ODEs with random parameters

Brian Glucksman and Shan Zhong

Under the Supervision of Dr. James Nolen and Oliver Kelsey Tough

The Models

The starting point for our project is a collection of three models: one to describe the processing of acetaminophen (APAP), another to describe the synthesis of glutathione, and the last to describe the movement of glutamate in the body. We combined these three models to form what we call the *full model*. In some instances, it is advantageous to only use a combination of the first two, which we call the *small model*.

Each of these models are a series of ordinary differential equations. Each differential equation describes the concentration of a certain chemical in a specific organ. For example, the first model contains a differential equation for APAP in the liver and a differential equation for APAP in the gut. The differential equations are populated by three kinds of expressions. First, there is the transport of one chemical from one part of the body to another. Exclusively, these are linear terms. The exact values depend only on the concentration in the place of origin and a transport constant. These equations, to a good approximation, are biological constants, not varying much from individual to individual. Second, there is the decay of certain biological products. Mostly, these are linear terms, but cysteine is a notable exception, which we will talk about later. Like transport terms, these are biological Third, there are reactions. For the most part, these follow Michaelis-Menten kinetics, but there are also a few instances of bi-bi kinetics.

For the sake of our project, the third kind of term is the most important because the speed of these reactions depend on the maximum reaction velocity, which vary from person to person. In fact, on average, the reaction velocities between two people vary by 25%. In each of the three models that we started with, however, reaction velocities had been assumed to be constant. So that we could better model populations, we examined what would happen if we created a population of *virtual individuals* by varying all maximum reaction velocities uniformly by 25% to reflect population-level variation.

The first model comes from the paper “The biochemistry of acetaminophen hepatotoxicity and rescue: a mathematical model” by Ben-Shachar et al. The basic premise of the model is that, when it comes to APAP, the body can be thought of as four compartments: the gut, the liver, the plasma, and the muscle tissue (urine is considered to be a fifth compartment but it has no effect on the other compartments). In the model, a fixed amount of APAP enters the body through the gut, where it is transported to the liver. In the liver, APAP undergoes the processes of sulfation and glucuronidation, which produce harmless byproducts; these processes also happen at slower rates in the tissue. The liver is also where APAP is turned into N-acetyl-p-benzoquinone imine (NAPQI), which inhibits the functionality of liver cells when covalently bonded. If fewer than 30% of liver cells are functioning, the model considers somebody dead, which is a key fact because, given enough time, the model would otherwise predict that everyone will recover to full health regardless of the severity of the liver poisoning. Importantly, the liver is also where glutathione (GSH), a molecule that turns NAPQI into a harmless byproduct, is produced.

The first model was designed to work in conjunction with the second model, which comes from “A mathematical model of glutathione metabolism” by Reed et al. In fact, the first

model does not have any terms for GSH synthesis at all, opting instead to use the second model to describe how GSH is produced within the liver. The glutathione model is complicated and, largely, immaterial to the processing of APAP. However, a small portion of the model is extremely relevant to the project. This portion begins with the assumption that cells produce a constant amount of cystathionine, which is turned into the ingredient for GSH production cysteine. Cysteine is essential for GSH synthesis. In fact, the concentration of cysteine, which is in relative short supply within the body, is the only thing that is targeted by the current treatment of APAP poisoning, an infusion of N-acetylcysteine (NAC) into the body.

As important as cysteine is glutamate, which is relatively abundant in the body; cysteine must bind with glutamate before the body is able to synthesize GSH. Although GSH synthesis happens within the cytosol of liver cells, cysteine and glutamate also can be found in the bloodstream. In fact, cysteine, glutamate, and their derivatives can be found in low concentrations all throughout the body, but, within this second model, cysteine and glutamate are either in the blood or in the cytosol.

In order to more accurately capture the movement and transformation of cysteine and of glutamate in the body, we decided to use the model from “The role of skeletal muscle in liver glutathione metabolism during acetaminophen overdose” by Bilinsky et al. Like the previous model, this model starts with the assumption that a constant amount of cystathionine is being formed in the body. This cystathionine binds to glutamate, and GSH is produced in the exact same way. However, whereas the previous model only considers cysteine and glutamate, this model considers two derivative molecules: cystine and glutamine. In the blood, cysteine can reversibly change into cystine, the oxidized form of cysteine. Only two major facts differentiate cystine and cysteine: one, each molecule of cystine can only enter liver cells if a molecule of glutamate is exiting at the exact same time, but cysteine moves free into the liver; two, the only cells that cystine can enter are liver cells while cysteine can be found everywhere. Similarly, glutamine and glutamate are very similar molecules with one key difference: glutamine can enter the liver, but glutamate cannot. If glutamate from the blood is to enter the liver, it must turn into glutamine, which can only happen in muscle tissue and change back into glutamate.

The third model is very similar to the other two. Like the other models, this model considers the liver, the blood, and the tissue. Additionally, the model describes two processes that are part of the second model: glutamate synthesis and circulation. However, as similar as this model is, this model has one glaring difference from the others that we used: it was based upon rats. The biological processes involving glutamate in humans and in rats are identical, but there is no guarantee that the reaction velocities should be the same. To make the models work as a reasonable model of a human, we needed to be sure that the reaction velocities that we choose produced reasonable results. To find suitable numbers, we randomly varied each of the reaction velocities. The reaction velocities we settled upon (along with the original values) are available in the appendix. Since the first two models were designed to model humans, we varied these by up to 99%, a relatively conservative choice. We were much less certain about the proper values for the reaction velocities from the third model, so, to scan a wider set of parameters, we divided the reaction velocities by a random number (sampled from a lognormal distribution with mean 0.99 and standard deviation 1).

The strength and the weaknesses of the parameters that we chose are apparent from the subsequent figures. The top graph, which comes directly from the paper that outlines the first model, shows the effects of a dose of twenty-two gram APAP dose followed by thirty-six millimolar of NAC. The second graph conveys the same information for the full model.

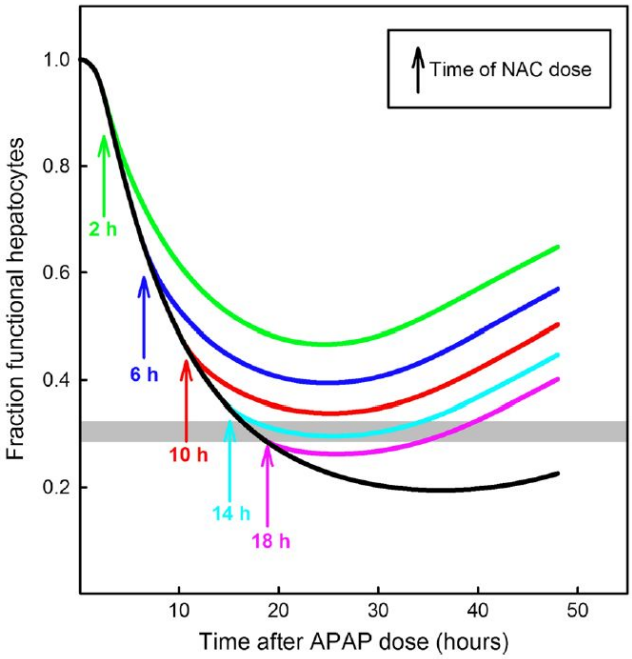

Credit: Ben-Shachar et al.

