b_1^s, \dots, b_q^s be the outcomes of p
independent in-vitro, and q independent in-silico, tests for b^* .

Because in-vitro tests are likely to be more accurate than in-silico tests, it is
reasonable to assume that the random variables a_i^v have a smaller variance than

275 the a_i^s (and similarly for b). Explicitly, we define constants l_i and r_j (for $i = 1 \dots n$
and $j = 1 \dots m$) and δ_i and ε_j (for $i = 1 \dots q$ and $j = 1 \dots q$) which capture the
accuracy of these tests, such that each $|a_i^v - a^*| \leq l_i$ and each $|a_i^s - a^*| \leq r_i$,
and similarly each $|b_j^v - b^*| \leq \delta_j$ and each $|b_j^s - b^*| \leq \varepsilon_j$. Note that the variances
of the l_i and δ_j are smaller than those of the r_i and ε_j , respectively.

280 2.4. The cost structure

There are several types of costs associated with toxicological testing. First of all,
carrying out each assay is costly; although hopefully much cheaper than in vivo
studies, even in silico approaches require computing infrastructure and human
resources. There are also costs associated to the misclassification of chemicals.
285 For example, a "false negative" classifying an unsafe chemical as safe (see (32)
for the definition of safety) will involve costs from lawsuits and negative public-
ity. On the other hand, a "false positive" classifying a safe chemical as unsafe
would yield costs in terms of wasted R&D and the loss of income which could
be generated from selling the product. In what follows we shall be using the
290 following notation:

- Ca_i^v = cost of in vitro assay i related to estimating a^*
- Ca_i^s = cost of in silico assay i related to estimating a^*
- Cb_j^v = cost of in vitro assay j related to estimating b^*
- Cb_j^s = cost of in vitro assay $\#i$ related to estimating b^*
- 295 $MC_{falseneg}$ = cost of classifying an unsafe chemical as safe
- $MC_{falsepos}$ = Cost of classifying a safe chemical as unsafe

300 Within the framework developed here, we will show how these costs can, and
indeed should, become a part of the development of optimal testing strategies.

2.5. Markov Decision Problems

Our ultimate goal is to find optimal policies for sequential decision making under
uncertainty. This is likely to involve the decision about which tests are

necessary to depend on the outcomes of earlier tests on the same, or a similar,
305 chemical (Bayer-Zubek (2004)). The optimal testing policy will be the one that
minimises the expected total cost, including the costs of misclassification. Due
to limited resources it is clear that the decision of whether to declare the chem-
ical safe or unsafe has to be reached after some finite number of tests have been
performed. The theory of (finite horizon) Markov Decision Problems provides
310 the necessary mathematical machinery, by describing the interaction between a
decision maker and the environment.

The general set up for Markov Decision Problems is as follows (see Bellman
(1957) for a more detailed treatment). At each (discrete) time step, the testing
315 process is in some state s , say, and the decision maker chooses one of several
available actions, a , say. This action moves the system to some new state s'
and gives the decision maker reward $R(s, a, s')$. In our case in costs which can
simply be thought of as negative rewards. The Markov part of the name comes
from the demand that the probability of transition from state s to state s' upon
320 choosing action a depends only on the state s , and not on the previous history
of the process. This function is denoted by $P(s'|s, a)$. The ultimate goal of the
decision maker is to maximise the expected reward (minimise the expected loss)
by choosing her actions optimally.

325 Formally speaking, a policy π is simply a function that maps states to actions.
The value function of a state s under some policy π , denoted $V^\pi(s)$ is formally
defined by:

$$V^\pi(s) = E[\sum R(s, a, s') | s_0 = s]. \quad (6)$$

In other words, the values function is the expected sum of future costs/rewards
on the assumption that we start at state s . Moreover, we introduce the special
330 state $s_{terminal}$ which is the absorbing state, meaning that when the process

reaches this state it remains there forever. We demand that

$$V^\pi(s_{terminal}) = 0 \tag{7}$$

This relation must hold for all policies π simultaneously.

