Title Page

Virtual Experiments Enable Exploring and Challenging Explanatory Mechanisms of Immune-Mediated P450 Down-Regulation

Authors: Brenden K. Petersen[1], Glen E. P. Ropella[2], and C. Anthony Hunt[1,3*]

[1] Joint Graduate Group in Bioengineering, University of California, Berkeley, California, and University of California, San Francisco, California, United States of America

[2] Tempus Dictum, Inc., Portland, Oregon, United States of America

[3] Department of Bioengineering and Therapeutic Sciences, University of California, San Francisco, California, United States of America

* Corresponding author

E-mail: a.hunt@ucsf.edu (CAH)

Abbreviations: acetaminophen (APAP), antipyrine (ANT), chlorzoxazone (CZN), cytochrome P450 (P450), in silico hepatocyte culture (ISHC), in silico liver (ISL), interleukin (IL), iterative refinement protocol (IRP), lipopolysachharide (LPS), standard uniform distribution ($U[0,1]$), systemic clearance (CL_{sys}), tumor necrosis factor-α (TNF-α)

Abstract

Hepatic cytochrome P450 levels are down-regulated during inflammatory disease states, which can cause changes in downstream drug metabolism and hepatotoxicity. Long-term, we seek sufficient new insight into P450-regulating mechanisms to correctly anticipate how an individual's P450 expressions will respond when health and/or therapeutic interventions change. To date, improving explanatory mechanistic insight relies on knowledge gleaned from in vitro, in vivo, and clinical experiments augmented by case reports. We are working to improve that reality by developing means to undertake scientifically useful virtual experiments. So doing requires translating an accepted theory of immune system influence on P450 regulation into a computational model, and then challenging the model via in silico experiments. We build upon two existing agent-based models—an in silico hepatocyte culture and an in silico liver—capable of exploring and challenging concrete mechanistic hypotheses. We instantiate an in silico version of this hypothesis: in response to lipopolysaccharide, Kupffer cells down-regulate hepatic P450 levels via inflammatory cytokines, thus leading to a reduction in metabolic capacity. We achieve multiple in vitro and in vivo validation targets gathered from five wet-lab experiments, including a lipopolysaccharide-cytokine dose-response curve, time-course P450 down-regulation, and changes in several different measures of drug clearance spanning three drugs: acetaminophen, antipyrine, and chlorzoxazone. Along the way to achieving validation targets, various aspects of each model are falsified and subsequently refined. This iterative process of falsification-refinement-validation leads to biomimetic yet parsimonious mechanisms, which can provide explanatory insight into how, where, and when various features are generated. We argue that as models such as these are incrementally improved through multiple rounds of mechanistic falsification and validation, we will generate virtual systems that embody deeper credible, actionable, explanatory insight into immune system-drug metabolism interactions within individuals.

Introduction

Hepatic cytochrome P450 (P450) is the major family of drug-metabolizing enzymes in the liver. Changes in P450 levels are common among many disease states, giving rise to the concern that a patient may experience an imbalance in drug exposure when the disease alters P450 levels and downstream drug metabolism. Though a small subset of P450s are induced by inflammation, most inflammatory states down-regulate hepatic P450, reducing drug clearance and elevating plasma drug levels, thus increasing the risk of adverse effects—especially for low therapeutic index drugs [1,2]. P450 down-regulation can also protect against toxicity caused by reactive metabolites [2,3]. For example, pretreatment with an inflammatory stimulus protects against acetaminophen-induced hepatotoxicity [4].

Inflammatory-induced P450 down-regulation is mediated by proinflammatory cytokines, including interleukin (IL)-1β, IL-6, and tumor necrosis factor-α (TNF-α), that specifically regulate different yet overlapping subsets of P450s in both humans and rats [5,6]. Many of these cytokines are derived from Kupffer cells. While some cytokines down-regulate P450 in primary hepatocytes cultures, others are dependent upon the presence of Kupffer cells [7]. Kupffer cells can be activated by bacterial endotoxin (lipopolysachharide, LPS). An LPS stimulus causes Kupffer cells to release proinflammatory cytokines, triggering P450 down-regulation. For more information, we refer the reader to four reviews on immune-mediated P450 down-regulation [1-3,8].

Long-term, we seek sufficient new insight into P450-regulating mechanisms to correctly anticipate how an individual's P450 expressions will respond when health and/or therapeutic interventions change. To date, improving explanatory mechanistic insight relies on knowledge gleaned from in vitro, in vivo, and clinical experiments augmented by case reports. We are working to improve that reality by developing means to undertake scientifically useful virtual experiments [9,10]. To be scientifically useful, the computational models employed must demonstrate credibility, in part by meeting

demanding representational requirements. For example, not only must the simulated phenomena be quantitatively similar to available wet-lab data, but the software mechanisms—the actual events occurring during execution—should also be demonstrably biomimetic. Making key aspects of both model and experiment increasingly analogous to past or future real-world counterparts further enhances credibility. Advances in agent-based modeling and simulation (M&S) methods have now made it feasible to begin achieving such requirements [11]. We report further progress.

A prerequisite for achieving our objective is to translate a currently accepted theory of immune system influence on P450 regulation (such as that cited above) into a computational model, and then challenge the model via in silico experiments to generate phenomena that are measurably similar to preselected data reported in the literature. An agent-based model's mechanisms will be a concretized hypothesis: these components interacting in these spaces, following these rules, under these constraints will, upon execution, produce material system changes, which when measured will be within (for example) ± 10% of the corresponding wet-lab values. When the initial hypothesis fails—and it almost always does—we posit explanations, revise both hypothesis and model, and repeat the challenge. In this report, we begin with two existing agent-based models—an in silico hepatocyte culture (ISHC) [12] and an in silico liver (ISL) [13,14]—that have already achieved many validation targets related to drug metabolism and hepatotoxicity. We then repurpose both models to support the additional use case of exploring mechanisms related to immune system involvement in the liver. Specifically, we instantiate and challenge the following mechanistic hypothesis in the ISHC and ISL: in response to LPS, Kupffer cells down-regulate hepatic P450 levels via inflammatory cytokines, thus leading to a reduction in metabolic capacity.

Knowledge about interactions among hepatic P450-regulating mechanisms and immune system components comes from both in vitro and in vivo experiments. An essential, demanding requirement herein is thus that the same mechanism components must be utilized in models simulating in vitro and in vivo environments. We report quantitative validation evidence supporting an in vitro (ISHC) and in vivo model (ISL). By employing modularization and integration techniques [12], we enable reuse of

components and mechanisms between ISHC and ISL. We conceptually deconstruct the above hypothesis into three "stages" and achieve degrees of validation for each stage: 1) Kupffer cells produce cytokines upon LPS stimulus, 2) cytokines down-regulate hepatic P450 levels, and 3) P450 down-regulation reduces drug clearance. We achieve multiple in vitro and in vivo validation targets gathered from five wet-lab experiments, including a LPS-cytokine dose-response curve, time-course P450 down-regulation, and changes in several different measures of drug clearance spanning three drugs: acetaminophen (APAP), antipyrine (ANT), and chlorzoxazone (CZN). During the validation process, we falsify several mechanistic details and other ISHC and ISL components, to which we respond by iteratively refining model aspects until validation targets are achieved. This iterative process of falsification-refinement-validation ensures that model components are increasingly biomimetic yet parsimonious. Such models are perpetual works in progress. We argue that as these models are incrementally improved through multiple future rounds of mechanistic challenge and validation against an expanding set of measured attributes, we will generate virtual systems that embody deeper credible, actionable, explanatory insight into immune system-drug metabolism interactions within individuals.

Methods

To avoid ambiguity between in silico components and their referent biological counterpart, we use small caps when referring to the former, e.g. HEPATOCYTE. Parameter names are italicized. Analogs were written in Java, utilizing the MASON multi-agent simulation toolkit [15]. In silico experiments were run using 2 or 16 node virtual machines on Google Compute Engine, running 64-bit Debian 7. For longer simulations, Monte Carlo trials were run in parallel.

Synthetic and agent-based modeling

Exploring the causal mechanisms underlying P450 down-regulation can be facilitated using synthetic M&S methods. Synthetic M&S is a developing method that fundamentally differs from conventional equation-based models (i.e. physiologically based pharmacokinetic/pharmacodynamic models) in several ways. In synthetic M&S, autonomous software objects representing components such as drugs, enzymes, cells, and tissues are plugged together to form a coherent whole. The resulting multi-scale model is called a *biomimetic analog*. Analog components have several mechanisms—composed of rules, equations, and/or other operating principles—that specify how to interact with other components. At the software level, analog components are dynamic data structures, containing state information that changes as the simulation evolves; analog mechanisms are sets of governing logic that manipulate component state information.

An analog is an experimental apparatus that can be used to challenge mechanistic hypotheses about referent phenomena. When an analog is executed, its components and mechanisms are instantiated—represented by a concrete instance—in silico [16]. The analog produces phenomena, some of which are intended to mimic those of its referent system. If analog phenomena are acceptably similar to wet-lab and/or clinical validation data, we support the following hypothesis: events that transpired in silico may have biological counterparts. Thus, a validated analog stands as a challengeable theory about mechanistic events that may have occurred in the referent system. Alternative analog mechanisms can be

tested in parallel. As we continue to refine and expand surviving analog components and mechanisms to mimic an ever-increasing set of validation data, we grow increasingly confident that the analog behaves analogously to the referent under specified conditions.

The main objective for synthetic M&S is to build better working hypotheses about the mechanisms of interest. Thus, an important requirement is that analogs are suitable for virtual experimentation and, therefore, hypothesis testing. An experiment on an analog is an in silico experiment, precisely analogous to a wet-lab or clinical experiment. As such, analogs include components and mechanisms that map to concrete, relevant aspects of a wet-lab or clinical experiment, including biological components (e.g. hepatocytes), wet-lab and/or clinical environments (e.g. in vitro, in vivo), experimental procedures (e.g. intravenous drug injection), and measurements (e.g. plasma concentration profile). Thus, a virtual experiment using an analog such as the ISL is not simply a model of (say) "drug metabolism" or some other particular biological process. Rather, an instantiated ISL is a concrete software analog of (say) a whole rat intravenously injected with 20 mg/kg acetaminophen and measured at 15 min intervals via external jugular cannulation. An advantage of such virtual experimentation is that variations of a single analog can be used to mimic a wide range of experimental protocols, treatments, disease states, and measurements [10]. When new use cases require including new aspects of the referent experiment (e.g. an alternative hypothesized mechanism, a new in silico measurement, a different experimental protocol, or an altered disease state), it is straightforward to add that level of detail.

There are several differences between synthetic M&S and traditional inductive methods, which arise largely from differences in model use case and goals. Whereas synthetic M&S focus on challenging concrete mechanisms, inductive approaches (e.g. equation-based, continuous mathematics models) typically employ equations that describe patterns in data and are used to make precise predictions about future patterns in data. For these uses, conventional methods are unsurpassed. Conventional models can also be used to challenge mechanistic hypotheses (for example, by changing parameters that map to different routes of administration); however, it is more difficult to challenge a series of alternative mechanisms, possibly at different levels of granularity. While it is straightforward for synthetic analogs to

change mechanistic detail or challenge multiple mechanisms in parallel, conventional equation-based models often require significant model re-engineering—or a completely new model—when adding new variables, changing granularity, or switching use cases. Validation targets herein span multiple wet-lab platforms (in vitro and in vivo) and measurements types (e.g. three different measures of drug clearance); thus, synthetic approaches are more appropriate for purposes herein. Other key features of the similarities and differences between synthetic M&S and conventional inductive models have been detailed elsewhere [11,16,17].

Note that it would be straightforward to employ traditional modeling approaches (e.g. pharmacokinetic/pharmacodynamic modeling) to describe the biological processes involved in the immune-mediated P450 down-regulation pathway. Given sufficient P450 isozyme-specific validation data, such models may be useful, for example, in precisely predicting levels of P450 down-regulation following inflammatory stimulus. However, the ability to explore, challenge, and refine mechanistic hypotheses related to P450 down-regulation is better suited for synthetic analogs. Further, other pharmacologically relevant phenomena are difficult or problematic to describe using traditional multi-compartment models—e.g. liver zonation (i.e. lobule location-dependent effects) or leukocyte recruitment and localization—but are straightforward to reproduce using spatially explicit synthetic analogs. Since long-term objectives include exploring mechanisms of such attributes, synthetic analogs are the more appropriate choice.

Agent-based modeling is one model type that can be used to achieve the above synthetic M&S goals. An agent-based model contains agents, which are software objects capable of scheduling their own events [11]. Each agent senses, interacts with, and is a part of a virtual environment. They each follow a set of governing rules or operating principles, the logic of which is specified by the modeler. In biological M&S, agents typically map to (represent) biological components (e.g. cells), and their operating principles map to cellular interactions and processes. We refer the reader to [11] for an in-depth review of the role of synthetic M&S utilizing agent-based modeling in pharmaceutical research.

Iterative refinement protocol

A core principle of biomimetic M&S is to develop software models that provide increasingly credible mechanistic explanations of their referent complex biological processes. Such increases require the strategic use of parsimony. Abstract models will be falsified when challenged against specific validation data. When a model fails, the appropriate response is to refine it: either by increasing complexity (e.g. by adding parameters or switching to finer-grained components) or testing alternative mechanisms and/or structures before repeating the challenge. In support of this core principle, our model development process is governed by an iterative refinement protocol (IRP), a scientific method for falsifying, refining, and validating multi-scale biomimetic analogs. Described in Fig 1A, the IRP focuses on model falsification and subsequent refinement, following a strict parsimony guideline. The goal of stepping through an IRP cycle is to formulate a mechanistic hypothesis by meeting a set of validation targets. When *prespecified* similarity criteria are attained for a targeted attribute, we have achieved a validation target. Similarity criteria can range from qualitative (e.g. event *X* occurs before event *Y*) to quantitative (e.g. in silico points fall within ± 1 standard deviation of the corresponding wet-lab value). An analog mechanism is falsified when it cannot achieve a given validation target, either because the target phenomenon cannot be generated or we fail to find parameter set values that satisfy all similarity criteria.