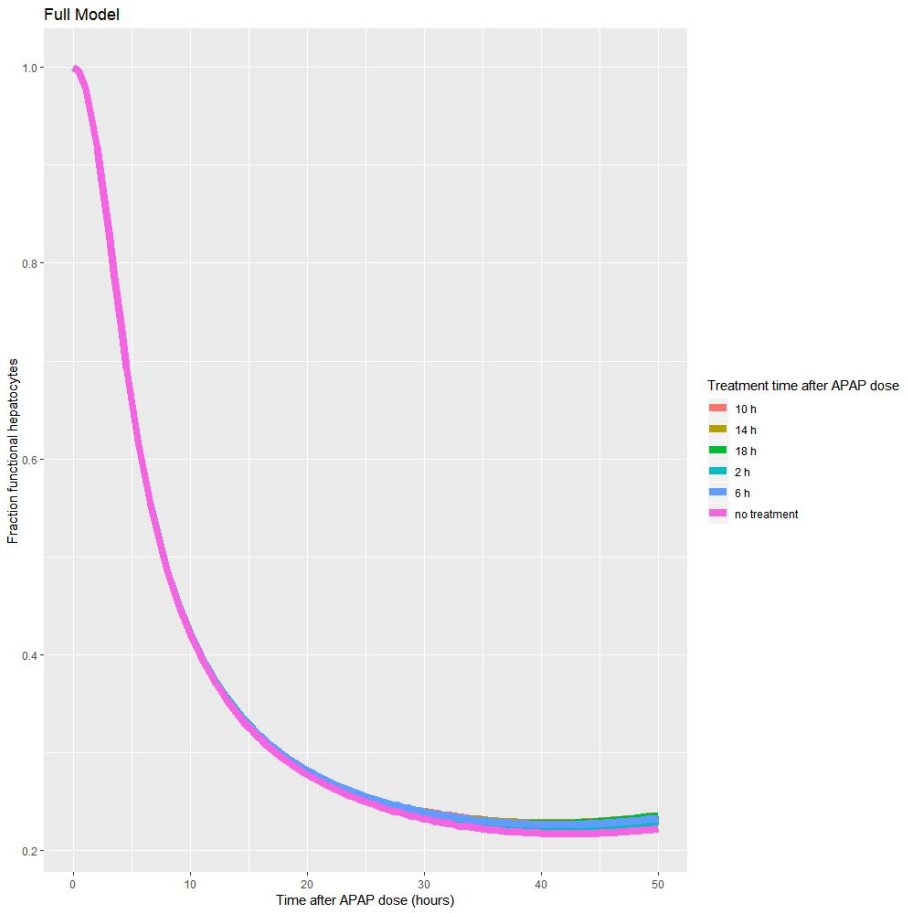

On the positive side, the shapes of the curves without treatment are almost identical: after a precipitous decrease, the curve flattens out. This demonstrates that, in both models, the rate of cellular regeneration eventually catches up to the rate of cell loss. Moreover, both curves cross the thirty percent functionality threshold, which is the dividing line between life and death, at approximately the same time.

Ideally, the rest of the lines would also look identical. However, that is not at all the case. The full model is much less responsive to an infusion of cysteine to the system than the small model is. Certainly, in the small model, all treatments are a little bit better than no treatments, there does not seem to be a big difference like there is in the small model. Without doubt, this is due to the parameters we chose. As demonstrated by the two subsequent graphs, a big difference between the small model and the full model is that GSH supplies are almost immediately exhausted in the small model, and that is not the case with the full model.

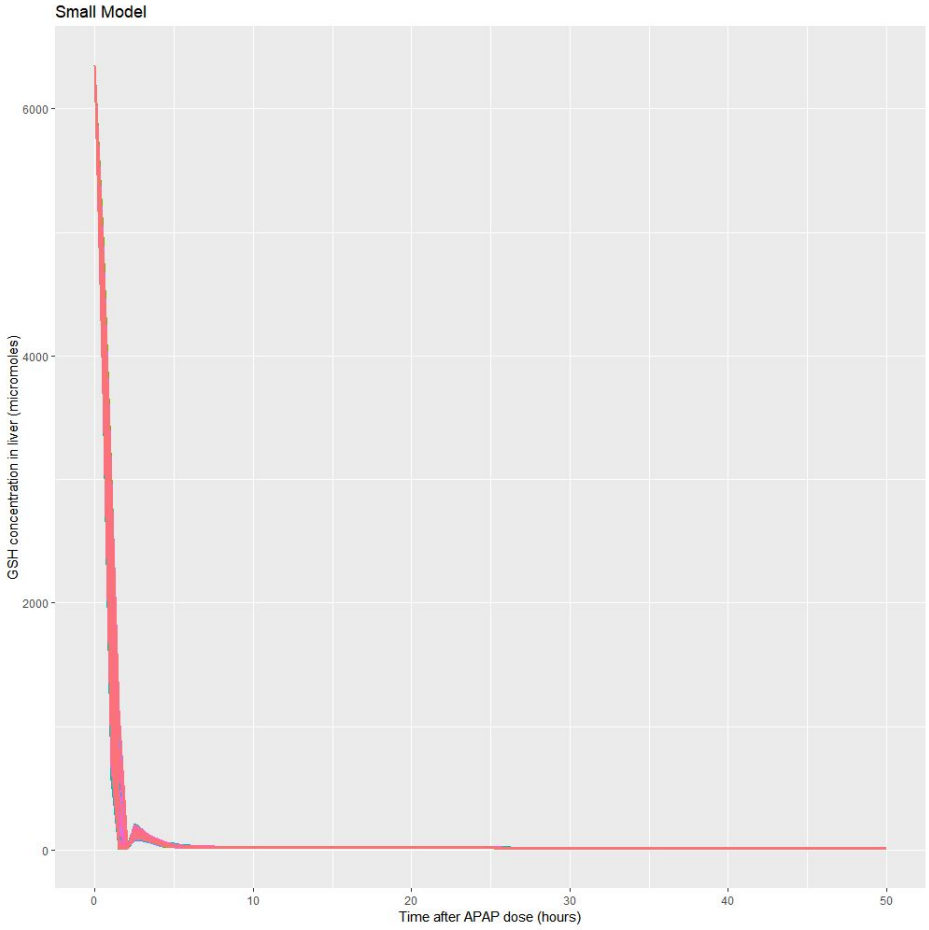
Small Model
GSH concentration in liver (micromoles)
6000
4000
2000
0
0
10
20
30
40
50
Time after APAP dose (hours)

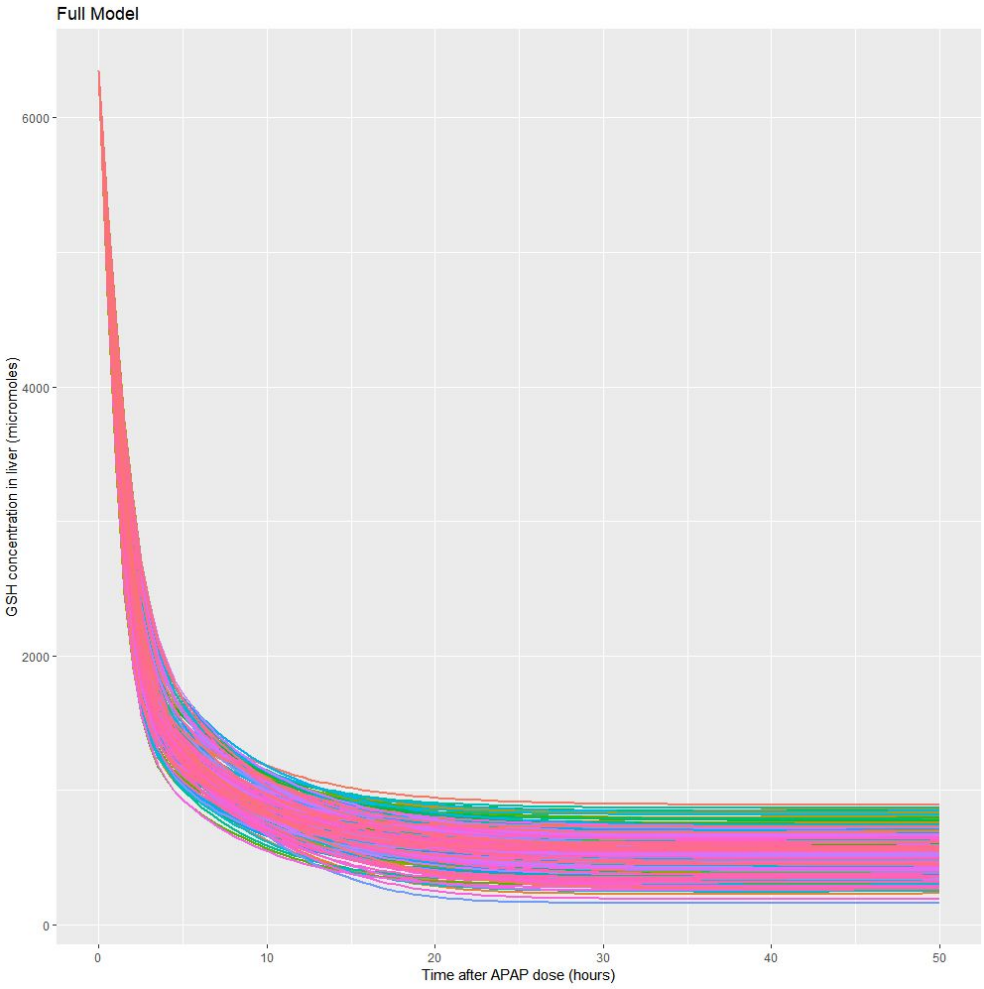

Each line represents the liver GSH concentration of a randomly generated virtual individual who was given twenty-two grams of acetaminophen. Within the small model, every virtual individual is practically out of acetaminophen by the fifth hour after taking the drug. This is not at all true for the full model. No virtual individual ever runs out of acetaminophen in the full model. Although there is a precipitous decline in the amount of GSH, the concentration eventually levels out around one millimolar.