The optimal value function V^* of the Markov Decision Process specified above

335 is:

$$V^*(s) = \min_{\pi} V^\pi(s), \text{ for all states } s \tag{8}$$

It can be demonstrated that V^* satisfies the following equation (Bayer-Zubek (2004)):

$$V^*(s) = \min_{\text{all actions } a} \sum_{\text{all states } s'} P(s'|s; a)[R(s, a, s') + V^*(s')] \tag{9}$$

340 Again, this relation has to be specified for all states simultaneously. The optimal policy π^* is defined state-by-state according to the following relation:

$$\pi^*(s) = \arg \min_{\text{all states } a} \sum_{\text{all states } s'} P(s'|s, a)[R(s, a, s') + V^*(s')] \tag{10}$$

2.6. Chemical Risk Classification Problem as a Markov Decision Problem

Recall that, in our general framework, the chemical is safe if and only if:

$$F(a^*, b^*) \leq 10^{-6} \tag{11}$$

where a^* and b^* are the parameters we seek to estimate using various tests.

345 Recall also that we have n in vitro and m in silico tests for estimating a^* . Similarly, we have p in vitro and q in silico tests for estimating b^* . The accuracies and costs of these tests are as defined in sections 2.3 and 2.5.

Essentially, our goal is to come up with a sequential testing strategy which tells

350 us whether (12) holds or not while incurring the smallest expected cost. In what

follows we shall describe how to write this as a Markov Decision Problem. To this end we follow the reasoning from (Bayer-Zubek (2004)) very closely indeed. We elaborated on the analogy between problem solved there and our problem in the Introduction. The only fundamental difference is that we have to propose
 355 the way of defining the transition probabilities whereas in (Bayer-Zubek (2004)) the authors estimate these are estimated empirically from data.

We have seen in the previous section that a Markov Decision Problem is fully specified by its actions, states and transition probabilities. We first describe the
 360 actions of our Markov Decision Problem. These comprise the $n + m + p + q$ available tests and two classification actions ("classify as safe" and "classify as unsafe"). States will represent the complete history of previously observed values upon testing, namely the results of performing various in vitro and/or in silico tests. Indeed, let s stand for the current state of the system. Moreover,
 365 assume that the next test to be done is G , say, where G can be any of the tests available to us. Then the system will move to the state $s' = s \cup \{G = g\}$ where g is one of the possible values of G . The probability of this transition will therefore depend on all the values previously obtained by measurements which is the information kept in s . In other words:

$$\mathbb{P}(s'|s, G) = \mathbb{P}(G = g|s) \tag{12}$$

370 Observe that this ensures Markovian nature of the problem. There will be an initial state which will correspond to the observed values obtained by performing two cheapest in silico tests (or alternatively, expert knowledge), one for each of the parameters. This provides us with a (potentially very wide due to imprecision) initial range of values which a^* and b^* can take, at a lowest cost possible.

375

We note that, because the same test cannot usefully be performed indefinitely, our Markov Decision Problem is acyclic (Bayer-Zubek (2004)), and that this implies important computational consequences (see sections 3.1 and 3.2). Finally,

we introduce a special state called the terminal state, denoted by $s_{terminal}$. This
380 has the property that every classification action (independently of the state at
which it has been performed) forces transition to $s_{terminal}$ with probability one.
In other words, this state is absorbing, indicating that an evidence-based clas-
sification decision has been reached.

385 All that remains is to introduce the costs of our available classification actions.
We introduce the quantity $MC(\omega, \zeta)$ to stand for the misclassification cost for
classifying a chemical as ω once its true class is ζ ; $\omega, \zeta \in \{safe, unsafe\}$. Since
we do not know the true risk class of the chemical at the time of testing (hence
the testing), a natural way of thinking of misclassification cost is to treat it as a
390 random variable. We will say that the cost of classification action ω performed
at some state s takes value $MC(\omega, \zeta)$ with probability $P(\zeta|s)$, where this last
quantity stands for the chance that the true risk class of the chemical is ζ given
that the system is at state s . The expected cost of the action is then:

$$R(s, \omega) = \sum_{\zeta} P(\zeta|s) MC(\omega, \zeta) \quad (13)$$

395 Putting all this together, and following Bayer-Zubek (2004), we can deduce that
the equation (9) for the optimal value function collapses to

$$V^*(s) = \min_{a \in \text{all actions available at state } s} [R(s, a) + \sum_{s'} P(s'|s, a) V^*(s')] \quad (14)$$

3. Calculation

The preceding work was general and technical, but showed that the problem of
400 finding optimal testing policies for toxicological testing can be mapped to the
existing mathematical framework of Markov Decision Processes. To show the
value of this result, and to explore its general consequences, it is expedient to
proceed via examples.