Fig 1. Analog methods and structure. A. An iterative protocol for refining biomimetic analogs. **B.** Key features of ISHC structure. The ISHC contains two grids: CELL SPACE and MEDIA SPACE (only portions of each grid are shown). SOLUTES can move laterally within a grid, or between CELL SPACE and MEDIA SPACE, subject to the parameters *pExitMedia* and *pExitCell*. Select HEPATOCYTE and KUPFFER CELL components are shown. **C.** Key features of ISL structure. There are two major components: BODY and LOBULE. SOLUTES are injected into BODY, where they distribute to the PORTAL VEIN (PV) of the LOBULE. SOLUTES percolate through a network of SINUSOID SEGMENTS (SS) toward the CENTRAL VEIN (CV), from

which they return to BODY. SOLUTES can also move radially within a SINUSOID SEGMENT through various SPACES. Select HEPATOCYTE and KUPFFER CELL components are shown.

When a mechanism is falsified, we specify a revision—a new hypothesis—which usually involves adding and/or revising analog components and/or mechanisms, perhaps at a finer level of granularity. To be scientifically meaningful, refinements should be parsimonious: we seek the simplest analog that still ("just barely") achieves validation targets. If an analog mechanism is too simple for a given set of validation targets, it will be falsified, and iterative refinement will lead to a new mechanism. However, for an analog that is too complex (over-mechanized, analogous to an equation-based model being overparameterized), we waste computational resources, and it becomes difficult to identify where and why a future mechanism is falsified, rendering reengineering problematic.

ISL and ISHC use cases

Both the ISL and ISHC are biomimetic analogs used to challenge mechanistic hypotheses related to drug metabolism and hepatotoxicity. As primarily exploratory devices, their focus shifts away from precise prediction of pharmacological values—for which traditional continuous mathematics and/or statistical methods are unsurpassed—and instead toward developing flexible methods to simulate mechanistic scenarios. It is easy to add mechanistic details that improve apparent realism of an ABM (e.g. [18]). However, to retain scientific usefulness, we strive to keep analog mechanisms parsimonious.

The ISL simulates drug clearance and hepatotoxicity experiments in an in vivo setting. It is a highly flexible platform that can mimic many experimental use cases. One use case configuration maps to a portion of an in situ isolated, perfused rat liver, in which simulations mimic the multiple indicator dilution method to measure hepatic outflow profile [19]. This iteration of the ISL includes a whole-rat configuration that can simulate either oral or intravenous drug administration. Measurements on the ISL range from measures of metabolism (e.g. hepatic outflow profile, extraction ratio, and intrinsic clearance) to necrosis (e.g. the time and location of cell death events). The current use case expands the ISL to

include analogs of immune system components and mechanisms; experiments are conducted to test the entire P450 down-regulation pathway from LPS stimulus to decreased clearance. In particular, ISL simulations here mimic several wet-lab validation experiments in which rats are pretreated with LPS before administering drug (see Results for experiment details) [20-22]. Measurements include changes in drug clearance and P450 amounts.

The ISHC simulates drug clearance and hepatotoxicity experiments in an in vitro setting. A simulation maps to a portion of a monolayer culture of isolated rat hepatocytes and/or Kupffer cells. While much simpler than the ISL, the ISHC is useful for mimicking small parts of the immune-mediated P450 down-regulation pathway. In particular, the current ISHC use cases mimic wet-lab experiments in which cultured hepatocytes produce cytokine in response to LPS [23] and cultured Kupffer cells down-regulate P450 in response to cytokine (see Results for experiment details) [24]. Measurements include a LPS-cytokine dose-response curve and time-course enzyme levels.

Selection of drugs and similarity criteria

We selected validation data for three drugs (APAP, ANT, and CZN) having different physiochemical properties and spanning a ~4-fold range in half-life in rats. Additional drugs can be added to the validation dataset to further improve model confidence; we selected three to demonstrate P450 down-regulation mechanisms. Indeed, earlier versions of the ISHC and ISL have been used for different use cases spanning additional drugs: see [12,25]. When introducing new analog drug objects into the simulation, we would first refer to literature reports of fraction metabolized. For the three chosen drugs, fraction metabolized in rats is nearly 100% [26-28]. We would next consider what fraction of metabolic clearance is due to P450 isoforms—where our attention is focused. For ANT and CZN, that is nearly 100% [27,28]; it is only a minor pathway (though important to toxicity) for APAP, dominated by cytochromes P450 2E1 and 1A2 [26,29]. Given fraction metabolized via specific enzymes, we refine the ISL and ISHC to include enzyme- and cytokine-specific functions (see “Mechanism granularity and flexibility”).

The choice of similarity criteria is governed by model use case. For example, a model intended to be used for precise prediction requires more stringent similarity criteria than a model used to explore explanatory mechanisms. When possible, similarity criteria were chosen to reflect variability in wet-lab measurements. Thus, a good starting point for moderately stringent similarity criteria is that in silico values fall within ± 1 standard deviation of the corresponding wet-lab values. When standard deviation is not given or cannot be determined from the validation data, an alternative is to specify an arbitrary percentage range. In these cases, we stipulated a range of ± 25%, which we found to be reflective of most related, reported standard deviation values. As an exception, we chose a range of ± 10% for ANT half-life to better reflect the smaller standard deviations of half-lives measures for other drugs. For validation targets that include many values (e.g. time-course data), we stipulate that at least 50% of in silico values must fall within the prespecified range. So doing prevents present but statistically insignificant abnormalities in a plot's microstructure from prohibiting model development in early stages of validation and reduces the risk of overfitting to one particular dataset. For an example of such irregularities in microstructure, validation data from [20] at 80 min shows a statistically insignificant nonlinear dip in the semi-logarithmic drug disappearance curve for APAP, which is not found in similar studies (e.g. [30,31]).

ISL and ISHC components

Upon execution, time advances through simulation cycles. Agents are scheduled to execute once each simulation cycle, with the exception of some events that are separately scheduled. Components common to both the ISL and ISHC include SOLUTES, ENZYMES, and CELLS. SOLUTES are mobile objects that map to a group of small molecules. They percolate through SPACES, influenced by various flow parameters (see Table 1). SOLUTES can have any number of properties that are specified offline as part of the parameter list (see Table 2). For example, only "bindable SOLUTES" (SOLUTES for which the Boolean parameter *bindable* is true) may bind to ENZYMES. Each SOLUTE is assigned a type that controls how other objects may interact with it. Herein, SOLUTE types include LPS, CYTOKINE, DRUG (APAP, ANT, or CZN), and METABOLITE (APAP-METABOLITE, ANT-METABOLITE, or CZN-METABOLITE); only the three

DRUG types are bindable. For current use cases, only a single CYTOKINE type was necessary; thus, CYTOKINE maps to the set of all cytokines that may cause hepatic P450 down-regulation. However, when required in future use cases, CYTOKINE may be replaced with objects that map to specific cytokines (e.g. INTERLEUKIN-1), each with unique parameter values. There are no restrictions on the type or amount of information that can be "attached" to each SOLUTE; additional properties are added based on use case and as additional validation targets are achieved.

292 **Table 1. ISL and ISHC parameters descriptions and values for validating experiments.**

Parameter name	Type/Range	ISHC value(s)	ISL value(s)	Description
Simulation control parameters				
cycleLimit	positive integer	{2880; 1440}	variable	Number of simulation cycles after which simulation stops.
monteCarloTrials	positive integer	16	16	Number of Monte Carlo trials to execute.
Dosing parameters				
doseTime (SOLUTE type)	positive integer	{1 (LPS); 1 (CYTOKINE)}	{1 (DRUG); 1 (LPS), 86401 (DRUG)}	Simulation cycle at which to administer SOLUTE dose. Multiple doses are separated by commas.
dosage	positive integer	{variable; 2000}	{62500; 125000}	Number of SOLUTE objects to administer each dose.
BINDING HANDLER parameters				
pBind	[0.0, 1.0]	N/A	0.25	Base probability for an unbound SOLUTE to bind to an unbound ENZYME.
bindCycles	positive integer	N/A	10	Number of a simulation cycles a bound SOLUTE remains bound to an ENZYME.
bindExponent	positive, real	N/A	see Table 3	Exponent that controls the degree to which increasing the number of bound ENZYMES decreases the probability of a binding event.
bindable	Boolean	see Table 2	see Table 2	If TRUE, this SOLUTE may bind to an ENZYME.
acceptedSolutes	list	see Table 3	see Table 3	List of SOLUTE types that can bind to this ENZYME.
METABOLISM HANDLER parameters				
pMetabolize	[0.0, 1.0]	N/A	see Table 2	Probability that a bound SOLUTE is metabolized by its bound ENZYME.
metabolic	Boolean	see Table 3	see Table 3	If TRUE, this ENZYME may metabolize bound SOLUTES.

metabolicProduct	list	see Table 2	see Table 2	List of METABOLITES (if any) to produce upon metabolism.
INFLAMMATION HANDLER parameters				
inflammatoryThreshold	non-negative integer	{1; N/A}	3	Threshold number of INFLAMMATORY STIMULI above which a CYTOKINE may be produced.
cytokineThreshold	non-negative integer	{2; N/A}	2	Threshold number of CYTOKINES above which no more CYTOKINES may be produced.
cytokineExponent	positive, real	{3.0; N/A}	3.0	Exponent that controls the degree to which increasing INFLAMMATORY STIMULI increases *pCytokine*.
inflammatory	Boolean	see Table 2	see Table 2	If TRUE, this SOLUTE may cause KUPFFER CELLS to produce CYTOKINES.
DOWN-REGULATION HANDLER parameters				
pRemove	[0.0, 1.0]	{N/A; 0.05}	see Table 3	Probability that a CYTOKINE removes an ENZYME.
delay	non-negative integer	{N/A; 30}	600	Number of simulation cycles to delay before an ENZYME scheduled for removal is actually removed.
pReplenish	[0.0, 1.0]	{N/A; 0.007}	0.0001	Probability that an ENZYME is created when there are no CYTOKINES, the removal queue is empty, and there are fewer than the starting number of ENZYME.
downRegulated	Boolean	see Table 3	see Table 3	If TRUE, this ENZYME may be down-regulated by CYTOKINES.
DEGRADATION HANDLER parameters				
pDegrade	[0.0, 1.0]	see Table 2	see Table 2	Probability that an unbound SOLUTE is degraded.
CELL parameters				
enzymesPerCellMin	non-negative integer	4	4	The minimum number of ENZYME of each type to create in each CELL at

				the start of the simulation.
enzymesPerCellMax	non-negative integer	8	8	The maximum number of ENZYME of each type to create in each CELL at the start of the simulation.
ecDensity	[0.0, 1.0]	N/A	0.66	Probability that a grid point in ENDOTHELIAL SPACE contains an ENDOTHELIAL CELL.
kcDensity	[0.0, 1.0]	{1.0; 0.0}	0.33	Probability that a grid point in CELL SPACE (ISHC) or ENDOTHELIAL SPACE (ISL) contains a Kupffer Cell.
hepDensity	[0.0, 1.0]	{0.0; 1.0}	0.9	Probability that a grid point in CELL SPACE (ISHC) or HEPATOCYTE SPACE (ISL) contains a HEPATOCYTE.
expressingCellTypes	list	see Table 3	see Table 3	List of which CELL types contain this ENZYME type.
SOLUTE flow parameters				
sampleRatio	[0.0, 1.0]	N/A	0.00115	Base fraction of SOLUTES in BODY that are transferred to PORTAL VEIN each simulation cycle.
sampleRatioFactor	positive, real	N/A	see Table 2	The fraction of SOLUTES in BODY that are transferred to PORTAL VEIN each simulation cycle is multiplied by this factor.
$V_{d,change}$	positive, real	N/A	see Table 2	The fraction of SOLUTES in BODY that are transferred to PORTAL VEIN each simulation cycle is reduced by this factor in LPS experiments. Further, the original amount of administered SOLUTES in reduced by this factor in LPS experiments.
forwardBias	[0.0,1.0]	N/A	0.2	Weight given to forward movement of a SOLUTE object, modifying the otherwise "Brownian" motion.

flowRate	positive integer	N/A	2	Number of grid points SOLUTE in the CORE is moved forward each simulation cycle.
lateralBias	[0.0, 1.0]	N/A	0.6	Weight given to lateral movement of a SOLUTE object, modifying the otherwise "Brownian" motion.
pExitMedia	[0.0, 1.0]	see Table 2	N/A	Probability that a SOLUTE can move from MEDIA SPACE to CELL SPACE.
pExitCell	[0.0, 1.0]	see Table 2	N/A	Probability that a SOLUTE can move from CELL SPACE to MEDIA SPACE.

When ISHC parameter values differ between dose-response experiments (used to generate Fig 2A) and time-course experiments (used to generate Fig 2B), the different values are shown in brackets separated by semicolons: e.g. {dose-response value; time-course value}. Similarly, when ISL parameter values differ between control and LPS experiments, the different values are shown in brackets separated by semicolons: e.g. {control value; LPS value}. "N/A" values denote that the parameter is either not included in that simulation (e.g. an ISHC-specific parameter in an ISL simulation) or is not relevant in that simulation (e.g. METABOLISM HANDLER parameters in ISHC simulations without DRUG). If a value states "see Table 2," that parameter differs based on SOLUTE type; see Table 2 for SOLUTE-specific values. Similarly, if a value states "see Table 3," that parameter differs based on ENZYME type; see Table 3 for ENZYME-specific values. The ISL values for *cycleLimit* are variable: for control experiments, the values are 6000 (for APAP), 18000 (for ANT), and 7200 (for CZN); for LPS experiments, the values are 92400 (for APAP), 104400 (for ANT), and 93600 (for CZN). The ISHC dose-response values for *dosage* are also variable: the dose-response curve was measured at 0, 70, 700, 7000, and 700000 LPS objects.

Table 2. SOLUTE-specific parameter values for validating experiments.