The differences between the two models comes from the fact that, while the analysis was being done, there was no term in the full model to account for the fact that GSH is used up when it is bound to NAPQI while a corresponding term exists within the small model. This error has been corrected in the final code. If this does not explain all the difference, then some of that can likely be explained by the fact that maximum reaction velocity for the binding of GSH to NAPQI is smaller than that of the full model. Virtual individuals in the full model, therefore, have more GSH in reserves and, consequently, are less in need of cysteine. Moreover, since GSH cannot be used as efficiently, cysteine is just less useful in the full model. While not perfect, as far as we know, the full model is the only model that allows for the study of glutamate as a treatment for acetaminophen poisoning in humans. Therefore, we present results from both models.

Variation

The Population Model

Both the small model and the full model had three sets of parameters: initial cytosolic, blood and tissue concentrations of various chemicals; transport rates; reaction velocities. We set initial concentrations to be the normal values reported in both papers (Reed et al.; Bilinsky et al.). Transport rates are also assumed to be nearly constant among the population. However, the greatest variability comes from reaction velocities. For most of the velocities, we assume that their dependence on substrates has Michaelis-Menten form with one or two substrates. Under Michaelis-Menten, the maximum rate of reaction (Vmax) and the Michaelis constant (Km) are the two parameters which define the kinetic behavior of an enzyme. Km depends only on the structure of the enzyme and is independent of enzyme concentration. In our model Km is taken to be the same as reported in previous papers (Reed et al.; Bilinsky et al.). However, the maximum rate of reaction (Vmax) could vary from person to person, and a search through literature and the BRENDA online enzyme database confirmed our intuitions about such variation. And to build upon the existing models (Reed et al.; Bilinsky et al.) where kinetic parameter values for reaction velocities are fixed, we decided to explore the population model where all the Vmax's are set to be random with 25% variability.

We considered the variability of minimum proportion of living liver cells among the population at different initial acetaminophen (APAP) dosages. Both models show similar results. With less than 5g initial gut APAP dosage, almost the entire population survives. We start to see a bifurcation among the population at around 80,000-100,000 micromolar (10-16g) gut APAP levels.

We also implemented binary search to find the initial gut APAP level that leads to a minimum proportion of living liver cells of 29-31%. Under both the small and the full model, the distribution of such threshold gut APAP levels is unimodal, with a mode at around 90,000 micromolar (14 g).

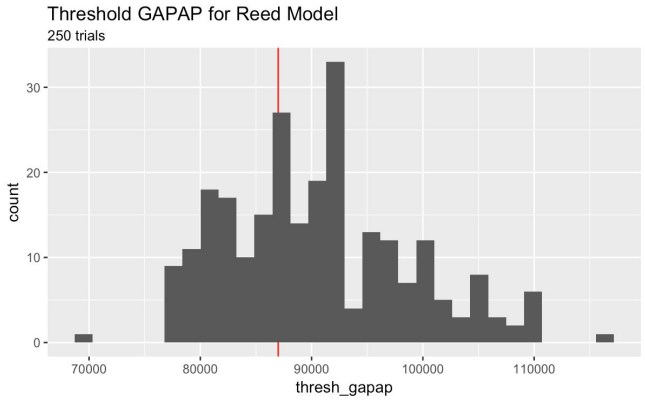
Threshold GAPAP for Reed Model
250 trials
count
thresh_gapap
0
10
20
30
70000
80000
90000
100000
110000

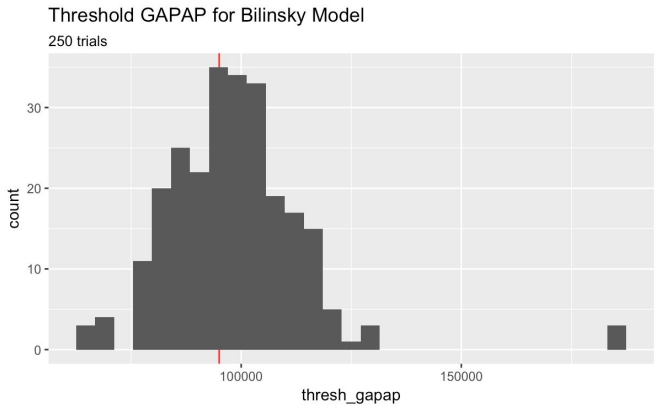
Threshold GAPAP for Bilinsky Model
250 trials
count
thresh_gapap
0
10
20
30
100000
150000

Small Model

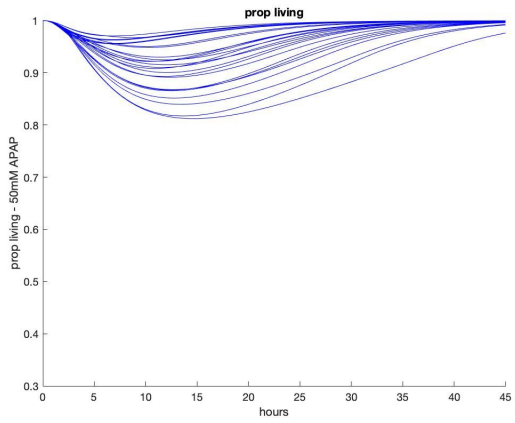
prop living
prop living - 50mM APAP
hours
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
5
10
15
20
25
30
35
40
45

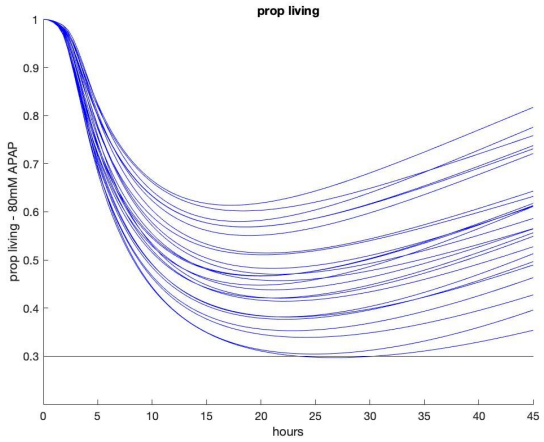
prop living
prop living - 80mM APAP
hours
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
5
10
15
20
25
30
35
40
45

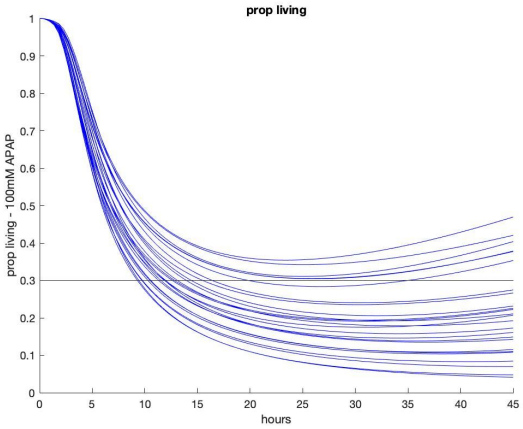
prop living
prop living - 100mM APAP
hours