405 In these examples it is assumed that the outcome of each test, and the values
of the unknown parameters a^* and b^* , are all integers. This is not a practical
limitation, because these integers may refer to measurements at an arbitrarily
fine scale; this assumption simply allows methods from discrete (as opposed to
continuous) mathematics to be applied and exact results to be computed. As
410 described above, we start by performing the cheapest in silico tests for a^* and b^*
(one for each). For sake of simplicity and illustration, we assume the outcomes
of these tests are

$$a_1^s = 6, b_2^s = 6. \tag{15}$$

and that the accuracy of these tests is such that $|a_1^s - a^*| \leq 3$ and $|b_1^s - b^*| \leq 3$.
This immediately implies that $a^* \in \{3, 4, 5, 6, 7, 8, 9\}$ and $b^* \in \{3, 4, 5, 6, 7, 8, 9\}$.

415

Moreover, assume relation (10) defines the following points in the (a^*, b^*) -plane:

$$\begin{aligned} K := \{(a^*, b^*) \in \mathbb{R}^2 : F(a^*, b^*) \leq 10^{-6}\} = & \{(3, 6), (4, 5), (4, 6), (4, 7), (5, 4), \\ & (5, 5), (5, 6), (5, 7), (5, 8), (6, 3), (6, 4), (6, 5), (6, 6), (6, 7), (6, 8), \\ & (6, 9), (7, 4), (7, 5), (7, 6), (7, 7), (7, 8), (8, 4), (8, 6), (8, 7), (9, 6)\} \end{aligned}$$

25 "safe points" of region K are indicated in the Figure 1 below via red points.
Remaining 24 points on the grid represent the "unsafe points".




Figure 1: region K

In this set up, we know from our first tests that $(a^*, b^*) \in \{3, 4, 5, 6, 7, 8, 9\} \times$
 420 $\{3, 4, 5, 6, 7, 8, 9\}$ (where \times denotes Cartesian product). The aim of the testing
 strategy is to simply to determine, with maximum efficiency, whether or not
 (a^*, b^*) is an element of K .

3.1. Example 1: The case of independent evidence

Consider a simple situation where, on top of the two initial in silico tests, we
 425 have one more in vitro test available for each of our parameters a^* and b^* .
 These will be denoted by a^v and b^v , respectively. We also have two possible
 classification actions, namely to classify our chemical as either safe or unsafe.
 Recall that the goal is to come up with the optimal sequential testing strategy,
 that is the one that minimises the expected cost. Evidence is independent in
 430 the sense that a^* -related tests tell us nothing on b^* and vice versa (we shall
 deal with the case of dependent evidence on the next example). In the spirit
 of Section 2.6 we introduce states, actions and transition probabilities needed
 for specifying the corresponding Markov Decision Problem. There will be the

initial state corresponding to the observation (15). We will denote this state by $\{\}$. Moreover, we will assume in vitro tests are more precise than their in silico counterparts, i.e. that they have a lower absolute measurement error. In order to capture this idea we will assume that:

$$|a^v - a^*| \leq 2 \text{ and } |b^v - b^*| \leq 2 \quad (16)$$

Using the above one deduces that

$$a^v \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}, b^v \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\} \quad (17)$$

This together with the idea the fact that a^* and b^* related tests are independent, gives rise to $11 + 11 + 121 = 143$ more states of the system, all listed below:

$$\begin{aligned} & \{a^v = 1\} \dots \{a^v = 11\} \\ & \{b^v = 1\}, \dots \{b^v = 11\} \\ & \{a^v = 1, b^v = 1\} \dots \{a^v = 1, b^v = 11\} \\ & \qquad \qquad \qquad \vdots \\ & \{a^v = 11, b^v = 1\}, \dots \{a^v = 11, b^v = 11\} \end{aligned}$$

Finally, we have the terminal state, $s_{terminal}$, defined exactly as before.

We now move on to describing the available actions and corresponding transition probabilities. At the initial state $\{\}$, there are four actions available: classify as safe, classify as unsafe, do the a^v test or do the b^v test. This means we must specify the following quantities: $P(is\ safe|\{\}, \text{any classification action})$, $P(a^v = i|\{\}, \text{do } a^v \text{ test})$, $P(b^v = i|\{\}, \text{do } b^v \text{ test})$; where $i \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$.