SOLUTE type	*bindable*	*inflammatory*	*pMetabolize*	*metabolicProduct*	*pDegrade*	*sampleRatioFactor*	$V_{d,change}$
ISL simulations							
APAP	TRUE	FALSE	[0.35, 0.95]	APAP-METABOLITE	N/A	1.0	N/A
ANT	TRUE	FALSE	[0.35, 0.95]	ANT-METABOLITE	N/A	0.26	N/A
CZN	TRUE	FALSE	[0.35, 0.95]	CZN-METABOLITE	N/A	0.52	1.69
LPS	FALSE	TRUE	N/A	N/A	0.0005	N/A	N/A
CYTOKINE	FALSE	FALSE	N/A	N/A	0.002	N/A	N/A
APAP-METABOLITE	FALSE	FALSE	N/A	N/A	N/A	N/A	N/A
ANT-METABOLITE	FALSE	FALSE	N/A	N/A	N/A	N/A	N/A
CZN-METABOLITE	FALSE	FALSE	N/A	N/A	N/A	N/A	N/A
SOLUTE type	***bindable***	***inflammatory***	***pMetabolize***	***metabolicProduct***	***pDegrade***	***pExitCell***	***pExitMedia***
ISHC simulations							
LPS	FALSE	TRUE	N/A	N/A	N/A	1	{0.5; 0.1}
CYTOKINE	FALSE	FALSE	N/A	N/A	{0.01; 0.002}	{0.01; 0.2}	0.02

When ISHC parameter values differ between dose-response experiments (used to generate Fig 2A) and time-course experiments (used to generate Fig 2B), the different values are shown in brackets separated by semicolons: e.g. {dose-response value; time-course value}. "N/A" values denote that the parameter is not relevant in that simulation (e.g. METABOLISM HANDLER parameters in ISHC simulations without DRUG). Note ISL *pMetabolize* values are given as a range; CELLS nearest the PORTAL VEIN exhibit the minimum value, CELLS nearest the CENTRAL VEIN exhibit the maximum value, and the value is linearly interpolated for CELLS in between.

ENZYMES are objects within CELLS that can bind and metabolize SOLUTES. Note an ENZYM object named for convenience. It does not represent actual metabolic enzymes. Rather, an ENZYME to a portion of material within a cell that can influence metabolism of the SOLUTE'S counterparts wi simulation cycle. So doing is a necessary consequence of adherence to our strong parsimony guideli which includes using single object types as placeholders for what in the future may be a set of distin different objects. Like SOLUTES, ENZYMES have a number of properties, including ENZYME type. In ISHC, there is a single type of ENZYME that interacts with all bindable SOLUTES. In this iteration of ISL, there are four ENZYME types: APAP-ENZYME, ANT-ENZYME, CZN-ENZYME, and NONSPECIFIC. Different ENZYME types have type-specific properties, including a list of which SOLUTE types can b that ENZYME type (see Table 3). Different CELL types can also contain ("express") different ENZYM (expanded upon below). For simplicity, we specified ENZYME types according to the corresponding DRUG. For example, APAP-ENZYMES exclusively bind and metabolize APAP objects. NONSPECIFIC c bind all bindable SOLUTES, but cannot metabolize them. While finer-grain knowledge of which P45 isoforms metabolize APAP, ANT, and CZN in both rats and humans is available [29,32-35], as imp above, parsimony dictates that we do not include that level of granularity until validation targets car be achieved without doing so. Similarly, [21] measured levels of CYP3A2 and CYP2C11 separately however, relative changes after LPS pretreatment were extremely similar for both isozymes. Thus, including two distinct ENZYME types for ANT experiments was unnecessary.

Table 3. ENZYME-specific parameter values for validating experiments.

ENZYME type	*expressingCellTypes*	*acceptedSolutes*	*metabolic*	*downRegulated*	*pRemove*	*bindExponent*
APAP-ENZYME	HEPATOCYTE	APAP	TRUE	TRUE	0.01	1.0
ANT-ENZYME	HEPATOCYTE	ANT	TRUE	TRUE	0.025	2.0
CZN-ENZYME	HEPATOCYTE	CZN	TRUE	TRUE	0.02	1.5
NONSPECIFIC	KUPFFER CELL, ENDOTHELIAL CELL	APAP, ANT, CZN	FALSE	FALSE	N/A	1.0

"N/A" values denote that the parameter is not relevant in that simulation (e.g. METABOLISM HANDLER parameters in CELLS without metabolizing ENZYMES).

CELLS are agent objects that maintain state information and can contain ENZYMES and SOLUTES. The three CELL types used in simulations here are HEPATOCYTES, ENDOTHELIAL CELLS (ISL only), and KUPFFER CELLS. HEPATOCYTES contain enzymes. ENDOTHELIAL CELLS contain only NONSPECIFICS. KUPFFER CELLS do not contain ENZYMES. Each type of CELL contains various physiomimetic mechanism modules [10] that are executed once per simulation cycle in pseudo-random (simply random hereafter) order. When executed, mechanism modules act upon CELL contents (SOLUTES and ENZYMES) and other CELL state information (see "ISL and ISHC mechanisms" for details).

ISL and ISHC structure

Full details of ISL structure are provided elsewhere [13,14]. For this work, we plugged the cited LOBULE object into a simple BODY compartment (Fig 1C). Together, they map to a rat. BODY maps to plasma plus all other drug-accessible tissues; LOBULE maps to the rat liver. Injected SOLUTES (i.e. LPS and/or DRUG) are added to BODY during the simulation, and BODY transfers a fraction of its SOLUTES to the PORTAL VEIN of the LOBULE each simulation cycle. In silico measurements of SOLUTE are sampled from BODY. LOBULE is a directed graph—or sinusoid network—of interconnected nodes and edges. Each node is a SINUSOID SEGMENT, which maps to a portion of sinusoid. Edges map to direction of blood flow, from the portal vein to the central vein. Once transferred from BODY to PORTAL VEIN, SOLUTES percolate through SINUSOID SEGMENTS and the edges connecting them. Surviving SOLUTES reach CENTRAL VEIN and return to BODY. SINUSOID SEGMENTS contain an innermost CORE, an outermost BILE CANAL, and concentric, cylindrical SPACES. Cellular SPACES contain grid points that can contain at most one CELL. Acellular SPACES can only contain SOLUTES. Moving radially outward from CORE, SOLUTES may enter BLOOD-CELL INTERFACE (acellular), ENDOTHELIAL SPACE (contains ENDOTHELIAL CELLS and KUPFFER CELLS), SPACE OF DISSE (acellular), HEPATOCYTE SPACE (contains HEPATOCYTES), and BILE CANAL (acellular).

ISHC structure is composed of two stacked, rectangular grids (Fig 1B), each mapping to different in vitro spaces [10]. CELL SPACE maps to the monolayer of cells; each grid point contains at most one

CELL (HEPATOCYTE or KUPFFER CELL). MEDIA SPACE maps to culture media. Both SPACES may contain SOLUTES, which can move between SPACES or laterally within a SPACE. To account for the much greater volume in culture media compared to the cell monolayer, between-grid movement is asymmetrical: a SOLUTE can only move "up" from CELL SPACE to MEDIA SPACE with probability *pExitCell*; a SOLUTE moving "down" uses the parameter *pExitMedia*, which is typically much smaller than *pExitCell*. Injected SOLUTES are randomly assigned to MEDIA SPACE and CELL SPACE grid points. SOLUTE measurements are sampled from the entire system.

ISL and ISHC mechanisms

Logic governing ISL and ISHC mechanism modules is outlined below, and flowcharts of mechanism logic are illustrated in S1 Fig. There are five mechanism modules used by CELLS in both the ISL and ISHC for current use cases. Mechanism names include the suffix HANDLER to emphasize the fact that they are modules that directly act upon model components and their state information. In fact, outside of the function of the mechanism modules, SOLUTES do nothing other than percolate through SPACES, and ENZYMES do nothing other than exist within CELLS. Each mechanism is probabilistic in nature. When an event occurs with probability p, a random draw from the standard uniform distribution, $U[0,1)$, determines whether the event occurs. If the random draw is less than p, the event occurs. Event execution is handled using a scheduler that conducts the timing and ordering of events.

BINDING HANDLER maps to both specific and nonspecific binding processes in the cell. This mechanism module is present in all HEPATOCYTES and ENDOTHELIAL CELLS. When the mechanism is executed, each unbound SOLUTE inside the CELL has a chance to bind to an unbound ENZYME. With probability *pBind*, a SOLUTE binds to an unbound ENZYME. While a SOLUTE is bound, it cannot leave that CELL, bind to another ENZYME, or be removed via degradation; however, it can be metabolized (see below). While an ENZYME is bound, it cannot be removed (e.g. via down-regulation) or bind another SOLUTE. Upon binding, the SOLUTE is scheduled to be released (unbound) from the ENZYME after *bindCycles* simulation cycles.

METABOLISM HANDLER maps to metabolism via P450 enzymes; it is unique to HEPATOCYTES. When executed, there is a chance for each bound SOLUTE to be metabolized with probability *pMetabolize*. When a SOLUTE is metabolized, two events occur. 1) The ENZYME that metabolized the SOLUTE is no longer bound and is thus free to bind and metabolize again. 2) The metabolized SOLUTE is removed from the system and replaced with its appropriate METABOLITE. The SOLUTE'S state information specifies which METABOLITE type is produced. If the SOLUTE specifies multiple METABOLITE types, one is randomly selected. If the SOLUTE does not specify a METABOLITE product, it is simply removed from the system and is not replaced.

A different value of *pMetabolize* is specified for each SOLUTE type, allowing for different types of SOLUTES to be metabolized at different rates. In simulations here, a particular SOLUTE type has only one corresponding ENZYME type that can metabolize it. However, some use cases may require multiple ENZYME types that can metabolize a particular SOLUTE type. In such cases, a different value of *pMetabolize* can be specified for each pairwise combination of SOLUTE and ENZYME types. Thus, metabolism can be differentially controlled by individual SOLUTE and ENZYME types.

INFLAMMATION HANDLER maps to cytokine production in response to an inflammatory stimulus like LPS; it is unique to KUPFFER CELLS. When executed, the mechanism has a chance to produce a CYTOKINE; if so, the CYTOKINE is added to the KUPFFER CELL. To determine whether to produce a CYTOKINE, it first counts how many INFLAMMATORY STIMULI are in the KUPFFER CELL. An INFLAMMATORY STIMULUS is any SOLUTE for which the Boolean property *inflammatory* is true. The only INFLAMMATORY STIMULUS is LPS. If the number of INFLAMMATORY STIMULI exceeds the value of the parameter *inflammatoryThreshold*, there is a chance to produce a CYTOKINE. With probability *pCytokine*, a CYTOKINE is produced. The value of *pCytokine* is variable, determined by several factors:

$$pCyotkine = 1 - \exp\left(-\frac{\#\text{INFLAMMATORY STIMULI} - inflammatoryThreshold}{cytokineExponent}\right)$$

Thus, *cytokineExponent* controls the degree to which increasing stimulus causes CYTOKINE formation. Lastly, to prevent excess CYTOKINE formation in the presence of significant INFLAMMATORY STIMULUS, CYTOKINE cannot be created if there are already more than *cytokineThreshold* CYTOKINES in the CELL.

The above equation is an example of a biomimetic rule. We lack sufficient detailed knowledge to describe exactly how a Kupffer cell in a particular lobule location determines whether to produce cytokine. The equation is a placeholder for yet to be specified fine-grain mechanisms. A risk of relying on such rules is committing inscription error, the logical fallacy of assuming the conclusion and programming in (consciously or not) aspects of the result we expect to see [36]. Cognizant of inscription error, we were careful not to explicitly encode a sigmoidal dose-response into INFLAMMATION HANDLER. We chose the above equation to mimic a Poisson distribution in which the probability of at least one binding event occurring within a simulation cycle depends on the extent to which the number of inflammatory stimuli exceeds a threshold value. The ability to generate a sigmoidal dose-response curve arises from a complex interplay among multiple analog mechanisms, facilitated by flexibility in the parameters *inflammatoryThreshold* and *cytokineExponent*.

DOWN-REGULATION HANDLER maps to P450 down-regulation processes in response to a cytokine signal; it is unique to HEPATOCYTES. This mechanism is more complex and can follow one of two pathways: when executed, it has a chance to either 1) remove an ENZYME (in response to sufficient CYTOKINE) or 2) create an ENZYME (when CYTOKINE is not present). ENZYME removal maps to P450 down-regulation; ENZYME creation maps to the gradual return to basal P450 levels. Importantly, each HEPATOCYTE contains a separate instance of DOWN-REGULATION HANDLER for each down-regulatable ENZYME type (call it *T*) it contains. In the descriptions that follow, each instance operates only on ENZYMES of type *T*. The instances operate independently.

When CYTOKINE is present in the HEPATOCYTE, an ENZYME has a chance to be removed. For each CYTOKINE, the probability to do so is *pRemove*. However, the ENZYME is not removed immediately; instead, ENZYME removal is scheduled to occur after *delay* simulation cycles. Until then, the ENZYME is marked for removal but is not actually removed. The more CYTOKINES are present, the more chances

there are to schedule ENZYME removal; however, to prevent abrupt changes in the number of ENZYMES, at most one ENZYME may be scheduled for removal each simulation cycle. Further, if additional ENZYMES are scheduled for removal before the previously scheduled ENZYME is actually removed, their delay is added to the tail end of the currently remaining delay. Thus, when CYTOKINE levels are sufficiently high, there is a growing "queue" of scheduled ENZYME removal events. Including the *delay* parameter and removal queue was necessary to prevent rapid ENZYME down-regulation over the course of very few simulation cycles.

Alternatively, ENZYMES may be created when there are no more CYTOKINES in the HEPATOCYTE and after the removal queue has cleared. Further, this pathway can only occur when the number of ENZYMES is less than the number at the start of the simulation. Thus, ENZYMES can gradually return to their original (basal) levels. If the three conditions are met—no CYTOKINES, empty queue, less than basal ENZYME—an ENZYME may be created with probability *pReplenish*.

We designed DOWN-REGULATION HANDLER to be sufficiently general so that, when needed, we can include specific subtypes of CYTOKINES (or other SOLUTES) that can specifically and differentially down-regulate different sets of ENZYME types. Currently, each ENZYME type can have a different value for *pRemove*. In the future, as mechanisms for individual CYTOKINE types are teased out, *pRemove* can be specific to a particular pair of ENZYME and CYTOKINE type. In other words, parameters like *pRemove* can take on pairwise values: e.g. $pRemove_{i,j}$, where *i* is a specific ENZYME type and *j* is a specific CYTOKINE type.

The final mechanism, DEGRADATION HANDLER, maps to non-P450 degradation of compounds; it is in all CELLS. The mechanism is simple: when executed, each unbound SOLUTE can, with probability *pDegrade*, be "degraded" (removed from the system). There is no limit to the number of SOLUTES that can be degraded each simulation cycle via this mechanism. SOLUTES that do not specify a value for *pDegrade* cannot be degraded.

Mappings between ISL/ISHC mechanism parameters (like *pMetabolize*) and kinetic, thermodynamic, and/or pharmacological properties exist, but need not be one-to-one. Quantitative

mappings can be made via methods including linear regression, artificial neural network, or fuzzy clustering algorithms. Indeed, these methods have each been shown to successfully predict ISL parameter values for several drugs, which, when plugged into an ISL and tested, produce simulated hepatic outflow profiles that closely match wet-lab data [37].