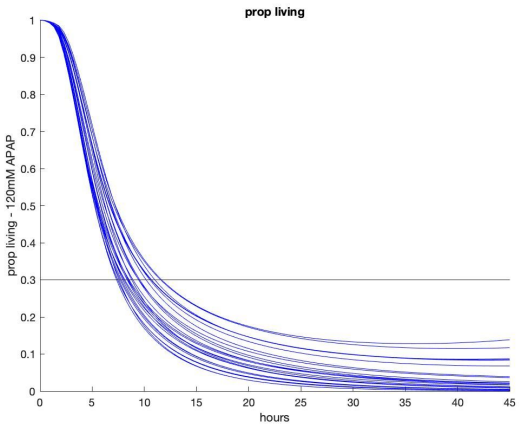
prop living
prop living - 120mM APAP
hours

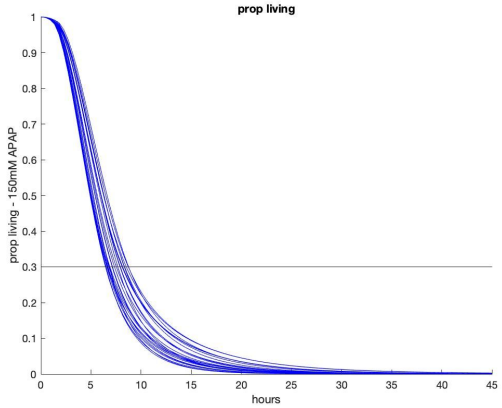
prop living
prop living - 150mM APAP
hours

Full Model

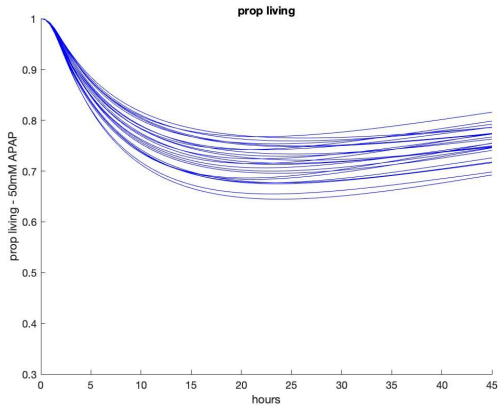

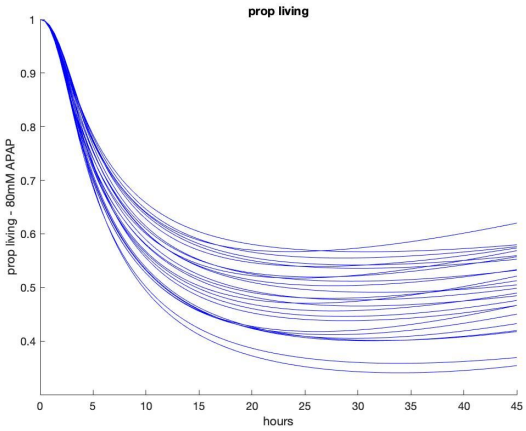

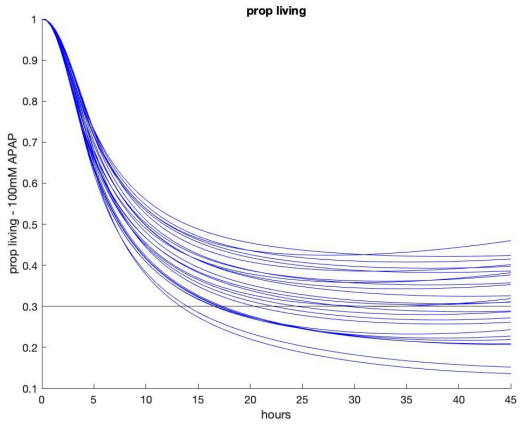

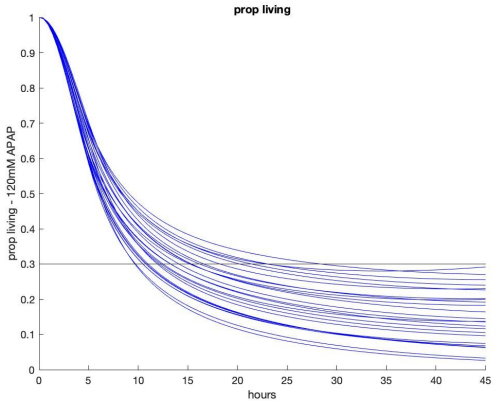

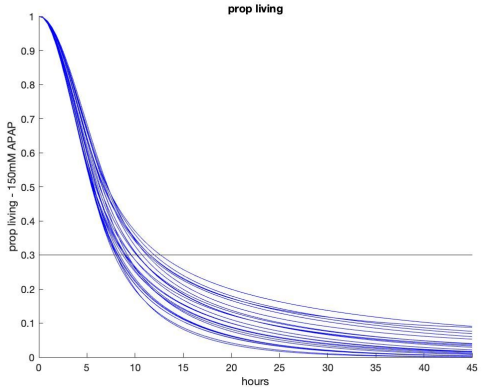

Principal Variables

Given such variability among the population, we wanted to dimension reduce in order to obtain a set of principal variables. We first estimated the Jacobian matrix at the point in parameter space where each of the coordinates are equal to the mean values of the random variables by bringing about small disturbances in each direction. This step gives a n x p matrix

where n is the number of simulated virtual individuals and p is the number of variables. We then decomposed the matrix using singular value decomposition. The largest singular values of this matrix parametrize the space, so we only look at singular values and singular vectors up to a preset accuracy (0.1).

Both models shared significant similarity in that the majority of the principal variables are from the original Tylenol model. These variables, i.e. VmaxCYP, VmaxlUGT, and VmaxGST explain how liver APAP affects NAPQI synthesis and how APAP is detoxified in the liver through glucoronidation and glutathione synthesis. Adding in more comprehensive models for liver glutathione metabolism (Reed) and glutamate metabolism (Bilinsky) seems to have only minor effects on the variability of the proportion of living cells among the population.

We also borrowed ideas from dimension reduction algorithms including locally linear embedding and diffusion maps (Coifman et al., 2005). In particular, diffusion maps use a distance metric, usually referred to as diffusion distance. Then the transition matrix for a Markov chain — specifically for a random walk on the graph of the data is constructed such that the random walk is biased towards making transitions that take it between similar points. And the eigenfunctions of Markov matrices can be used to construct coordinates of diffusion maps that generate efficient representations of the global geometry of the data. However, given the relatively simple structure of the manifold in our model, we think such algorithms might be overkill for the purpose of this project.