We start with the first quantity in the above. We have seen that we must have: $(a^*, b^*) \in \{3, 4, 5, 6, 7, 8, 9\} \times \{3, 4, 5, 6, 7, 8, 9\}$, so that there are 49 values the ordered pair (a^*, b^*) can possibly take, only 25 of which belong to K . Thus, in

the absence of any further knowledge, it makes sense to define:

$$P(\text{chemical is safe}|\{\}, \text{any classification action}) = 25/49 \quad (18)$$

We now present an argument for defining the quantities of the form $P(a^v = i|\{\}, \text{do } a^v \text{ test})$, $i \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$. Notice that, due to symme-
 455 try, the quantities $P(b^v = i|\{\}, \text{do } b^v \text{ test})$, $i \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$, can be defined in an analogous way.

For $j \in \{3, 4, 5, 6, 7, 8, 9\}$ define sets T_j by

$$T_j := \{j - 2, j - 1, j, j + 1, j + 2\} \quad (19)$$

Notice than that a^v must belong to at least one of the T_j 's since the absolute
 460 error of the in vitro measurement is 2. With this in mind we define our quantities of interest as follows:

$$P(a^v = i|\{\}, \text{do } a^v \text{ test}) := \frac{\sum_{j=3}^9 I\{i \in T_j\}}{\sum_{k=1}^{11} \sum_{j=3}^9 I\{k \in T_j\}} \quad (20)$$

The idea behind (20) is simple. Observe that if we knew the true value of a^* were j , say, (16) would immediately tell us that $a^v \in T_j$ (since the absolute error of this measurement is 2). The problem, clearly, is that we do not know the exact value of a^* , only that a^* is an element of $\{3, 4, 5, 6, 7, 8, 9\}$. First, we count all those intervals T_j that contain i . Secondly, we scale appropriately to get the normalising constant of the probability density. We then define the probability of interest as a function of the two. For example, notice there is only one suchlike interval that contains value 1, two of them that contain value 2, three of them that contain value 3 and so on. Moreover, let us note that this is only one of the possible ways to define the transition probabilities in a manner consistent with precision of measurements and non-contradictory nature of evidence and is by no means the only one. Thus, the random variable $a^v|\{\}$ has the following

distribution:

$$a^v |_{\{\}, do\ a^v\ test} = \begin{cases} 1 & \text{with prob } 1/35 \\ 2 & \text{with prob } 2/35 \\ 3 & \text{with prob } 3/35 \\ 4 & \text{with prob } 4/35 \\ 5 & \text{with prob } 1/7 \\ 6 & \text{with prob } 1/7 \\ 7 & \text{with prob } 1/7 \\ 8 & \text{with prob } 4/35 \\ 9 & \text{with prob } 3/35 \\ 10 & \text{with prob } 2/35 \\ 11 & \text{with prob } 1/35 \end{cases}$$

As promised before, the corresponding counterpart for b^v test is defined by:

$$b^v |_{\{\}, do\ b^v\ test} = \begin{cases} 1 & \text{with prob } 1/35 \\ 2 & \text{with prob } 2/35 \\ 3 & \text{with prob } 3/35 \\ 4 & \text{with prob } 4/35 \\ 5 & \text{with prob } 1/7 \\ 6 & \text{with prob } 1/7 \\ 7 & \text{with prob } 1/7 \\ 8 & \text{with prob } 4/35 \\ 9 & \text{with prob } 3/35 \\ 10 & \text{with prob } 2/35 \\ 11 & \text{with prob } 1/35 \end{cases} \quad (21)$$

Now, at the states of the form $\{a^v = i\}$, $i \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$ we have only 3 actions available, namely do b^v test or perform one of the two

classification actions available. This, due to the independence of measurements, simply means that one must define

$$P(\text{chemical is safe} | a^v = i) \text{ for } i \in \{1, \dots, 11\}.$$

Let $A_0 = B_0 = \{3, 4, 5, 6, 7, 8, 9\}$. Thus A_0 and B_0 can be associated with our knowledge of a^* and b^* at state $\{\}$, respectively. However, we now have
 465 new information, namely that $a^v = i$. This, by using the principle of non-contradictory nature of evidence, means that $a^* \in A_0 \cap T_i$. Since we have not done any further b^* -related tests at this stage it makes sense to define (although this can be done in a variety of different ways we believe the one below is the simplest) the following quantity:

$$P(\text{chemical is safe} | a^v = i) := \frac{\#((A_0 \cap T_i) \times B_0) \cap K}{\#((A_0 \cap T_i) \times B_0)} \quad (22)$$

470 Here $\#D$ stands for the number of elements in the set D . Indeed, numerator in (22) counts only those ordered pairs (a^*, b^*) that are consistent with the new information and correspond to the safe chemical, whereas the denominator corresponds to all ordered pairs that are consistent with the new information. Yet again, using symmetry and independence, actions in the states of the form
 475 $\{b^v = i\}$, $i \in \{1, \dots, 11\}$ are fully specified upon defining:

$$P(\text{chemical is safe} | b^v = j) := \frac{\#((A_0 \times (T_j \cap B_0)) \cap K)}{\#(A_0 \times (T_j \cap B_0))} \quad (23)$$

Finally, it remains to specify the actions and transition probabilities for the states of the form $\{a^v = i, b^v = j\}$ $i, j \in \{1, \dots, 11\}$.