ENZYME mechanism falsification and generalization

The ENZYME ontology and related binding and metabolism mechanisms were falsified by the observation that even a dramatic reduction in the number of ENZYMES does not significantly affect SOLUTE clearance. The explanation is that the probabilities of binding events—and therefore metabolism events—is neither a function of the number of ENZYMES nor the number of SOLUTES that are already bound (provided there is at least one unbound ENZYME). There are two exceptional cases: 1) the case in which all ENZYMES are bound and 2) the subset of this case in which there are zero ENZYMES. In both exceptional cases, the probability of a binding event drops to zero (S2A Fig). These cases were found to be rare or nonexistent except under extreme parameter settings. Thus, within typical parameter settings, effective binding probability is unaffected by changes in the number of ENZYME (S2C Fig). The result is that a change, such as simulating P450 down-regulation, has no significant effect on SOLUTE clearance, thereby falsifying those earlier mechanisms. Hereafter, whenever necessary to avoid ambiguity between ENZYMES used earlier and those described below, we refer to the former as generation 1 ENZYME (G1-ENZYME) and refer to the latter as generation 2 ENZYME (G2-ENZYME).

Achieving validation targets like those drawn from P450 down-regulation required a finer-grained, more generalized ENZYME ontology and mechanisms for binding and metabolism. We implemented the following refinements.

ENZYME-specific properties

Whereas all G1-ENZYMES were controlled by the same set of parameters, G2-ENZYMES are dynamic data structures that have ENZYME-specific parameters and maintain ENZYME-specific state information. Among important parameters is ENZYME type, which allows other components to recognize

and interact differently with different types of ENZYME. Each G2-ENZYME is assigned a type (e.g. "Phase 1"), which maps to some set of xenobiotic-metabolizing enzymes. This allows different model components to recognize and interact with ENZYMES differently based on type. G2-ENZYMES can also have type-specific properties. For example, each ENZYME type contains a list of which SOLUTE types can bind ENZYMES of that type. Other parameters include Booleans controlling whether that ENZYME can metabolize SOLUTE (as opposed to binding only), is induced by DRUG, is down-regulated by CYTOKINE, etc.

Each CELL contains G2-ENZYMES of each type required by the current use case. Different CELL types within a simulation can contain (or "express") a different—possibly overlapping—set of ENZYME types, which are specified offline via a parameter file (see Table 3). In ISL simulations here, HEPATOCYTES contain all four ENZYME types (APAP-ENZYME, ANT-ENZYME, CZN-ENZYME, and NONSPECIFIC), and ENDOTHELIAL CELLS contain only NONSPECIFIC.

Parameter-controlled binding probability

G2-ENZYMES can bind and metabolize SOLUTES similarly to G1-ENZYMES; however, binding need not be controlled by a single constant parameter, *pBind*. Rather, the number of unbound ENZYMES of a given type in a given CELL determines the probability of a binding event. The probability function used is controlled by a whole-model parameter, *bindingMode*. Currently, there are two binding modes: stepwise and variable (illustrated in S2 Fig). Stepwise mode is intended to recapitulate the previous G1-ENZYME mechanism behaviors, used primarily for verification purposes. Using the stepwise mode:

$$P(\text{binding event}) = \begin{cases} pBind & \#unbound > 0 \\ 0 & \#unbound = 0 \end{cases}$$

where *#unbound* is the total number of unbound ENZYMES. Thus, the binding probability is either zero (when all ENZYMES are bound) or *pBind* (when at least one ENZYME is unbound) (S2C Fig). Using the variable mode:

$$P(\text{binding event}) = pBind \cdot \left(-\frac{\#bound - total}{total_{\text{initial}}}\right)^{bindExponent}$$

where *#bound* is the total number of bound ENZYMES, *total* is the total number of ENZYMES (bound or unbound), and $total_{initial}$ is the total number of ENZYMES at the start of the simulation. Thus, the binding probability decreases as more ENZYMES bind (moving down the line) and/or as ENZYMES are removed from the system (shifting the curve downward) (S2D Fig). The parameter *bindExponent* controls the nonlinearity of this effect (S2B Fig). The probability reaches zero when all ENZYMES are bound (i.e. the number of bound ENZYMES reaches the current total number of ENZYMES). Further, the *y*-intercept (the probability of binding when there are no bound ENZYMES) is scaled by *total* normalized by $total_{initial}$. Thus, changes in the total number of ENZYMES are two-fold (S2D Fig): it changes both the probability of binding and the maximum number of SOLUTES that can be bound to the ENZYME (*x*-intercept). Because all metabolism events are preceded by a binding event, changes in binding probability will directly affect metabolism and clearance.

G2-ENZYMES are designed to be inherently tunable [38]. If a future use case requires addition of non-P450 enzymes or further delineation into (say) specific P450 isoforms, the G2-ENZYME structure can still be used. On the other hand, if a future use case does not require this level of granularity, the G1-ENZYME mechanisms can be restored by specifying only two ENZYME types (ENZYME and NONSPECIFIC) and setting the binding probability function to stepwise mode. In this way, the phenotype space of the G2-ENZYME mechanism subsumes that of the G1-ENZYMES mechanism; that is, the new mechanism is a generalization of the earlier mechanism.

Mechanism granularity and flexibility

It is well established that a particular drug may be metabolized by several enzymes, each of which may be differentially down-regulated (or perhaps up-regulated or unaffected) by different cytokines. Given that a G2-ENZYME object does not map to a particular P450 isoform, why should we compare ENZYME measurements to wet-lab data using a single, specific P450 isoform? A similar question can be asked of CYTOKINES. The answer is two-fold. Firstly, the chosen validation data is not sufficiently detailed to extract differential effects of multiple specific P450 isoforms or specific cytokines. Following

our strong parsimony guideline, analog mechanisms should not include that level of isoform- or cytokine-specific detail unless and until doing so is needed to achieve validation targets. For example, the validation data do not measure the regulation of more than one or two P450 isoforms. ([21] measures two isoforms, but they are down-regulated almost the same amount.) Thus, current use cases only require simulation and measurement of a single metabolizing ENZYME type per drug. Secondly, the current use case is to develop flexible methods for simulating immune-mediated P450 down-regulation. Future use cases may include validating against differential isoform- and/or cytokine-specific data. In these cases, the ISL and ISHC have built-in capabilities to further delineate ENZYMES and CYTOKINES using only changes in parameter values. Their differential effects on inflammation and P450 down-regulation can subsequently be specified using CYTOKINE-specific parameters (e.g. *inflammatoryThreshold*) and pairwise CYTOKINE-ENZYME parameters (e.g. *pRemove*), respectively.

Modularity and integration

We modularized the analog components and mechanisms using methods outlined in [10]. Modularization facilitates reusing and integrating components among different models. For example, immune system components (namely KUPFFER CELL) were initially developed as part of the ISHC. Since components were modularized, they were easily shared with the ISL with minimal code refactoring.

Selecting parameter values

Where applicable, we began with parameter values used to achieve validation targets in previous ISL and ISHC publications (e.g. [10,13]). For parameters introduced in this work, we began with modest values: e.g. 0.5 for probabilistic parameters like *pRemove* with range [0,1], 1 for threshold parameters like *inflammatoryStimuli* based on number of objects, and 1.0 for weighting parameters like *cytokineExponent*. From there, parameter value options are evaluated iteratively, following Steps 5 - 7 of the IRP (see Fig 1A). Simultaneous, small changes (e.g. 5 - 10%) in several parameter values can offset each other and may produce no detectable change in a measured phenomenon. Thus, conventional linear sensitivity studies are less informative and meaningful than complete location changes in analog

parameter space. When needed, we use batch parameter space sampling (as in [39,40]) to identify small subsets of parameter values that are most influential for particular attributes. We also often rely on heuristics and trial-and-error to arrive at validating parameterizations. For the exploratory studies herein, we only present one validating parameter set for each experiment (see Table 1). Narrative explanations for parameter value choices are provided at the end of each experiment in Results.

Results

A full list of parameters descriptions and their values used for validation experiments is provided in Table 1. Some parameters are specific to the SOLUTE or ENZYME types used in each simulation; these parameter values are provided in Tables 2 and 3, respectively. Validation targets and associated validation data are summarized in Table 4.

Table 4. Validation targets achieved for immune-mediated P450 down-regulation attributes.

Targeted attribute	Validation data	Similarity criteria
Kupffer cells produce cytokine upon LPS stimulus in vitro (Fig 2A).	Dose-response curve between LPS dose and TNF-α response using an in vitro Kupffer cell culture [23].	In silico values fall within ± 1 standard deviation of wet-lab values.
Cytokines down-regulate hepatic P450 levels in vitro (Fig 2B).	Time-course drop in P450 levels after IL-1 stimulus using an in vitro hepatocyte culture [24].	In silico values fall within ± 1 standard deviation of wet-lab values.
LPS reduces APAP clearance in rats (Figs 3-5).	1) Disappearance curves and 2) half-life values with/without LPS pretreatment in rats [20].	1) >50% in silico values fall within ± 25% of wet-lab values. 2) In silico values fall within ± 1 standard deviation of wet-lab values.

LPS reduces ANT clearance (Figs 3-5) and CYP3A2/2C11 levels (Fig 2C) in rats.	1) Disappearance curves, 2) relative CYP3A2/2C11 levels, 3) half-life values, and 4) relative systemic clearance with/without LPS pretreatment in rats [21].	1) >50% in silico values fall within ± 25% of wet-lab values. 2) In silico values fall within ± 1 standard deviation of wet-lab values.
LPS reduces CZN clearance (Figs 3-5) and CYP2E1 levels (Fig 2C) in rats.	1) Disappearance curves, 2) CYP2E1 levels, 3) half-life values, and 4) relative intrinsic clearance with/without LPS pretreatment in rats [22].	1) >50% in silico values fall within ± 25% of wet-lab values. 2-4) In silico values fall within ± 1 standard deviation of wet-lab values.

ISHC experiments

We first aimed to achieve a degree of validation for the first part of the pathway: an LPS stimulus causes Kupffer cell-mediated cytokine release. The validation data, derived from [23], is a dose-response curve between LPS dose and TNF-α response using an in vitro Kupffer cell culture. TNF-α was measured after 48-hr LPS treatment using LPS concentrations of 0, 0.1, 1, 10, and 1,000 ng/ml. To simulate this in vitro Kupffer cell culture, we instantiated an ISHC with the only CELLS being KUPFFER CELLS (i.e. *kcDensity* = 1.0 and *hepDensity* = 0.0). We specified the analog-to-referent mapping of 700 LPS objects to 1 ng/ml LPS, and 1 simulation cycle to 1 min. We arrived at this LPS mapping after several iterations, as it provided a consistent sigmoidal dose-response curve over a wide range of parameter values. The time mapping was the same as previous ISHC experiments [10]; it is intentionally coarse-grain, as current use cases do not require mimicking temporally fine-grain phenomena.

To mimic the dose-response curve between 0 and 1,000 ng/ml LPS, a DOSE ranging from 0 to 700,000 LPS objects was injected at the start of the simulation. The total number of CYTOKINES was measured after 2,880 simulation cycles (maps to 48 hr). The resulting values were used to construct a dose-response curve, normalized by the maximum obtained number of CYTOKINES. We prespecified that the analog dose-response curve would be acceptably similar to the referent if each in silico point falls within ± 1 standard deviation of the corresponding wet-lab value. In accordance with the IRP, we sampled ISHC parameter space until similarity criteria were achieved. The resulting validated dose-response curve is shown in Fig 2A.

Fig 2. Validation targets for LPS treatment, before or without drug administration. A. Dose-response curve between LPS stimulus and normalized cytokine response. Values were measured after 48 hr (2,880 simulation cycles). Error bars: wet-lab standard deviation. In silico points are averages of 16 Monte Carlo trials. Wet-lab values are from [23]. **B.** Time-course levels of enzymes, normalized by the starting value. Error bars: wet-lab standard deviation. In silico points are averages of 16 Monte Carlo trials. Wet-lab values are from [24]. **C.** Wet-lab and in silico P450 levels relative to control values. Wet-lab values are relative measures of CYP3A2 (ANT) or CYP2E1 (CZN). In silico values are relative measures of the respective ENZYME type. Error bars: standard deviation. Wet-lab values are from [20] (APAP), [21] (ANT), and [22] (CZN). Note [20] did not provide P450 data for APAP, but we included in silico values for comparison.

When selecting parameter values and the LPS analog-to-referent mapping for this experiment, the resulting dose-response curve followed a sigmoidal shape for all tested parameter values. Thus, we first arrived at a robust choice for the LPS analog-to-referent mapping. The LPS mapping was chosen as a balance between granularity and computational efficiency. It is reasonable to strike that balance because both ISL and ISHC are analogies, not one-to-one models of their referents. When the mapping was too large (e.g. 1000 LPS objects maps to 1 ng/ml LPS), large DOSES (over one million LPS objects)

603 significantly increased computational costs. When the mapping was too small (e.g. 100 LPS objects maps
604 to 1 ng/ml LPS), the smallest non-zero DOSE (mapping to 0.1 ng/ml LPS) consisted of only 10 LPS objects
605 (in the presence of 625 grid points in CELL SPACE), which was deemed too coarse-grain to provide
606 meaningful insight. We arrived at the mapping of 700 LPS objects to 1 ng/ml LPS after several iterations,
607 then proceeded to test selected parameter values. The challenge in specifying parameter values was to
608 control the microstructure of the resulting sigmoidal curve (e.g. the slope of the inflection point). In this
609 case, *inflammatoryThreshold* was a sensitive parameter; it effectively controlled the DOSE at which the
610 CYTOKINE response begins to sharply increase. The value of *pDegrade* for CYTOKINE was also sensitive;
611 when too small (e.g. 0.001), the slope of the inflection point was too shallow; when too large (e.g. 0.1),
612 the inflection point would shift too far to the right.
613 We then aimed to achieve a degree of validation for the second part of the pathway: an increase in
614 cytokines results in hepatic P450 down-regulation. The targeted attribute was data from an in vitro
615 hepatocyte culture: specifically, time-course P450 levels after IL-1 stimulus [24]. We mimicked the in
616 vitro hepatocyte culture by instantiating an ISHC with only HEPATOCYTES (i.e. *kcDensity* = 0.0 and
617 *hepDensity* = 1.0). A DOSE of 2,000 CYTOKINES was injected at the start of the simulation. The total
618 number of G1-ENZYMES was measured at intervals mapping to corresponding wet-lab points (0, 12, and
619 24 hr), plotted against time, and normalized by the starting number of G1-ENZYMES. Similarity criteria
620 stipulated that each in silico time measurement fall within ± 1 standard deviation of the corresponding
621 wet-lab value. Fig 2B illustrates the time-course plot after identifying parameter values that achieved
622 validation.
623 Despite DOWN-REGULATION HANDLER being a finer grain mechanism, specifying parameter
624 values for Fig 2B was straightforward. Parameter value choices were influenced by the following
625 deduction: given enough time, the simulation would eventually return to basal levels of ENZYME. This is
626 because the CYTOKINE signal would eventually die due to DEGRADATION HANDLER, and ENZYMES will
627 eventually be replenished (provided *pReplenish* is greater than zero). That scenario is consistent with
628 what one would expect physiologically. Thus, a moderate value for *delay* (mapping to 30 minutes) and

small value for *pReplenish* (0.007) produced the relatively slow time-course P450 down-regulation as found in the wet-lab validation data.