Reed Model

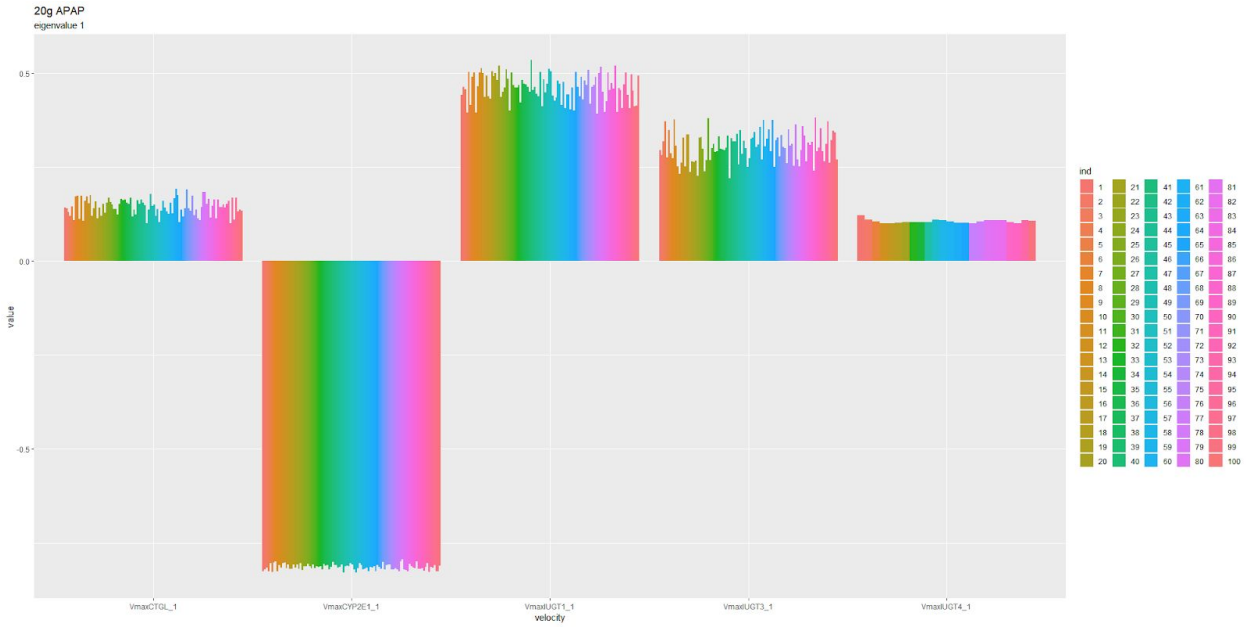

Bilinsky Model

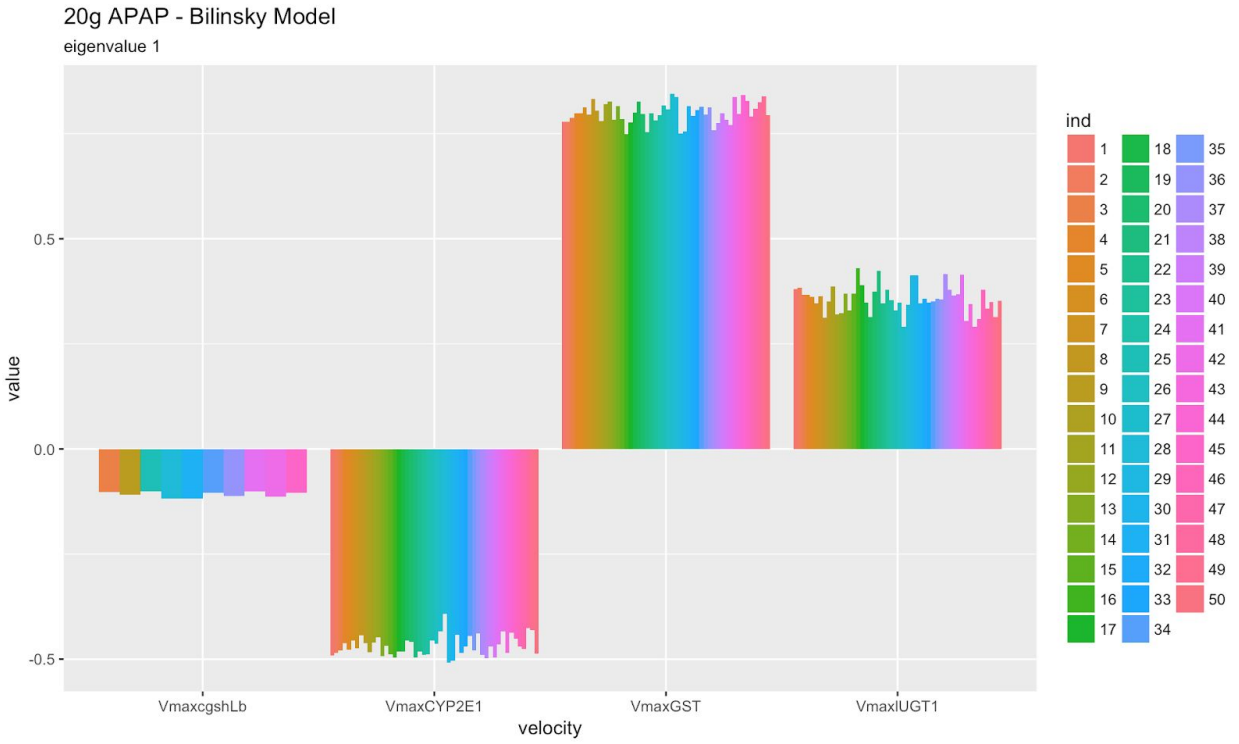

Holding the Principal Variables Constant

To verify the effect of the sets of key variables for both models, we held the key variables constant and varied the rest of Vmax's. As expected, the viability in simulation results reduced tremendously. For example, as shown for 20g gut APAP dosage under the small model, with all Vmax's random, the minimum proportion of living cells range from about 10% to 50%. However, if we hold VmaxCTGL, VmaxCYP2E1, VmaxlUGT1, VmaxlUGT3, and VmaxlUGT4 constant, the range shrinks to 40% to 42% for 2500 simulated trials.