Clearly, there are only two actions available in such states, and these are to
 480 classify as safe or unsafe. As before, this boils down to computing $P(\text{safe} | a^v = i, b^v = j)$. Using the same notation and ideas as before we define these as follows:

$$P(\text{safe} | a^v = i, b^v = j) := \frac{\#(((A_0 \cap T_i) \times (B_0 \cap T_j)) \cap K)}{\#((A_0 \cap T_i) \times (B_0 \cap T_j))} \quad (24)$$

We are now ready to compute the corresponding optimal policies. Because our Markov Decision Problem is acyclic, equation (14) for the optimal value function can be solved in a single sweep through the states. This means we start by computing the optimal value function at the so called "leaf states", that is the states of the form $\{a^v = i, b^v = j\}$ and then work our way upwards to the states of the form $\{a^v = i\}, \{b^v = j\}$ (notice we need the information above to compute these). Finally, we finish with the computation of $V^*(\{\})$ having computed all other necessary quantities along the way. The optimal policy is then read from top down. The actual computations were implemented using MATLAB, code is available upon request.

By manipulating the attribute and misclassification costs we were able to obtain several classes of fundamentally different optimal policies. This clearly demonstrates that, even in this apparently simple case, the optimal policy emerges from a complex interplay between the geometry of the feasible region K and the costs present in the problem. Moreover, this shows that simple non-adaptive policies are unlikely to be optimal, that is, having changed your population model and the cost structure you have no reason whatsoever to believe the optimal policy will not change as well. Indeed, one would have to repeat all the steps in the above and compute the optimal policy corresponding to one's favourite interplay between the ingredients. Finally, let us note that this gives the problem a strongly Bayesian flavour. The policies thus obtained, and some motivation for what they emerge, are summarized below. For convenience we introduce the "cost vector" Ω to represent all the costs associated with the problem. Mathematically:

$$\Omega := (MC_{falsepos}, MC_{falseneg}, Ca^v, Cb^v) \quad (25)$$

First we look at the following situation:

$$\Omega = (200, 100, 80, 50) \tag{26}$$

510 The optimal policy is $\pi^*(\{\}) = \text{"classify as unsafe"}$. In plain English means that it is optimal to classify the chemical as unsafe without doing any further test at all. The cost of testing, and of misclassification, mean that the chemical's development should be abandoned.

515 Secondly, we look at the following cost structure:

$$\Omega = (100, 200, 80, 50) \tag{27}$$

In this case we obtain $\pi^*(\{\}) = \text{"classify as safe"}$. In contrast to Case 1, the change in the balance between the consequences of false negatives versus false positives means that we should immediately classify our chemical as safe.

520 In reality, the cost of testing is likely to be dwarfed by the cost of misclassification, with false negatives being particularly costly. This motivates the following example captured via respective cost structure:

$$\Omega = (10000, 20000, 80, 50) \tag{28}$$

In this case we obtain the following:

525 $\pi^*(\{\}) = \text{"do } b^v \text{ test"}$
 $\pi^*(\{b^v \in \{4, 5, 6, 7, 8\}\}) = \text{"do } a^v \text{ test"}$
 $\pi^*(\{b^v \in \{1, 2, 3, 9, 10, 11\}\}) = \text{"classify as unsafe"}$
 $\pi^*(\{b^v = 4, a^v = 6\}) = \text{"classify as safe"}$
 $\pi^*(\{b^v = 4, a^v = i\}) = \text{"classify as unsafe" for all } i \in \{1, 2, 3, 4, 5, 7, 8, 9, 10, 11\}$
530 $\pi^*(\{b^v = 5, a^v \in \{5, 6, 7\}\}) = \text{"classify as safe"}$
 $\pi^*(\{b^v = 5, a^v \in \{1, 2, 3, 4, 8, 9, 10, 11\}\}) = \text{"classify as unsafe"}$
 $\pi^*(\{b^v = 6, a^v \in \{4, 5, 6, 7, 8\}\}) = \text{"classify as safe"}$