Single-DRUG ISL experiments

The next three sets of validation targets were designed to achieve degrees of validation for the entire coarse-grain pathway: an LPS stimulus results in hepatic P450 down-regulation and downstream changes in measures of drug clearance. We drew validation targets from three wet-lab experiments, each of which followed a similar experimental protocol but used a different drug and measure of drug clearance. Each experiment pretreated rats with LPS (experimental) or saline (control) for 24 hr before injecting a drug. Clearance was measured at prespecified time points following drug injection. Enzyme measures were taken after the 24-hr pretreatment. All experiments measured disappearance curves from blood (APAP) or plasma (ANT, CZN) for their respective drug. In addition, [20] provide APAP half-life values; [21] measure ANT half-life and systemic clearance (CL_{sys}) and CYP3A2/CYP2C11 levels; [22] measure CZN half-life and create a scatterplot of CZN intrinsic clearance versus CYP2E1 levels. Similarity criteria were prespecified as follows. For drug disappearance curves, at least 50% of in silico values must fall within ± 25% of the corresponding wet-lab value. For enzyme measurements, in silico values must fall within ± 1 standard deviation of the wet-lab value. For all clearance measures except ANT half-life, the in silico value must fall within ± 1 standard deviation of the corresponding wet-lab value. Standard deviation for ANT half-life was neither given nor able to be determined from the data given, so we specified the similarity criteria that in silico values fall within ± 10% of the corresponding wet-lab value.

Simulating whole rat experiments required switching from experimenting on an ISHC to an ISL. We first instantiated an ISL for each of the six experiments: two treatment groups (control and LPS) for each of three drugs. For ISL simulations, we specified the analog-to-referent mapping of 1 simulation cycle to 1 sec, consistent with previous ISL experiments [37]. This constitutes a 60-fold finer temporal resolution than the ISHC mapping of 1 simulation cycle to 1 min. Thus, the ISL is intended to be a

654 temporally finer-grained analog, as required to capture hepatic transit times on the order of seconds. (Note
655 we did not need analog-to-referent mappings for drug or enzyme concentrations because measurements
656 were normalized.) At the start of each simulation in the LPS group, an initial DOSE of 125,000 LPS objects
657 was injected into BODY. After 86,400 simulation cycles (maps to 24 hr), a second DOSE of 125,000 DRUG
658 objects (APAP, ANT, or CZN) was injected into BODY. At this point we also measured the total number of
659 ENZYMES of the type corresponding to the injected DRUG. At time points corresponding to the respective
660 wet-lab data, we measured the number of DRUG objects in BODY.
661 For control group simulations, we bypassed the 86,400-simulation cycle phase. So doing spared
662 computational costs without altering simulation results, because without LPS there are no ISL mechanisms
663 that could alter ENZYME levels or affect DRUG clearance; further, these simulations do not require
664 "priming" from initialization. Thus, there was only a single DOSE of 125,000 DRUG objects at the start of
665 each control group simulation. After DRUG injection, control and LPS group experiments operated
666 identically.
667 Various parameter value combinations were evaluated for each experiment until validation targets
668 were achieved. We imposed several constraints when selecting those parameters. Namely, we ensured
669 that whole-model parameter values (which are neither SOLUTE- nor ENZYME-specific) could not change
670 among the six experiments. Furthermore, DRUG and ENZYME properties could not change between control
671 and LPS experiments of the same DRUG type. The reasoning behind these constraints was to mimic
672 experimenting on the same (or similar) rats. For experiments involving a particular DRUG (say APAP), the
673 DRUG-specific parameter values of other DRUGS (i.e. ANT and CZN) were still included in the parameter
674 set, but their DOSE was set to zero. Similarly, all three ENZYME types are included in APAP simulations,
675 though ANT-ENZYMES and CZN-ENZYMES do not affect APAP results. Thus, the only parameter values that
676 actually changed between simulation experiments were in silico experiment design parameters (i.e. those
677 having no biological counterpart): those specifying DOSE (i.e. whether or not to inject LPS; which DRUG to
678 inject), experiment duration (i.e. how many simulation cycles), and measurements (i.e. when to take
679 measurements). As a consequence of the above constraints, validation is properly limited to those cases in

680 which each experiment achieved validation targets using a common set of in silico experiment design
681 parameters.
682 Meeting the above set of strict validation targets required several rounds of iterative refinement.
683 The most significant refinement was the need to generalize ENZYME mechanisms, outlined in Methods.
684 While ISHC experiments achieved validation targets using G1-ENZYMES and associated mechanisms,
685 achieving ISL validation targets required G2-ENZYMES. Additional refinements are detailed in
686 Discussion.
687 The following figures recapitulate the selected wet-lab data using in silico results from validating
688 experiments. Fig 3 shows the disappearance curves with wet-lab data overlaid. Sixty percent of values fall
689 within acceptable similarity (± 25% of the wet-lab value) for APAP data; 100% of values fall within
690 acceptable similarity for ANT and CZN data. Fig 2C shows the relative change in P450 levels after 24 hr
691 LPS pretreatment. While [20] did not provide P450 data (and thus no validation targets can be drawn for
692 APAP experiments), in silico values for ANT and CZN experiments fall within the prespecified similarity
693 criteria of ± 1 standard deviation of the wet-lab value. In fact, they fall within a more stringent ± 0.5
694 standard deviation similarity criteria. Fig 4 shows the half-lives for control experiments, relative half-lives
695 for LPS experiments, and relative clearance measures (half-life, systemic clearance, and intrinsic
696 clearance for APAP, ANT, and CZN, respectively) alongside validation data. All values fall within
697 acceptable similarity (± 1 standard deviation of the wet-lab value). Lastly, Fig 5 shows scatterplots of
698 relative enzyme versus clearance measures, with validation data overlaid.
699
700 **Fig 3. Wet-lab and in silico normalized drug disappearance curves. A.** APAP [20]; **B.** ANT [21]; **C.**
701 CZN [22]. Closed circles: in silico averages of 16 Monte Carlo trials. Gray circles: wet-lab averages.
702 Red/blue lines: additional in silico values between wet-lab time points. The initial spike in drug
703 corresponds to the administered DOSE. All drug values are normalized by the control value at the first
704 time point. Error bars: ± 25% of the wet-lab value (the similarity criteria).
705

Fig 4. Wet-lab and in silico measures of drug clearance. A. Wet-lab and in silico half-life measures without LPS pretreatment (control). **B.** Wet-lab and in silico half-life given 24 hr LPS pretreatment, relative to control. **C.** Wet-lab and in silico clearance measures given 24 hr LPS pretreatment, relative to control. Error bars: standard deviation. Wet-lab values are from [20] (APAP), [21] (ANT), and [22] (CZN).

Fig 5. Scatterplots between enzyme measurements and clearance measurements for both control and LPS experiments. A. APAP [20]; **B.** ANT [21]; **C.** CZN [22]. Gray circles: wet-lab data points (when provided). Red/blue circles: in silico data points. Error bars: in silico standard deviation, extending from the mean of 16 Monte Carlo trials. Blue box: area of acceptable similarity (± 1 standard deviation of wet-lab value). Since [20] did not provide enzyme data, there is no associated validation target (**A**). Only [22] provided values for individual wet-lab trials (**C**).

Whereas ISHC experiments focused on small parts of the P450 down-regulation pathway—and thus selecting parameter valuess required focus on the one or two relevant mechanisms—the process for ISL experiments required attention to all mechanisms. For example, changes in INFLAMMATION HANDLER parameter values affect CYTOKINE levels, which directly affects downstream DOWN-REGULATION HANDLER behavior even if its parameters do not change. We first focused on INFLAMMATION HANDLER because it contains no SOLUTE- or ENZYME-specific parameters. Most notably, *inflammatoryThreshold* was chosen higher than the ISHC value (see Table 1); this was necessary to avoid excessive CYTOKINE levels. DOWN-REGULATION HANDLER required significant parameter changes compared to the ISHC. The DOWN-REGULATION HANDLER parameters are all sensitive to time: an ENZYME removal or replenish event may be scheduled each simulation cycle. Thus, it is not surprising that the ISL—with a 60-fold finer grain temporal mapping than the ISHC—required smaller values for *pRemove* and *pReplenish* and a larger value for *delay* compared to the ISHC. The simulation must eventually return to basal levels of ENZYME if left running indefinitely. However, if down-regulation were too sensitive (e.g. *pRemove* too high or

CYTOKINE'S *pDegrade* too low), ENZYME levels could reach zero for long periods of time before replenishing. Thus, selection of parameter values often revolved around achieving sustained, but not excessive, levels of P450 down-regulation.

Multi-DRUG ISL experiments

Lastly, we repeated both control and LPS experiments for a new use case in which all three DRUGS are co-administered. Wet-lab validation data for such co-administration experiments in the context of P450 down-regulation are unavailable. However, in silico multi-DRUG experiments represent an important use case for several reasons. 1) Multi-DRUG simulations stand as challengeable predictions of hypothetical wet-lab co-administration experiment results. Given new wet-lab co-administration data that falsifies these predictions, we can then hypothesize model refinements in accordance with the IRP. Refinements may include analog mechanisms of drug-drug interactions; however, details of the path to revision would depend on the nature of the results. 2) We demonstrate a proof-of-concept that the ISL can be made robust to changes in analog DRUG, and thus able to support co-simulation of any number of DRUG types. Extensibility to multiple DRUG types is essential to support future simulations designed to better understand and anticipate drug-drug interactions. 3) We demonstrate that all three DRUG types can simultaneously achieve validation targets using a single ISL parameter set. Simultaneous validation is significant because, if achieved, we can then conclude that the DRUG- and ENZYME-specific parameterizations are sufficient to span the mechanistic changes needed to generate patterns in drug clearance and P450 down-regulation. In other words, values for whole-model parameters that are neither DRUG- nor ENZYME-specific (e.g. *hepDensity*) need not be specified independently for each DRUG type. 4) Most importantly, failure to achieve simultaneous validation could be taken as evidence that hypotheses and several aspects of the ISL are falsified.

Notably, for this work, multi-DRUG experiments were not intended to explore or mimic drug-drug interactions. However, we also did not expect identical results between single- and multi-DRUG simulations because the presence of additional DRUG objects (even without the inclusion of explicit drug-

drug interactions) may still indirectly affect ISL mechanisms. For example, the flow of SOLUTES depends on the total number of SOLUTE objects , so multi-DRUG experiments with a larger total number of SOLUTES may lead to subtle changes in the cascading of events. We expect more pronounced differences between single- and multi-DRUG experiments when individual cytokines and P450 isoforms are made explicit and/or when we explicitly add mechanisms that handle drug-drug interactions.

While the three wet-lab validation studies administered different amounts of LPS stimulus, we assumed an equipotent stimulus. We do not expect this equipotent assumption to be biologically realistic; however, in the absence of P450 down-regulation wet-lab data spanning both single and multiple drug experiments, we needed to make some assumption to simulate multi-drug experiments. Thus, BODY was injected with 125,000 LPS objects, similar to single-DRUG experiments. The second DOSE comprised 125,000 (for LPS experiments) or 62,500 (for control experiments) objects of each DRUG type (APAP, ANT, and CZN). Thereafter, single- and multi-DRUG experiments operated identically.

The validation target for the multi-DRUG experiments was to simultaneously achieve all validation targets for the single-DRUG experiments. Besides the in silico experiment design parameters (those controlling DOSE, experiment duration, and measurements), parameter values were not changed between single- and multi-DRUG experiments. Notably, all multi-DRUG validation targets were achieved without changing parameterizations that achieved the single-drug validation targets. Figs 2C and 4 include values for the multi-DRUG experiments alongside single-DRUG results and wet-lab validation data.

Discussion

We execute a series of virtual experiments that together achieve several quantitative validation targets involving immune system interactions in drug metabolism. The described mechanisms—those that finally achieved validation targets—are supportive of the mechanistic hypothesis that in response to LPS, Kupffer cells down-regulate P450 via inflammatory cytokines, thus leading to a reduction in metabolic capacity. Given the synthetic nature of the analogs, we can continue to iteratively refine them to test additional and/or alternative mechanistic hypotheses, validating an increasingly large set of targeted attributes along the way.

Trajectory of model falsification and refinement

Falsification is the primary source of knowledge generation, integral to the scientific method itself [41]. Unfortunately, published computational models tend to describe a finished product without detailing—or even mentioning—the many rounds of falsification and revision along the way. Describing this trajectory of iterative falsification and refinement requires detailed annotations of simulation experiments (including unsuccessful ones) and paying special attention to how, when, and why a mechanism fails to achieve given validation targets. Having that information is essential to the scientific process because it is falsification that provides new knowledge: specifically, the current (falsified) mechanisms are flawed—they are not a good analogy of the referent biological mechanisms [42]. Following falsification, analog mechanisms (the hypothesis) can be refined. Focusing on falsification as a scientifically *productive* endeavor helps one adhere to parsimony, as it curbs the tendency to commit inscription error and/or to succumb to the urge to incorporate "everything we know" into a model.

We describe the trajectory of three demonstrative falsification and refinement efforts encountered with the P450 down-regulation validation targets, and commentate on any mechanistic insight gained from this new knowledge. 1) We described earlier the G1-ENZYME mechanism falsification and the need for the finer-grain, generalized G2-ENZYME (see "ENZYME mechanism falsification and generalization").