Reed Model - 20g APAP

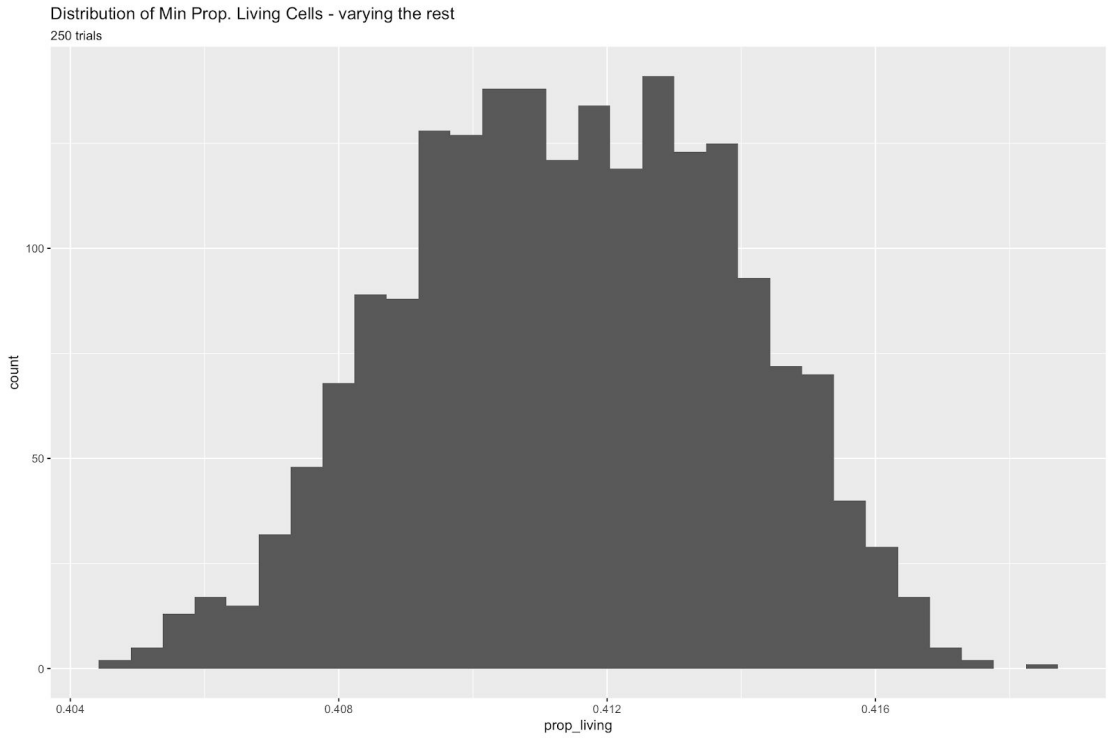

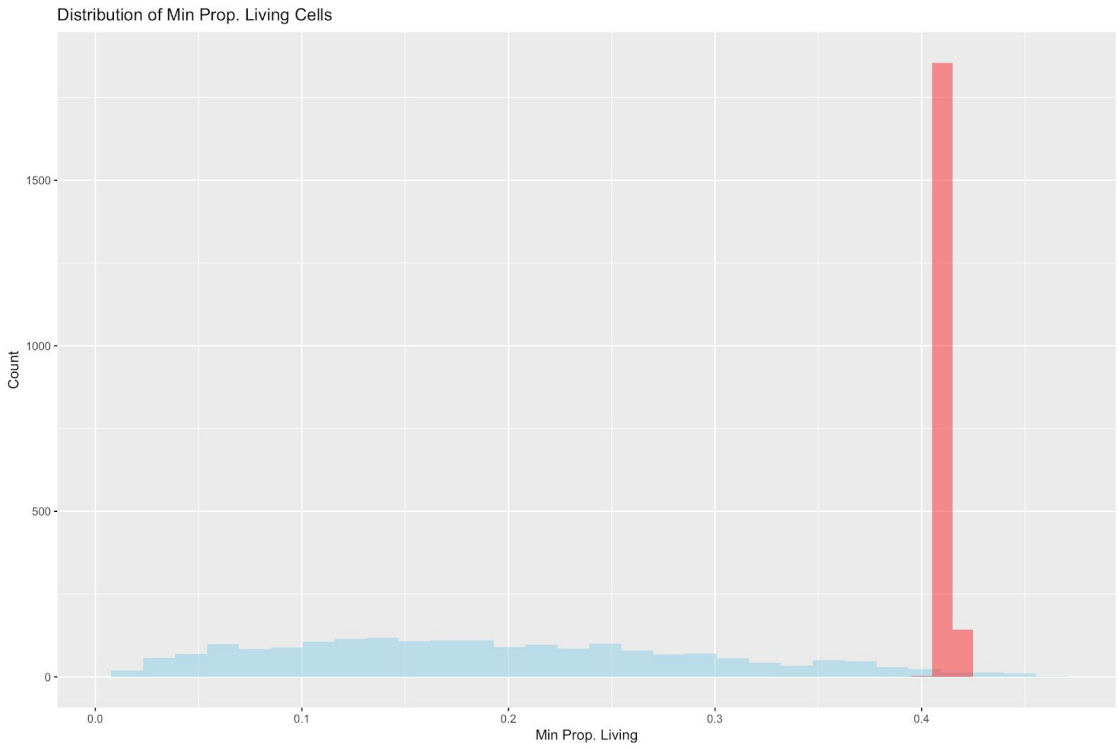

Bilinsky Model - 20g APAP

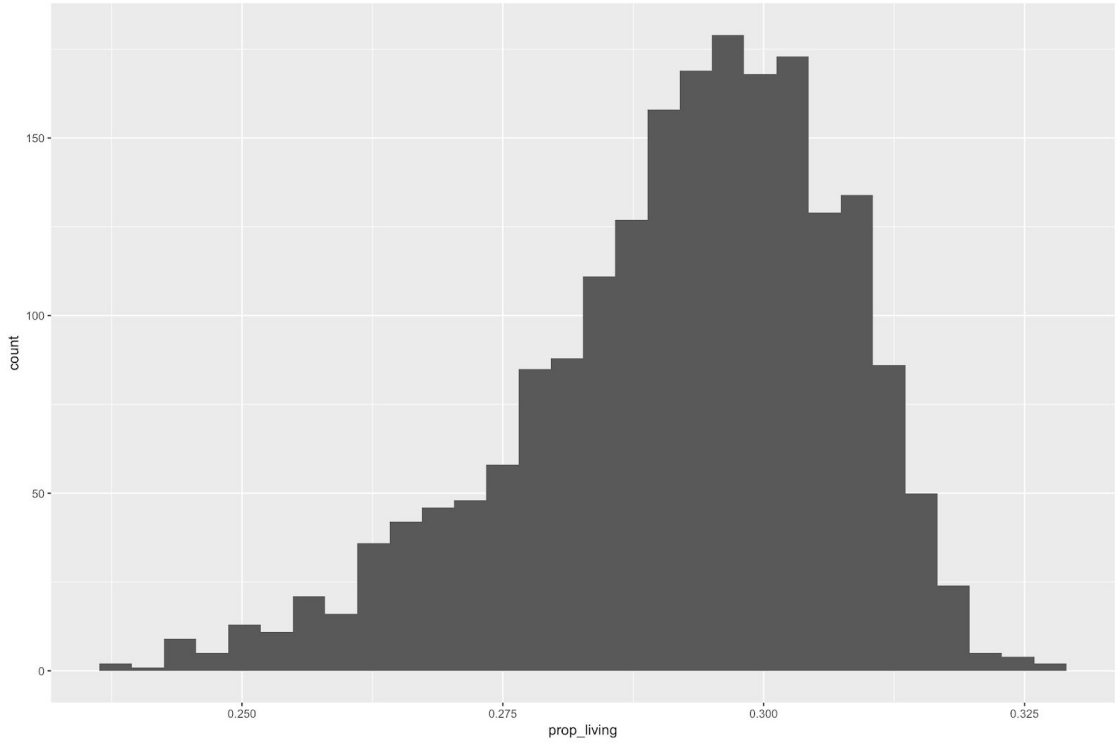

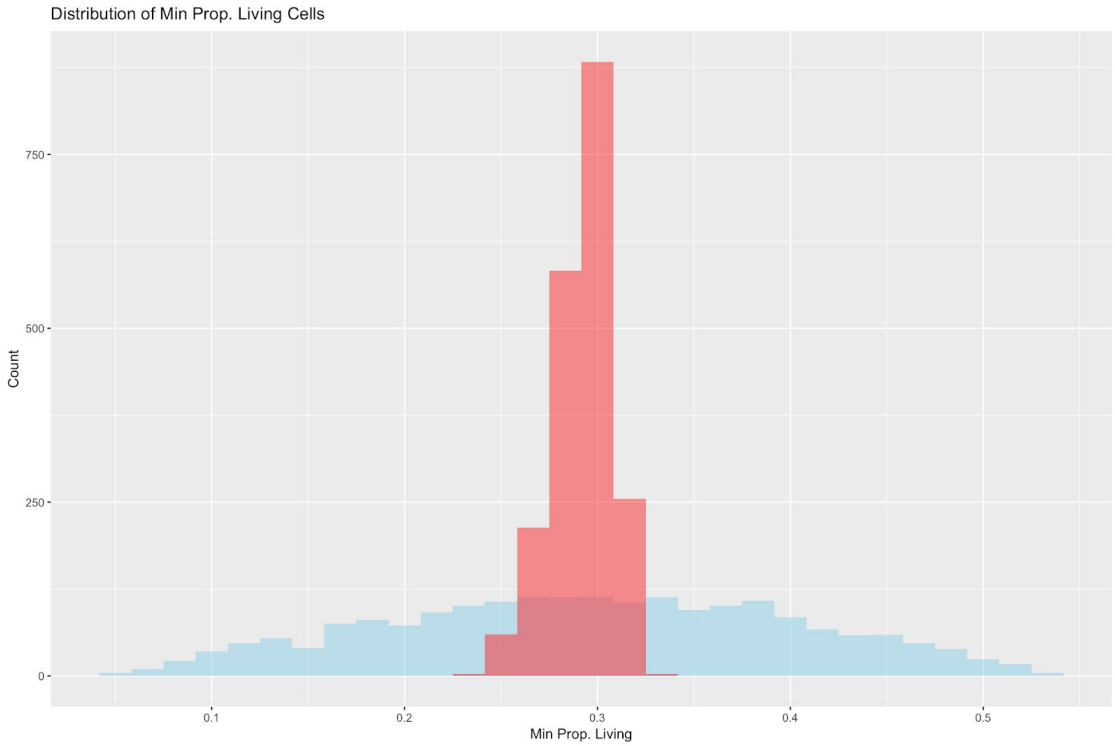

Optimization

The power of our research comes in part from the fact that we can try different treatments on a collection of random virtual individuals. For example, we were able to try out the treatment called “protocol 3” on over a hundred of random virtual individuals who were exposed to varying amounts of APAP. Specifically, protocol 3 requires giving half of the thirty-six millimolar NAC dose between two hours after APAP was taken to three hours after. Another six millimolar are given evenly over the next four hours. The last twelve millimolar of NAC is given over the following sixteen hours.

The effects can be seen below. Each box represents a mass of APAP that is taken. Each line is a different virtual individual (generated by varying reaction velocities in the small model). The importance of varying reaction velocities can be drawn out from the graph. For example, there is a large spread of results for virtual individuals who had been exposed to twenty grams of APAP. This variation, in fact, can account for the difference between life and death; some virtual individuals survive the APAP dosage while some succumb to its effects.