$\pi^*({b^v = 6, a^v \in \{1, 2, 3, 9, 10, 11\}}) = \text{"classify as unsafe"}$
 $\pi^*({b^v = 7, a^v \in \{5, 6, 7\}}) = \text{"classify as safe"}$
535 $\pi^*({b^v = 7, a^v \in \{1, 2, 3, 4, 8, 9, 10, 11\}}) = \text{"classify as unsafe"}$
 $\pi^*({b^v = 8, a^v = 6}) = \text{"classify as safe"}$
 $\pi^*({b^v = 8, a^v = i}) = \text{"classify as unsafe"}$ for all $i \in \{1, 2, 3, 4, 5, 7, 8, 9, 10, 11\}$

The above is a mathematical description of the optimal sequential testing policy.
540 In simple English, the first four lines in the above scheme can be interpreted
as follows: Start by doing the b^v test. If this test gives you values 1,2,3,9,10
or 11 then classify your chemical as unsafe, otherwise do the a^v test. If the b^v
test gave value 4, then if the a^v test gives value 6 then classify your chemical
as safe, but otherwise classify it as unsafe. The remaining statements in the
545 classification scheme are interpreted similarly.

The details of this example are not as important as the general picture which
emerges: even in this very simple scenario, the optimal policy is a set of adaptive
tests where the outcome of the earlier tests dictates what subsequent tests (if
550 any) are needed. For an adaptive testing policy to emerge from such a limited
range of possibilities, and for it to have such clear dependence on relative costs,
is evidence that adaptive testing policies may be expected to be optimal in more
realistic scenarios.

3.2. Example # 2: The case of dependent evidence

555 As in the previous example, we first carry out the two cheapest (and probably
therefore inaccurate) in silico tests, one for each parameter. Now suppose we
have no more b^* - related tests available, but have two more a^* - related in vitro
tests at our disposal, denoted a^{v1} and a^{v2} , respectively. We will assume these
two tests have different precision, and that both are more precise than their in
560 silico counterpart. Mathematically, we capture this by imposing that:

$$|a^{v1} - a^*| \leq 2 \text{ and } |a^{v2} - a^*| \leq 1 \quad (29)$$

Using the ideas from Section 3.1 we deduce that the corresponding Markov Decision Problem will have a total of 79 states. Detailed description of these and the corresponding transition probabilities can be found in the Appendix.

565 As in the previous example, it is possible to manipulate attribute and misclassification costs to obtain a range of optimal policies. Examples of these are summarized below via corresponding Ω vectors. We begin by looking at the cost structure given by:

$$\Omega = (200, 100, 80, 50) \tag{30}$$

570 In this case the optimal policy is to classify the chemical as unsafe without doing any test.

Secondly, we consider the following costs:

$$\Omega = (100, 200, 80, 50) \tag{31}$$

575 Then the optimal policy is to immediately classify the chemical as safe without doing any test.

Finally we look at the case capture via the following more realistic cost structure:

$$\Omega = (20000, 10000, 80, 50) \tag{32}$$

The optimal policy turns out to be a complex mixture, interested reader can find the details in the Appendix.

580

We finish this section by illustrating another useful feature of our model. Observe that if $a^{v1} = 1$ then (24) tells us that a^* is an element of $\{-1, 0, 1, 2, 3\}$. However, we actually know for sure that a^* is an element of $\{3, 4, 5, 6, 7, 8, 9\}$. Putting these two pieces of information together we deduce that $a^* = 3$, i.e.

585 we have thus identified a^* with absolute certainty. Therefore, common sense tells us that it should not be optimal to do any further a^* -related tests after that since we would be simply spending money for learning nothing we already know. Indeed, the above tells us that $\pi^*({a^{v^{t_1}} = 1}) = \text{"classify as unsafe"}$. This (and other similar situations) are not coincidental. Interested reader can
590 find the formal proof of this fact in the Appendix.

Discussion and Conclusion

This paper provides a mathematically rigorous framework for the computation of optimal sequential testing strategies for chemical risk classification. Its novelty lies in the fact that it does so by taking into account various important
595 and realistic ingredients. These include population model, errors in laboratory measurements and the corresponding cost structure. Mathematically, the heart of the argument comes from Machine Learning, more specifically the rich theory of Markov Decision Problems. Optimal Sequential Policies were computed in two simple but highly illustrative and easily generalisable examples, namely
600 tests that are independent and those whose outcomes are deeply correlated by the common truth they are revealing with possibly different precision. Indeed, in case of more parameters the corresponding Markov Decision Problem shall simply have more states, all other aspects will be identical.