Validation required three changes: 1) ENZYME concentration-dependent binding probability 2) ENZYME type-specific properties, and 3) the inclusion of various ENZYME types within a single simulation. Taken together, these changes suggest that differential properties of multiple groups of enzymes influence drug clearance events during inflammatory states. While this dependence might be expected, we anticipate providing additional insight when ISL experiments move toward explicitly modeling individual P450 isoforms.

2) The large drop in CZN clearance (drops to 29.2% of control value), coupled with a moderate drop in CZN enzyme levels (drops to 57.8% of control value), could not initially be reproduced in silico. Using original mechanisms, relative decreases in clearance values were always too weak compared to corresponding decreases in the number of ENZYMES. After falsification, we noted that the wet-lab validation experiment recorded a 1.69-fold increase in CZN volume of distribution after the addition of LPS. Thus, we posited a revision hypothesis that validation could be achieved by adding a coarse-grain mechanism for changes in volume of distribution. Recall that a fraction of SOLUTES in BODY are transferred to LOBULE each simulation cycle. This fraction is equal to *sampleRatio* (a whole-model parameter between 0 and 1) multiplied by *sampleRatioFactor* (a SOLUTE-specific parameter). We implemented the following coarse-grain change. In LPS experiments, the size of the initial DOSE and the value of *sampleRatioFactor* for CZN are reduced by a factor of $V_{d,change}$, where $V_{d,change}$ is the wet-lab fold change in volume of distribution. After implementing this revision, we achieved the original validation targets. Thus, accounting for changes in volume of distribution is necessary to mimic P450 down-regulation phenomena for some drugs. As new, finer-grained validation targets are considered, ISL experiments can flesh out mechanistic details driving changes in volume of distribution.

3) Initially, INFLAMMATION HANDLER did not include *cytokineThreshold*, a parameter that "turns off" CYTOKINE production if there is enough CYTOKINE in the CELL. However, this exclusion resulted in unchecked CYTOKINE formation (and thus a non-sigmoidal dose-response curve) when the amount of INFLAMMATORY STIMULUS was sufficiently high. Since our validation experiments included large amounts of CYTOKINE (up to 700,000 objects), that mechanism was falsified. We implemented the

826 *cytokineThreshold* parameter, with which we found sets of parameter values that achieved validation
827 targets. This cytokine threshold mechanism is coarse-grain. It may map to more complex cytokine-
828 suppression mechanisms including feedback inhibition and suppressor of cytokine signaling proteins [43].
829 Given sufficient validation data, we can add finer grain mechanistic features by introducing new SOLUTE
830 types (e.g. those that map to suppressor of cytokine signaling proteins) and behaviors in place of the
831 existing cytokine threshold mechanism.

832 Current model falsification

833 Reporting validation results without demonstrating subsequent falsification would provide an
834 incomplete picture of the IRP. We describe current model falsification to address both the limitations and
835 future direction of the ISHC and ISL. The current models are limited to the majority of P450s that are
836 down-regulated by an inflammatory state; however, a small subset of P450s are actually induced [1].
837 Clearance data for drugs primarily metabolized by this subset is expected to falsify existing mechanisms.
838 Given that scenario, we can iteratively refine mechanisms to simultaneously allow both inflammatory-
839 induced P450 down-regulation and induction.
840 The analogs are also falsified by additional immune system mechanisms other than P450 down-
841 regulation. For example, inflammatory states cause leukocyte recruitment and localization via
842 chemokines secreted from sites of necrosis [44,45]. Achieving validation targets drawn from these
843 phenomena would require in silico mechanisms for leukocyte circulation, recruitment, and extravasation,
844 as well as their downstream effects on hepatotoxicity. As we continue expanding the set of targeted
845 attributes to include more diverse yet still interconnected phenomena, we expect IRP cycles to continue
846 improving explanatory mechanistic insight into immune system involvement in drug metabolism and
847 toxicity.

Acknowledgments

BKP is the recipient of a UCSF Discovery Fellowship. We thank members of the BioSystems Group for helpful discussion and commentary.

References

1. Morgan ET. Regulation of cytochromes P450 during inflammation and infection. Drug Metab Rev. 1997;29(4):1129-88.
2. Morgan ET. Regulation of cytochrome p450 by inflammatory mediators: why and how? Drug Metab Dispos. 2001;29(3):207-12.
3. Aitken AE, Richardson TA, Morgan ET. Regulation of drug-metabolizing enzymes and transporters in inflammation. Annu Rev Pharmacol Toxicol. 2006;46:123-49.
4. Liu J, Sendelbach LE, Parkinson A, Klaassen CD. Endotoxin pretreatment protects against the hepatotoxicity of acetaminophen and carbon tetrachloride: role of cytochrome P450 suppression. Toxicology. 2000;147(3):167-76.
5. Aitken AE, Morgan ET. Gene-specific effects of inflammatory cytokines on cytochrome P450 2C, 2B6 and 3A4 mRNA levels in human hepatocytes. Drug Metab Dispos. 2007;35(9):1687-93.
6. Morgan E. Impact of infectious and inflammatory disease on cytochrome P450–mediated drug metabolism and pharmacokinetics. Clin Pharmacol Ther. 2009;85(4):434-8.
7. Sunman JA, Hawke RL, LeCluyse EL, Kashuba AD. Kupffer cell-mediated IL-2 suppression of CYP3A activity in human hepatocytes. Drug Metab Dispos. 2004;32(3):359-63.
8. Morgan ET, Li-Masters T, Cheng P-Y. Mechanisms of cytochrome P450 regulation by inflammatory mediators. Toxicology. 2002;181:207-10.
9. Peck SL. Simulation as experiment: a philosophical reassessment for biological modeling. Trends Ecol Evol. 2004;19(10):530-4.
10. Petersen BK, Hunt CA. Developing a vision for executing scientifically useful virtual biomedical experiments. In: Proceedings of the 2016 Spring Simulation Multiconference 2016 Apr 3. Society for Computer Simulation International.

11. Hunt CA, Kennedy RC, Kim SH, Ropella GE. Agent-based modeling: a systematic assessment of use cases and requirements for enhancing pharmaceutical research and development productivity. Wiley Interdiscip Rev Syst Biol and Med. 2013;5(4):461-80.
12. Petersen BK, Ropella GE, Hunt CA. Toward modular biological models: defining analog modules based on referent physiological mechanisms. BMC Syst Biol. 2014;8(1):95.
13. Yan L, Ropella GE, Park S, Roberts MS, Hunt CA. Modeling and simulation of hepatic drug disposition using a physiologically based, multi-agent in silico liver. Pharm Res. 2008;25(5):1023-36.
14. Smith AK, Ropella GE, Kaplowitz N, Ookhtens M, Hunt CA. Mechanistic agent-based damage and repair models as hypotheses for patterns of necrosis caused by drug induced liver injury. In: Proceedings of the 2014 Summer Simulation Multiconference 2014 Jul 6 (p. 16). Society for Computer Simulation International.
15. Luke S, Cioffi-Revilla C, Panait L, Sullivan K. Mason: A new multi-agent simulation toolkit. In: Proceedings of the 2004 Swarmfest Workshop 2004 May 9 (Vol. 8, p. 44).
16. Hunt CA, Ropella GE, Lam TN, Tang J, Kim SH, Engelberg JA, et al. At the biological modeling and simulation frontier. Pharm Res. 2009;26(11):2369-400.
17. Hunt CA, Ropella GE, Yan L, Hung DY, Roberts MS. Physiologically based synthetic models of hepatic disposition. J Pharmacokinet Pharmacodyn. 2006;33(6):737-72.
18. Pogson M, Smallwood R, Qwarnstrom E, Holcombe M. Formal agent-based modelling of intracellular chemical interactions. Biosystems. 2006;85(1):37-45.
19. Park S, Kim SH, Ropella GE, Roberts MS, Hunt CA. Tracing multiscale mechanisms of drug disposition in normal and diseased livers. J Pharmacol Exp Ther. 2010;334(1):124-36.
20. Laskin DL, Gardner CR, Price VF, Jollow DJ. Modulation of macrophage functioning abrogates the acute hepatotoxicity of acetaminophen. Hepatology. 1995;21(4):1045-50.
21. Ueyama J, Nadai M, Kanazawa H, Iwase M, Nakayama H, Hashimoto K, et al. Endotoxin from various gram-negative bacteria has differential effects on function of hepatic cytochrome P450 and drug transporters. Eur J Pharmacol. 2005;510(1):127-34.

902 22. Rockich K, Blouin R. Effect of the acute-phase response on the pharmacokinetics of chlorzoxazone
903 and cytochrome P-450 2E1 in vitro activity in rats. Drug Metab Dispos. 1999;27(9):1074-7.
904 23. Milosevic N, Schawalder H, Maier P. Kupffer cell-mediated differential down-regulation of
905 cytochrome P450 metabolism in rat hepatocytes. Eur J Pharmacol. 1999;368(1):75-87.
906 24. Morgan ET, Thomas KB, Swanson R, Vales T, Hwang J, Wright K. Selective suppression of
907 cytochrome P-450 gene expression by interleukins 1 and 6 in rat liver. Biochim Biophys Acta.
908 1994;1219(2):475-83.
909 25. Sheikh-Bahaei S, Hunt CA. Enabling clearance predictions to emerge from in silico actions of quasi-
910 autonomous hepatocyte components. Drug Metab Dispos. 2011;39(10):1910-20.
911 26. Watari N, Iwai M, Kaneniwa N. Pharmacokinetic study of the fate of acetaminophen and its
912 conjugates in rats. J Pharmacokinet Biopharm. 1983;11(3):245-72.
913 27. Velic I, Metzler M, Hege HG, Weymann J. Separation and identification of phase I and phase II [^{14}C]
914 antipyrine metabolites in rat and dog urine. J Chromatogr B Biomed Sci Appl. 1995;666(1):139-47.
915 28. Moon YJ, Lee AK, Chung HC, Kim EJ, Kim SH, Lee DC, et al. Effects of acute renal failure on the
916 pharmacokinetics of chlorzoxazone in rats. Drug Metab Dispos. 2003;31(6):776-84.
917 29. Patten CJ, Thomas PE, Guy RL, Lee M, Gonzalez FJ, Guengerich FP, et al. Cytochrome P450
918 enzymes involved in acetaminophen activation by rat and human liver microsomes and their kinetics.
919 Chem Res Toxicol. 1993;6(4):511-8.
920 30. Price VF, Jollow DJ. Increased resistance of diabetic rats to acetaminophen-induced hepatotoxicity. J
921 Pharmacol Exp Ther. 1982;220(3):504-13.
922 31. Naritomi Y, Terashita S, Kagayama A, Sugiyama Y. Utility of hepatocytes in predicting drug
923 metabolism: comparison of hepatic intrinsic clearance in rats and humans in vivo and in vitro. Drug
924 Metab Dispos. 2003;31(5):580-8.
925 32. Inaba T, Lucassen M, Kalow W. Antipyrine metabolism in the rat by three hepatic monooxygenases.
926 Life Sci. 1980;26(23):1977-83.

927 33. Sharer JE, Wrighton SA. Identification of the human hepatic cytochromes P450 involved in the in
928 vitro oxidation of antipyrine. Drug Metab Dispos. 1996;24(4):487-94.
929 34. Kobayashi K, Urashima K, Shimada N, Chiba K. Substrate specificity for rat cytochrome P450
930 (CYP) isoforms: screening with cDNA-expressed systems of the rat. Biochem Pharmacol.
931 2002;63(5):889-96.
932 35. Carriere V, Goasduff T, Ratanasavanh D, Morel F, Gautier JC, Guillouzo A, et al. Both cytochromes
933 P450 2E1 and 1A1 are involved in the metabolism of chlorzoxazone. Chem Res in Toxicol.
934 1993;6(6):852-7.
935 36. Smith BC. Inscription errors. In: On the origin of objects. Cambridge, MA: MIT Press; 1996.
936 37. Yan L, Sheihk-Bahaei S, Park S, Ropella GE, Hunt CA. Predictions of hepatic disposition properties
937 using a mechanistically realistic, physiologically based model. Drug Metab Dispos. 2008;36(4):759-
938 68.
939 38. Kirschner DE, Hunt CA, Marino S, Fallahi-Sichani M, Linderman JJ. Tuneable resolution as a
940 systems biology approach for multi-scale, multi-compartment computational models. Wiley
941 Interdiscip Rev Syst Biol and Med. 2014;6(4):289-309.
942 39. Ropella GE, Hunt CA. Cloud computing and validation of expandable in silico livers. BMC Syst
943 Biol. 2010;4(1):1.
944 40. Tang J, Hunt CA. Identifying the rules of engagement enabling leukocyte rolling, activation, and
945 adhesion. PLOS Comput Biol. 2010;6(2):e1000681.
946 41. Popper K. Science as falsification. In: Conjectures and refutations: London: Routledge and Kegan
947 Paul; 1963.
948 42. Ropella GE, Kennedy RC, Hunt CA. Falsifying an enzyme induction mechanism within a validated,
949 multiscale liver model. Int J Agent Technol Syst. 2012;4(3):1-14.
950 43. Nicola NA, Greenhalgh CJ. The suppressors of cytokine signaling (SOCS) proteins: important
951 feedback inhibitors of cytokine action. Exp Hematol. 2000;28(10):1105-12.

44. Li X, Klintman D, Liu Q, Sato T, Jeppsson B, Thorlacius H. Critical role of CXC chemokines in endotoxemic liver injury in mice. J Leukoc Biol. 2004;75(3):443-52.

45. Liu Z-X, Govindarajan S, Kaplowitz N. Innate immune system plays a critical role in determining the progression and severity of acetaminophen hepatotoxicity. Gastroenterology. 2004;127(6):1760-74.

Supporting Information Captions

S1 Fig. Flowcharts of mechanism logic. *U*[0,1) represents a random probability draw from the standard uniform distribution.

S2 Fig. Binding probabilities used by BINDING HANDLER. E1 – E4 represent four ENZYMES. S1 – S5 represent 5 SOLUTES. **A.** Binding probability using stepwise binding mode. **B.** Binding probability using variable binding mode. **C.** P450 up- or down-regulation causes the binding curve to shift right or left, respectively, which has no effect on binding probability in the yellow region (typical range of number of bound ENZYMES). **D.** P450 up- or down-regulation causes the binding curve to shift right or left, respectively, which affects binding probability within the yellow region.