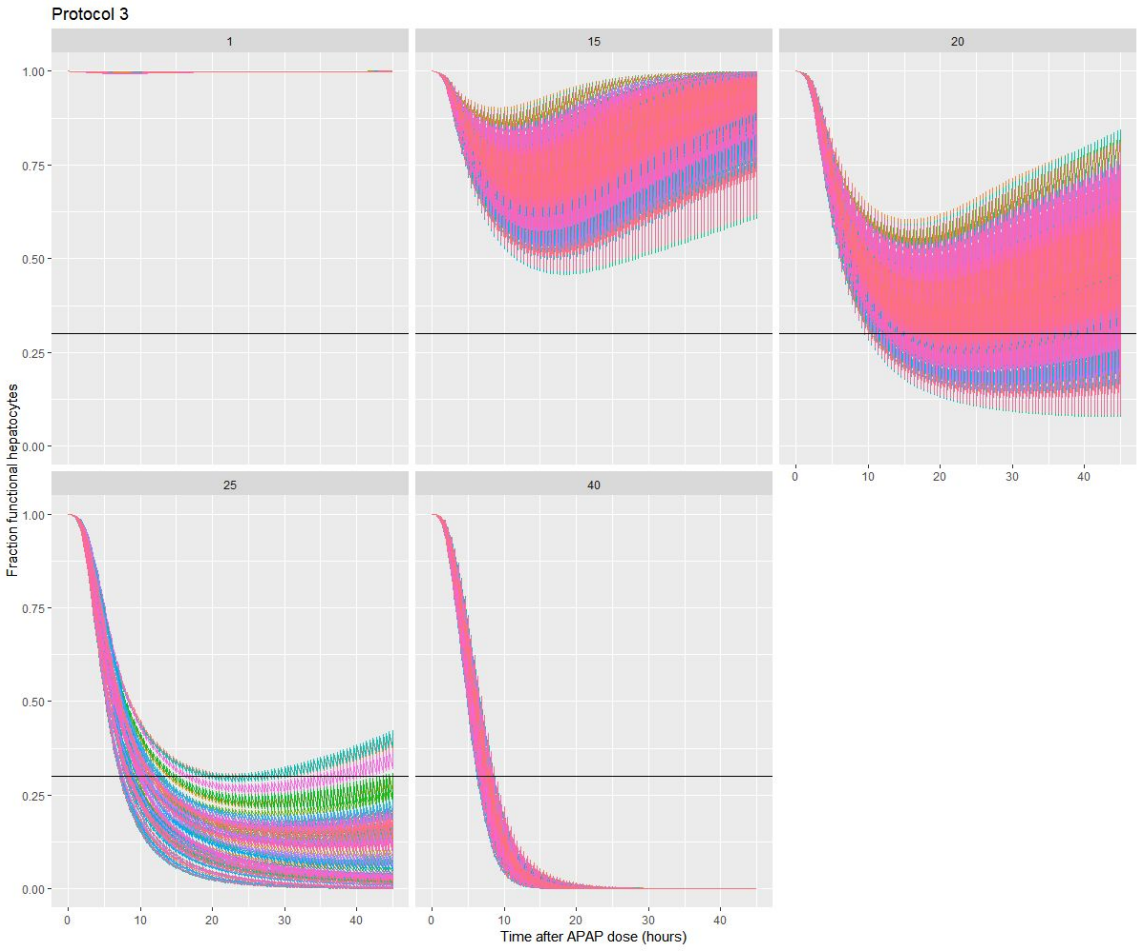

Since we can use the models to test how any given treatment affects many virtual individuals, we are also able to determine which treatment is optimal for a population at large. We did precisely this. For both models, we generated two hundred virtual individuals, which we exposed to twenty-two grams of APAP. Then, we allowed the model to change the proportion of NAC given (and, in the case of the full model, the proportion of the 5 millimolar of glutamate also given) through a random walk for at least one hundred steps, which did not really matter because the proportions converged really quickly. We took as a basis protocol 3, which Reed et al. established was the most effective treatment for APAP poisoning that they examined. We measured the number of patients living and the sum of the fraction of functional hepatocytes. If the number of living patients increased (which never happened) or if the sum of the minimum fraction of functional hepatocytes increased and the number of living patients remained the same, then we would direct the algorithm to remember those parameters. Otherwise, we told it to go back to the best-known treatment. The optimal treatments that we found are displayed in the following graph (see the appendix for details on dosing).

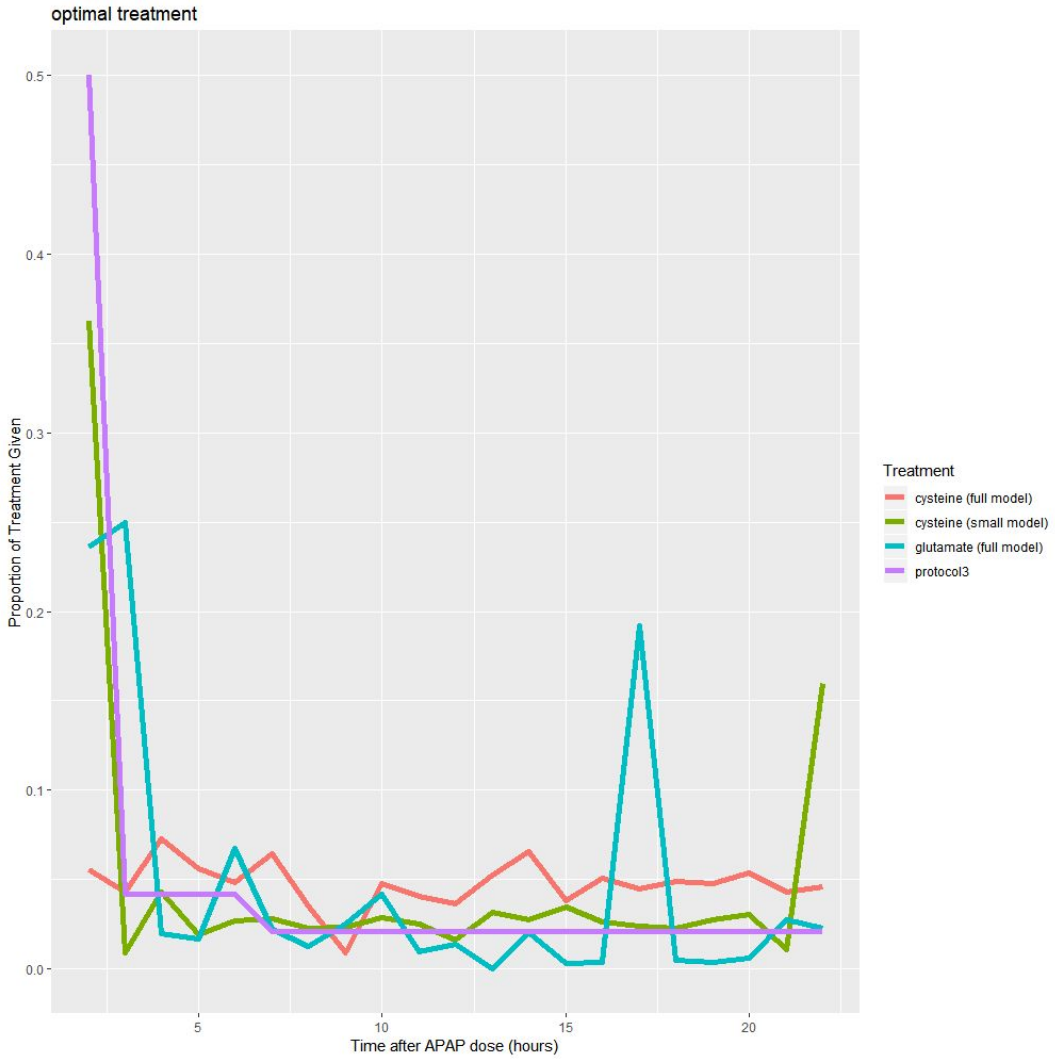

In many respects, the results are surprising. According to the full model, way too much cysteine is given in the first few hours in protocol 3. Instead, it should be given relatively consistently throughout the first twenty-four hours. There are many possible reasons the full model says this, but the most likely reason is that cysteine decays quadratically. Because cysteine is not in short supply in the full model, too much cysteine at any given time just leads to waste. In the small model, cysteine is in much shorter supply, so the concerns about loss to decay are less pressing. Still, the small model found that the optimal treatment involves giving less NAC immediately than protocol 3 suggests, opting instead to deliver an extra boost around hour 23. A similar phenomenon can be found in terms of glutamate. Similarly, according to the optimization process, protocol 3 allocates too much of the glutamate treatment to the very beginning, and a subsequent spike is crucial. Since the model is so complicated, it is hard to tell exactly what this true. However, this phenomenon is very interesting and might hold the secret to a better APAP treatment.