605 The key take-home message for toxicological testing is that it is possible, and indeed necessary, to search for optimal testing strategies. Markov Decision Problems provide a mathematical framework with which to achieve this. We show that the strategies which emerge must explicitly take into account the costs both of testing and of misclassification. Moreover, the optimal testing
610 policies are typically adaptive, where the outcome of any given test influences the decision as to which test (or classification) to apply next. This may be interpreted as a mathematical avocation of the recent interest in adaptive testing strategies from both industry and regulatory bodies.

References

- 615 Alagoz O, Maillart LM, Schaefer AJ, Roberts MS. Choosing among living-donor and cadaveric livers. *Management Science* 2007;53(11):1702–15. URL: <http://dx.doi.org/10.1287/mnsc.1070.0726>. doi:10.1287/mnsc.1070.0726. arXiv:<http://dx.doi.org/10.1287/mnsc.1070.0726>.
- Alterovitz R, Branicky M, Goldberg K. Motion planning under uncertainty
620 for image-guided medical needle steering. *International Journal of Robotics Research* 2008;27(11-12):1361–74. doi:10.1177/0278364908097661.
- Bayer-Zubek . Learning cost-sensitive diagnostic policies from data. Ph.D. thesis; Oregon State University; 2004.
- Bellman R. *Dynamic Programming*. 1st ed. Princeton, NJ, USA: Princeton
625 University Press, 1957.
- Chang M. *Monte Carlo Simulation for the Pharmaceutical Industry: Concepts, Algorithms, and Case Studies*. 1st ed. Boca Raton, FL, USA: CRC Press, Inc., 2010.
- Currey JD, Pitchford JW, Baxter PD. Variability of the mechanical properties
630 of bone, and its evolutionary consequences. *Journal of The Royal Society Interface* 2007;4(12):127–35. URL: <http://rsif.royalsocietypublishing.org/content/4/12/127>. doi:10.1098/rsif.2006.0166. arXiv:<http://rsif.royalsocietypublishing.org/content/4/12/127.full.pdf>.
- Embry MR, Bachman AN, Bell DR, Boobis AR, Cohen SM, Dellarco
635 M, Dewhurst IC, Doerrer NG, Hines RN, Moretto A, Pastoor TP, Phillips RD, Rowlands JC, Tanir JY, Wolf DC, Doe JE. Risk assessment in the 21st century: Roadmap and matrix. *Critical Reviews in Toxicology* 2014;44(sup3):6–16. URL: <http://dx.doi.org/10.3109/10408444.2014.931924>. doi:10.3109/10408444.2014.931924.
640 arXiv:<http://dx.doi.org/10.3109/10408444.2014.931924>.

- Fakih S. A learning approach to obtain efficient testing strategies in medical diagnosis. *IEEE Transactions on Information Technology in Biomedicine* 2006;:220–8.
- Gabbert S, van Ierland EC. Cost-effectiveness analysis of chemical testing for decision-support: How to include animal welfare? *Human and Ecological Risk Assessment: An International Journal* 2010;16(3):603–20. URL: <http://dx.doi.org/10.1080/10807031003788840>. doi:10.1080/10807031003788840. arXiv:<http://dx.doi.org/10.1080/10807031003788840>.
- Gabbert S, Weikard HP. A theory of chemicals regulation and testing. *Natural Resources Forum* 2010;34(2):155–64. URL: <http://dx.doi.org/10.1111/j.1477-8947.2010.01300.x>. doi:10.1111/j.1477-8947.2010.01300.x.
- Gabbert S, Weikard HP. Sequential testing of chemicals when costs matter: A value of information approach. *Human and Ecological Risk Assessment: An International Journal* 2013;19(4):1067–88. URL: <http://dx.doi.org/10.1080/10807039.2012.685810>. doi:10.1080/10807039.2012.685810. arXiv:<http://dx.doi.org/10.1080/10807039.2012.685810>.
- Jaworska J, Dancik Y, Kern P, Gerberick F, Natsch A. Bayesian integrated testing strategy to assess skin sensitization potency: from theory to practice. *Journal of Applied Toxicology* 2013;33(11):1353–64. URL: <http://dx.doi.org/10.1002/jat.2869>. doi:10.1002/jat.2869.
- Jaworska JS, Natsch A, Ryan C, Strickland J, Ashikaga T, Miyazawa M. Bayesian integrated testing strategy (its) for skin sensitization potency assessment: a decision support system for quantitative weight of evidence and adaptive testing strategy. *Archives of Toxicology* 2015;89(12):2355–83. URL: <http://dx.doi.org/10.1007/s00204-015-1634-2>. doi:10.1007/s00204-015-1634-2.
- Jaworska J Gabbert S AT. Towards optimization of chemical testing under reach: A bayesian network approach to integrated testing strategies. *Regu-*