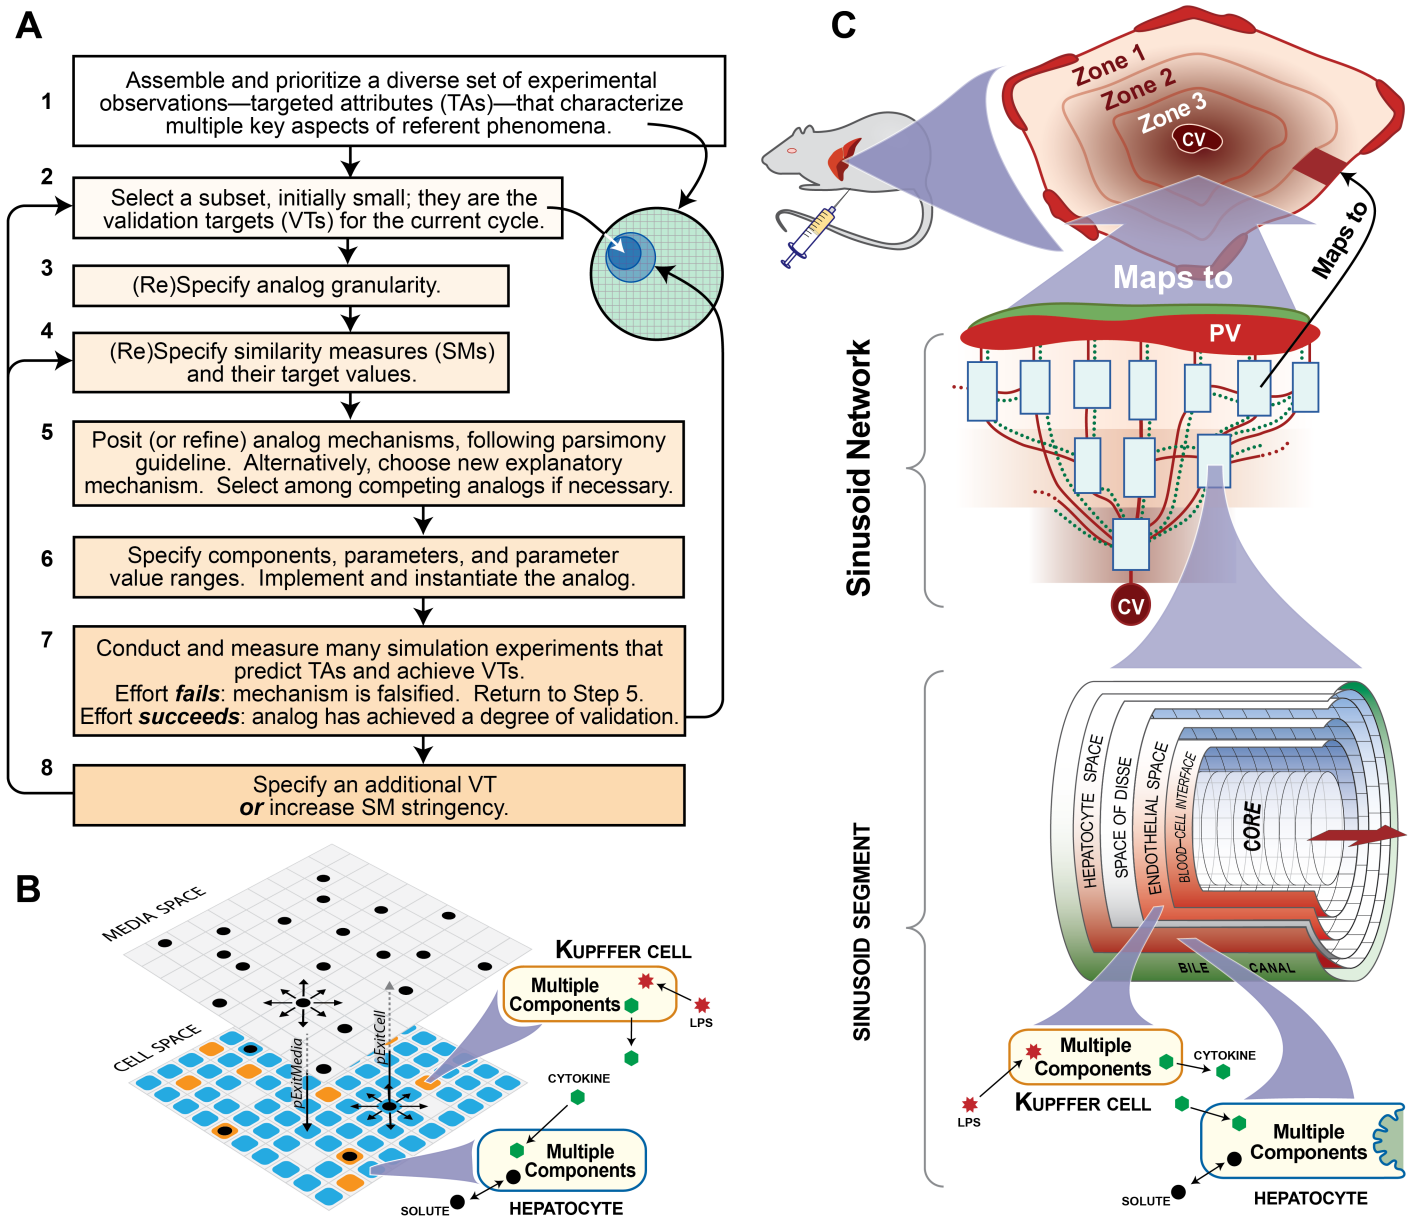
A
1
Assemble and prioritize a diverse set of experimental observations—targeted attributes (TAs)—that characterize multiple key aspects of referent phenomena.
2
Select a subset, initially small; they are the validation targets (VTs) for the current cycle.
3
(Re)Specify analog granularity.
4
(Re)Specify similarity measures (SMs) and their target values.
5
Posit (or refine) analog mechanisms, following parsimony guideline. Alternatively, choose new explanatory mechanism. Select among competing analogs if necessary.
6
Specify components, parameters, and parameter value ranges. Implement and instantiate the analog.
7
Conduct and measure many simulation experiments that predict TAs and achieve VTs.
Effort ***fails***: mechanism is falsified. Return to Step 5.
Effort ***succeeds***: analog has achieved a degree of validation.
8
Specify an additional VT
or increase SM stringency.
B
MEDIA SPACE
CELL SPACE
pExitMedia
pExitCell
KUPFFER CELL
Multiple Components
LPS
CYTOKINE
Multiple Components
SOLUTE
HEPATOCYTE
C
Zone 1
Zone 2
Zone 3
CV
Maps to
Maps to
PV
Sinusoid Network
CV
SINUSOID SEGMENT
HEPATOCYTE SPACE
SPACE OF DISSE
ENDOTHELIAL SPACE
BLOOD–CELL INTERFACE
CORE
BILE CANAL
Multiple Components
CYTOKINE
LPS
KUPFFER CELL
Multiple Components
SOLUTE
HEPATOCYTE

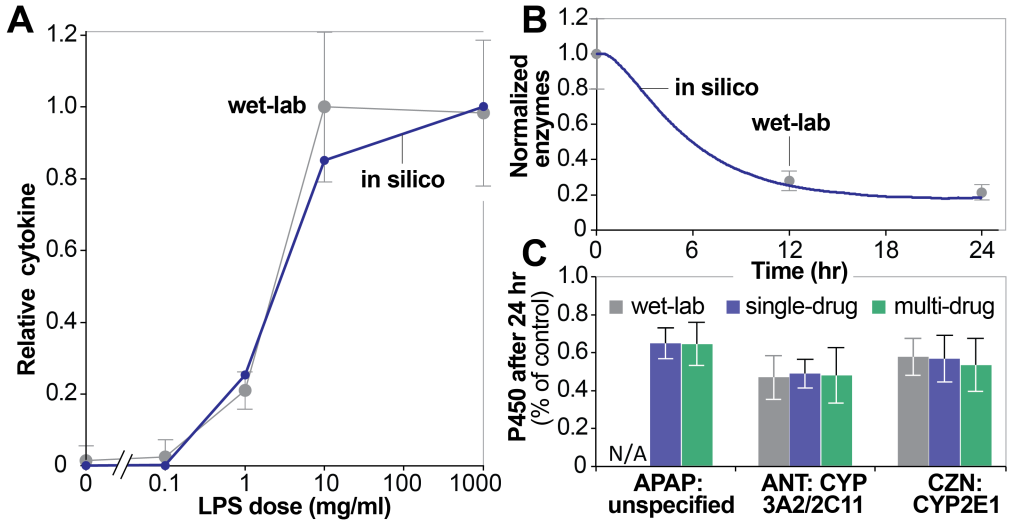
A
Relative cytokine
1.2
1.0
0.8
0.6
0.4
0.2
0
wet-lab
in silico
0
0.1
1
10
100
1000
LPS dose (mg/ml)
B
Normalized enzymes
1.2
1.0
0.8
0.6
0.4
0.2
0
in silico
wet-lab
0
6
12
18
24
Time (hr)
C
P450 after 24 hr (% of control)
1.0
0.8
0.6
0.4
0.2
0
wet-lab
single-drug
multi-drug
N/A
APAP: unspecified
ANT: CYP 3A2/2C11
CZN: CYP2E1

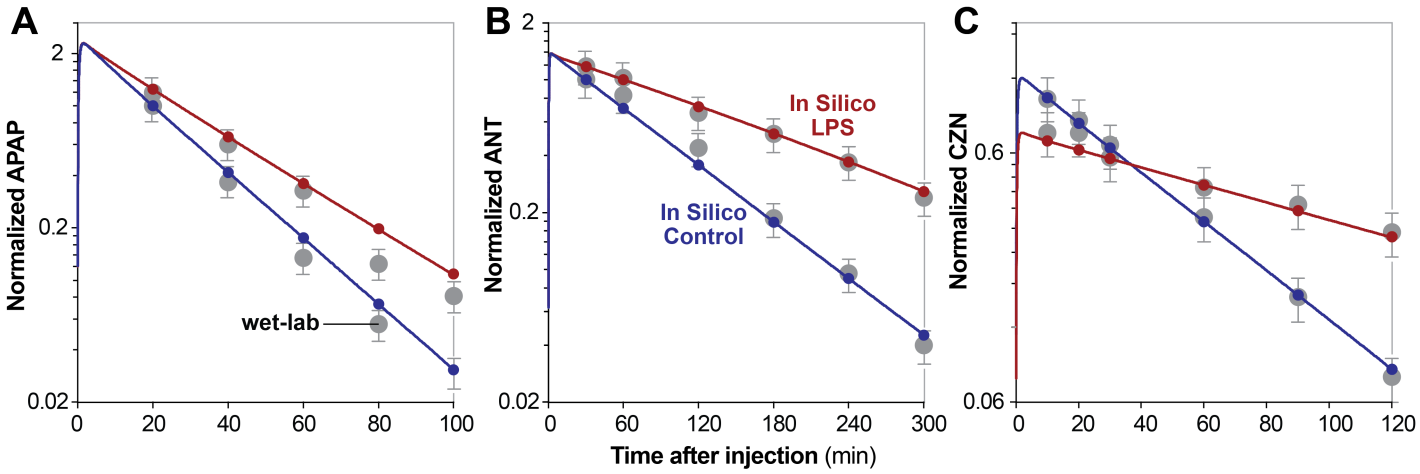
A
B
C
Normalized APAP
Normalized ANT
Normalized CZN
2
0.2
0.02
0.6
0.06
0
20
40
60
80
100
120
180
240
300
wet-lab
In Silico
LPS
In Silico
Control
Time after injection (min)

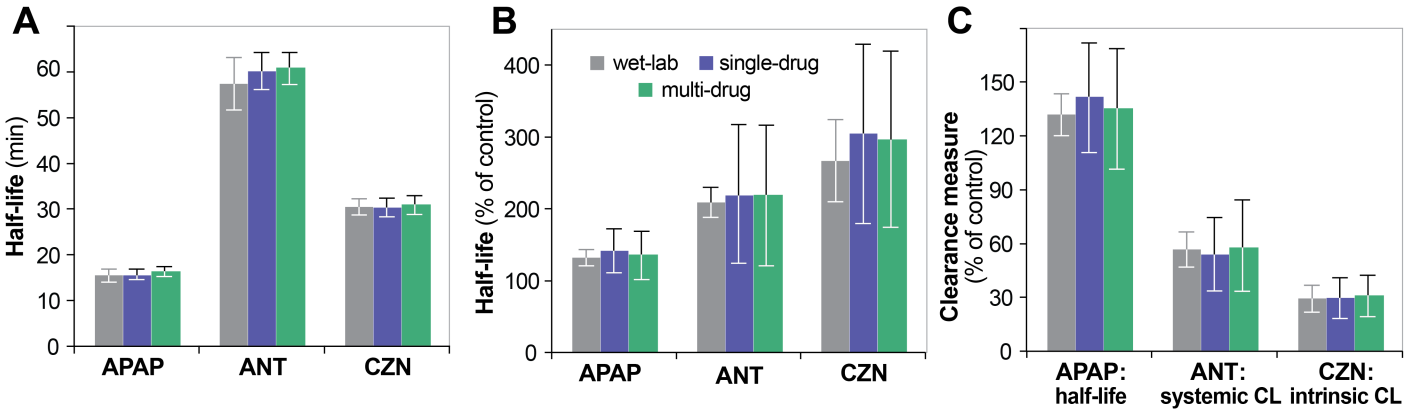
A
Half-life (min)
0
10
20
30
40
50
60
APAP
ANT
CZN
B
Half-life (% of control)
0
100
200
300
400
wet-lab
single-drug
multi-drug
APAP
ANT
CZN
C
Clearance measure (% of control)
0
30
60
90
120
150
APAP: half-life
ANT: systemic CL
CZN: intrinsic CL

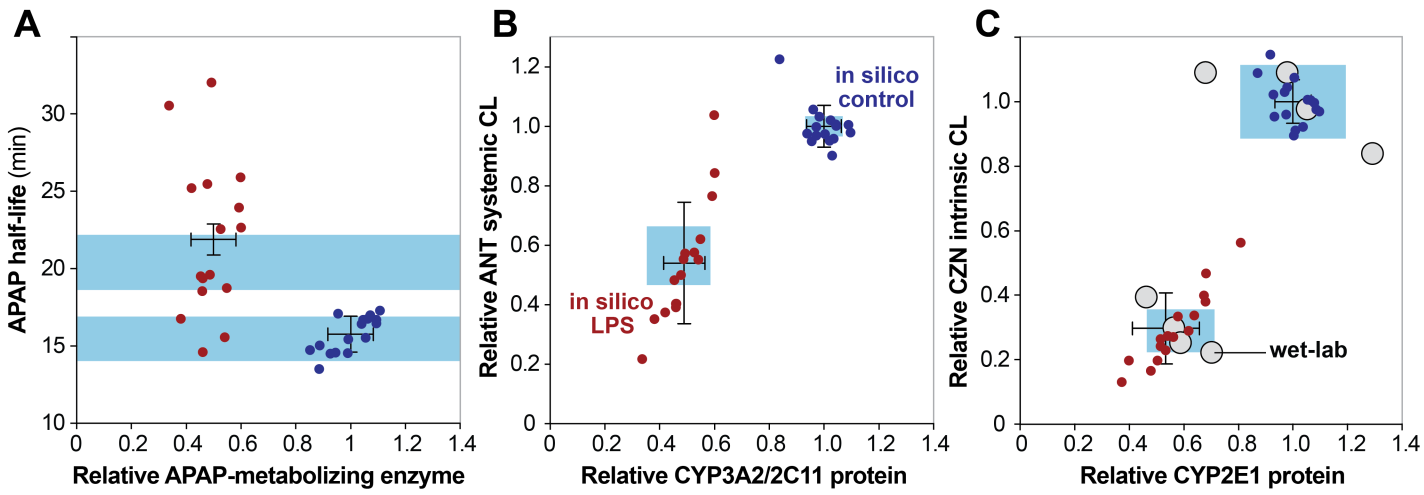
A
APAP half-life (min)
Relative APAP-metabolizing enzyme
B
Relative ANT systemic CL
Relative CYP3A2/2C11 protein
in silico control
in silico LPS
C
Relative CZN intrinsic CL
Relative CYP2E1 protein
wet-lab

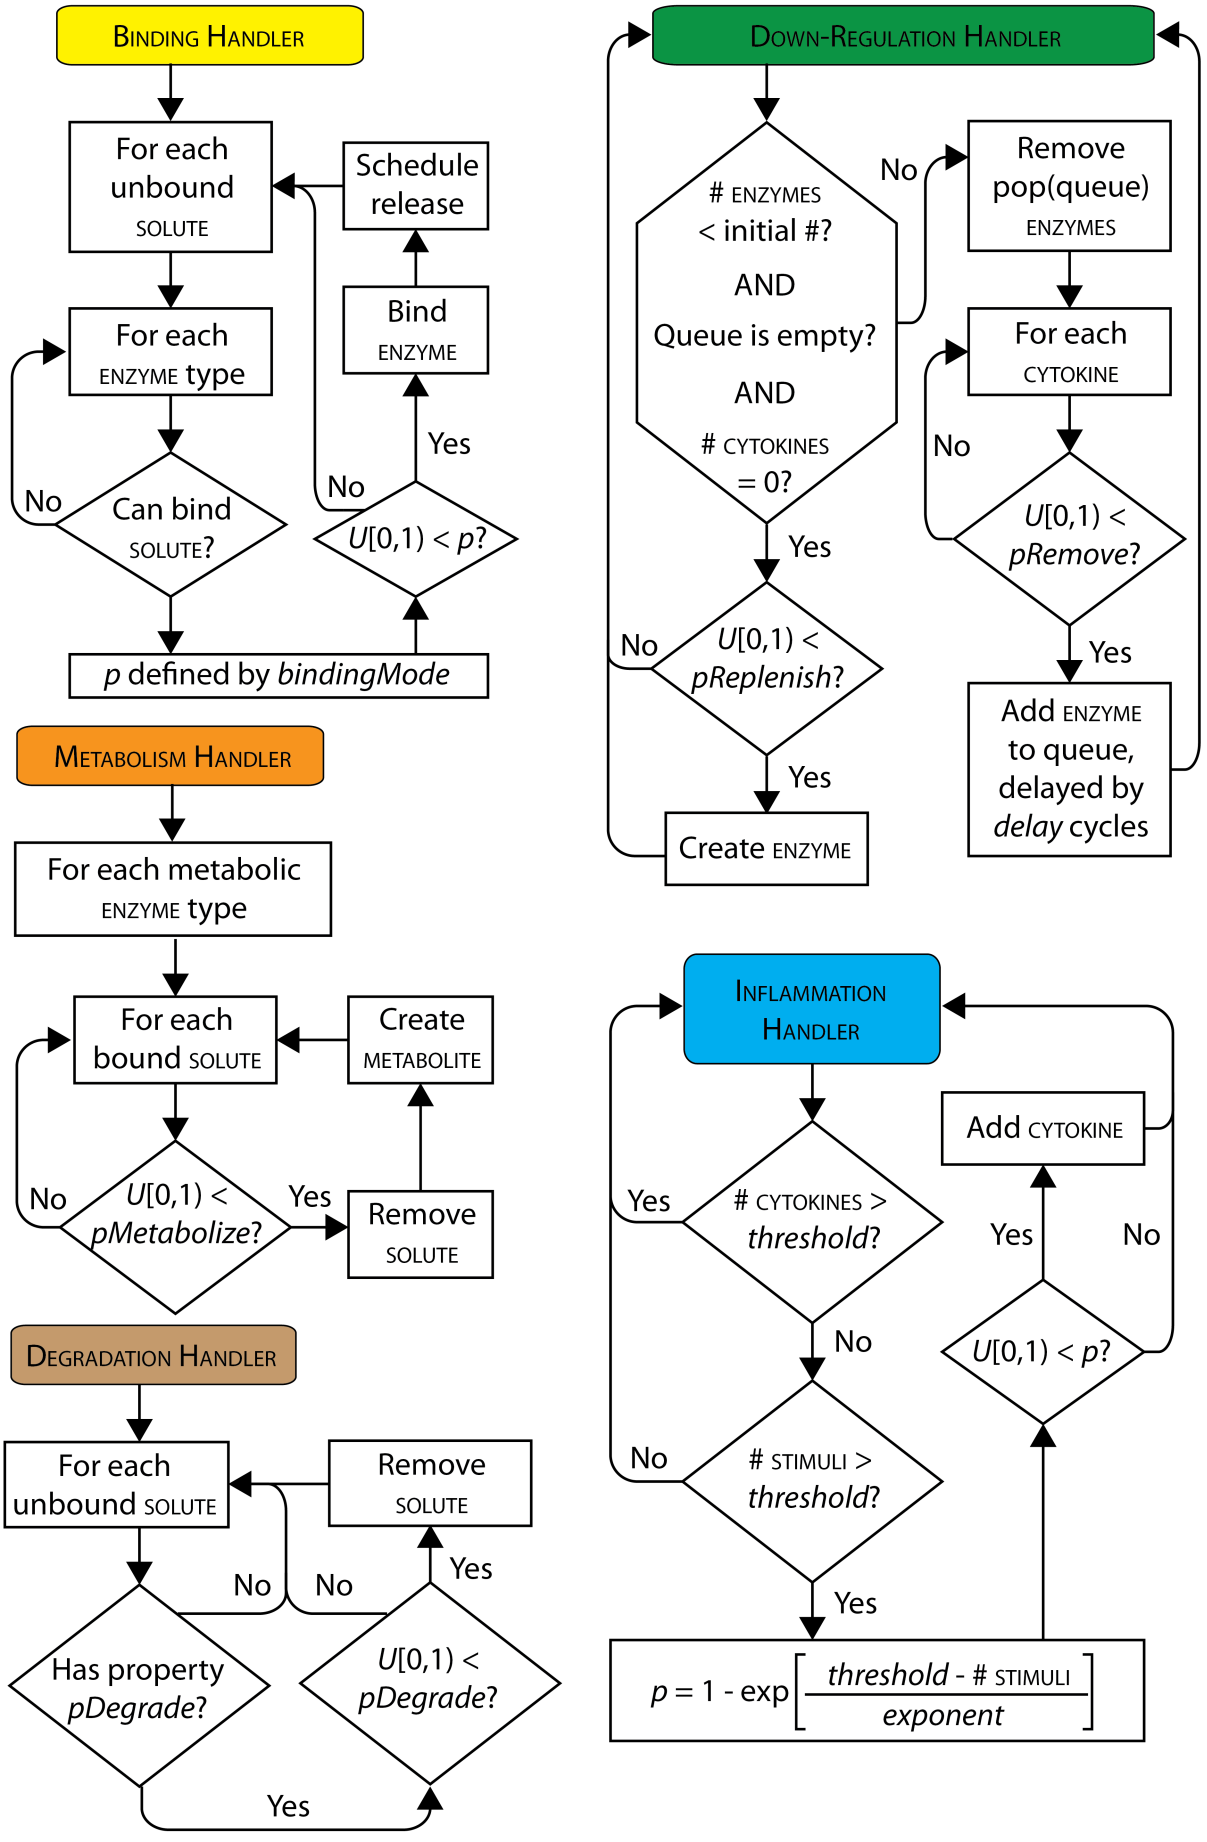
Binding Handler
For each unbound SOLUTE
For each ENZYME type
Can bind SOLUTE?
No
Yes
U[0,1) < p?
No
Bind ENZYME
Schedule release
p defined by bindingMode
Down-Regulation Handler
ENZYMES < initial #? AND Queue is empty? AND # CYTOKINES = 0?
No
Yes
U[0,1) < pReplenish?
No
Yes
Create ENZYME
Remove pop(queue) ENZYMES
For each CYTOKINE
No
U[0,1) < pRemove?
Yes
Add ENZYME to queue, delayed by delay cycles
Metabolism Handler
For each metabolic ENZYME type
For each bound SOLUTE
No
U[0,1) < pMetabolize?
Yes
Remove SOLUTE
Create METABOLITE
Degradation Handler
For each unbound SOLUTE
Has property pDegrade?
No
U[0,1) < pDegrade?
No
Yes
Yes
Remove SOLUTE
Inflammation Handler
Yes
CYTOKINES > threshold?
No
No
STIMULI > threshold?
Yes
p = 1 - exp [threshold - # STIMULI / exponent]
Add CYTOKINE
Yes
No
U[0,1) < p?

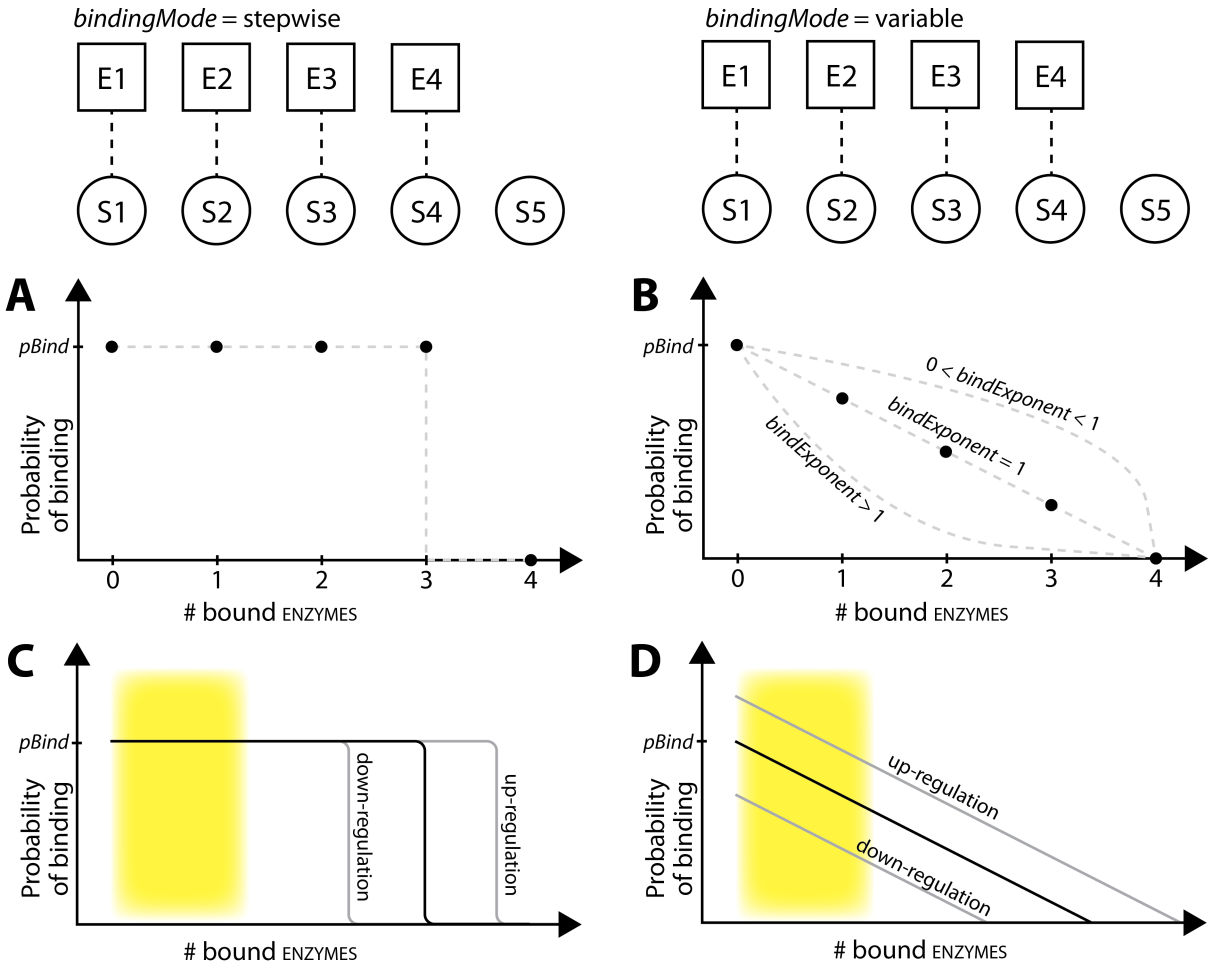
bindingMode = stepwise
E1
E2
E3
E4
S1
S2
S3
S4
S5
bindingMode = variable
E1
E2
E3
E4
S1
S2
S3
S4
S5
A
pBind
Probability of binding
0
1
2
3
4
bound ENZYMES
B
pBind
Probability of binding
0 < bindExponent < 1
bindExponent = 1
bindExponent > 1
0
1
2
3
4
bound ENZYMES
C
pBind
Probability of binding
down-regulation
up-regulation
bound ENZYMES
D
pBind
Probability of binding
up-regulation
down-regulation
bound ENZYMES