- latory Toxicology and Pharmacology 2010;57:157–67. doi:10.1016/j.yrtph.
2010.02.003.
- 670 Korthikanti VA, Viswanathan M, Agha G, Kwon Y. Reasoning about mdps
as transformers of probability distributions. In: 2010 Seventh International
Conference on the Quantitative Evaluation of Systems. 2010. p. 199–208.
doi:10.1109/QEST.2010.35.
- 675 Kurt M, Denton B, Schaefer A, Shah N, Smith S. The structure of optimal
statin initiation policies for patients with type 2 diabetes. IIE Transactions
on Healthcare Systems Engineering 2011;1(1):49–65. doi:10.1080/19488300.
2010.550180.
- Norlen H Worth A GS. A tutorial for analysing the cost-effectiveness of alter-
680 native methods for assessing chemical toxicity: the case of acute oral toxicity
prediction. Alternatives to laboratory animals 2014;42(2):115–27.
- Nunes LGN, de Carvalho SV, Rodrigues RdCM. Markov decision process ap-
plied to the control of hospital elective admissions. Artificial Intelligence in
Medicine 2017;47(2):159–71. URL: [http://dx.doi.org/10.1016/j.artmed.](http://dx.doi.org/10.1016/j.artmed.2009.07.003)
685 2009.07.003. doi:10.1016/j.artmed.2009.07.003.
- Pastoor TP, Bachman AN, Bell DR, Cohen SM, Dellarco M, De-
whurst IC, Doe JE, Doerrer NG, Embry MR, Hines RN, Moretto
A, Phillips RD, Rowlands JC, Tanir JY, Wolf DC, Boobis AR. A
21st century roadmap for human health risk assessment. Critical Re-
690 views in Toxicology 2014;44(sup3):1–5. URL: [http://dx.doi.org/](http://dx.doi.org/10.3109/10408444.2014.931923)
10.3109/10408444.2014.931923. doi:10.3109/10408444.2014.931923.
arXiv:<http://dx.doi.org/10.3109/10408444.2014.931923>.
- Program NT. The Murine Local Lymph Node Assay: A test method for As-
sessment of Allergic Contact Dermatitis Potential of chemicals/ compounds.
695 Princeton University Press, 1999.

- Shechter SM, Bailey MD, Schaefer AJ, Roberts MS. The optimal time to initiate hiv therapy under ordered health states. *Operations Research* 2008;56(1):20–33. URL: <http://dx.doi.org/10.1287/opre.1070.0480>. doi:10.1287/opre.1070.0480. arXiv:<http://dx.doi.org/10.1287/opre.1070.0480>.
- 700 Sloan TW. Safety-cost trade-offs in medical device reuse: a markov decision process model. *Health Care Management Science* 2007;10(1):81–93. URL: <http://dx.doi.org/10.1007/s10729-006-9007-2>. doi:10.1007/s10729-006-9007-2.
- Thomas RS, Black MB, Li L, Healy E, Chu TM, Bao W, Andersen ME,
705 Wolfinger RD. A comprehensive statistical analysis of predicting in vivo hazard using high-throughput in vitro screening. *Toxicological Sciences* 2012;128(2):398–417. URL: <http://toxsci.oxfordjournals.org/content/128/2/398.abstract>. doi:10.1093/toxsci/kfs159. arXiv:<http://toxsci.oxfordjournals.org/content/128/2/398.full.pdf+html>.
- 710 van der Veen JW, Rorije E, Emter R, Natsch A, van Loveren H, Ezen-
dam J. Evaluating the performance of integrated approaches for hazard identification of skin sensitizing chemicals. *Regulatory Toxicology and Pharmacology* 2014;69(3):371–9. URL: <http://www.sciencedirect.com/science/article/pii/S0273230014000816>. doi:<http://dx.doi.org/10.1016/j.yrtph.2014.04.018>.
- 715