If there is a big takeaway, however, from the treatment optimization process, it could be that protocol 3 is nearly the optimal treatment. In the case of the small model, forty-seven of the two hundred survive with the optimal treatment and forty-seven of the two hundred survive with protocol 3. The average minimum percent of functional hepatocytes increased only from 42.5834 to 42.7174. In the case of the full model, the results are only slightly more pronounced. While sixty virtual individuals survive both treatments, the average minimum percent of functional hepatocytes has the larger increase from 45.6666 to 46.5707.

As a comparison, we can think about protocol 1, which involves giving all the NAC at once between the second and third hours (as would happen when somebody arrives in a hospital today) and, in the case of the full model, also giving all the glutamate over the same short time period. For the small model, if all the NAC and glutamate was given between the second and third hours, then thirty-one virtual individuals would have survived; the average minimum percent of functional hepatocytes would be 33.7772. Likewise, in the case of the full model, if protocol 1 is followed, then forty-eight virtual individuals would have survived; the average minimum percent of functional hepatocytes would be 41.2965. Moreover, if no glutamate was given at all (as it would not be in a hospital today), then forty-seven virtual individuals would have survived; the average minimum percent of functional hepatocytes would be 41.0813. These statistics are visualized in the following charts.

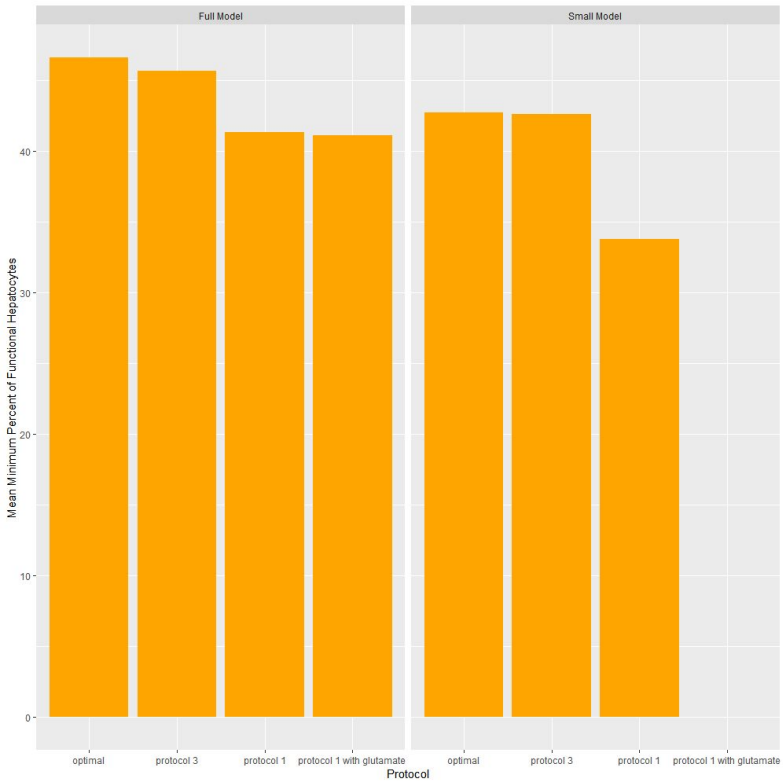
Full Model
Small Model
Mean Minimum Percent of Functional Hepatocytes
0
10
20
30
40
optimal
protocol 3
protocol 1
protocol 1 with glutamate
Protocol

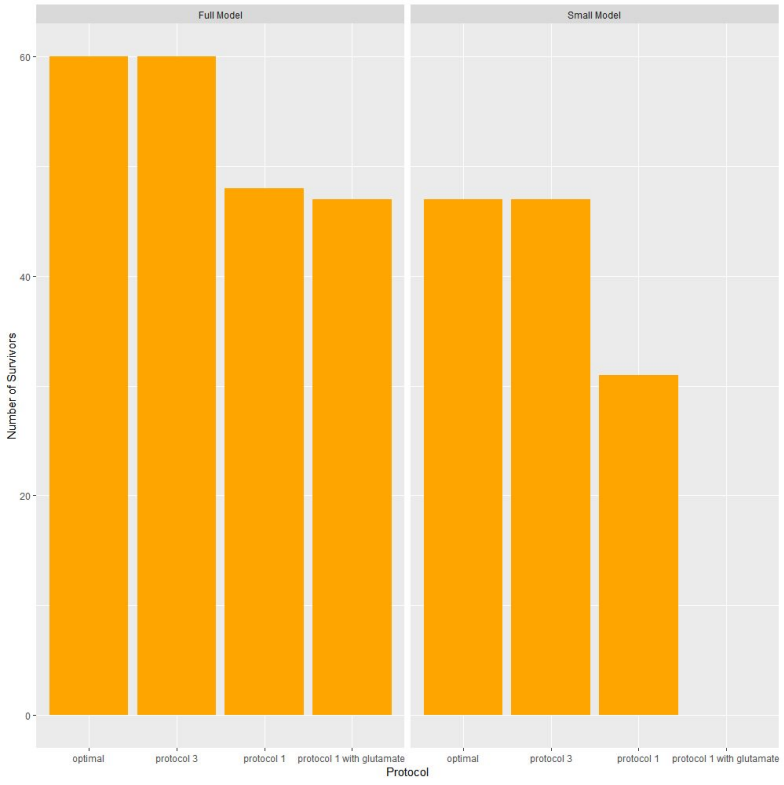
Full Model
Small Model
Number of Survivors
0
20
40
60
optimal
protocol 3
protocol 1
protocol 1 with glutamate
Protocol

References

[1] Ben-Shachar et al.: The biochemistry of acetaminophen hepatotoxicity and rescue: a mathematical model. *Theoretical Biology and Medical Modelling* 2012 9:55.

[2] Bilinsky et al.: The role of skeletal muscle in liver glutathione metabolism during acetaminophen overdose. *Journal of Theoretical Biology* 2015 118:133.

[3] Coifman and Lafon: Diffusion maps. *Applied and Computational Harmonic Analysis* 2006 5:30.

[4] Reed et al.: A mathematical model of glutathione metabolism. *Theoretical Biology and Medical Modelling* 2008, 5:8

Appendix

Reaction velocities

Maximum Velocity	Small Model	Full Model
VmaxCYP1A2	0.55	1.35
VmaxCYP2E1	345	215
VmaxCYP3A4	0.99	2.25
VmaxlSULT	1785	5.15
VmaxtSULT	357	26
Vlpaps	0.1	0.031
Vtpaps	0.01	0.015
VmaxGST	72000	1320
VmaxlUGT1	6370	11257
VmaxtUGT1	1274	3060
VmaxlUGT2	490	2
VmaxtUGT2	98	163
VmaxlUGT3	490	217.5
VmaxtUGT3	980	945
VmaxlUGT4	7350	673
VmaxtUGT4	1470	11.2
VmaxCTGL	1500	
VmaxGCS	3600	3430
VmaxbCYSc	14950	14750
VmaxbGLUc	28000	

VmaxGS	5400	1335
VmaxGPX	4500	1110
VmaxGR	8925	317
VmaxcgshLb	1005	1630
VmaxcgshHb	146	189
VmaxcgsgLb	4025	3160
VmaxcgsgHb	40	11.3
VmaxGLS		6240
Vmaxglutout		4170
Vmaxglyin		5750
VmaxGSsm		810
Vmaxglnin		190000
Vmaxglnout		13300
Vmaxsmglutin		8880
Vmaxsmglnin		395000
Vmaxsmglnouy		32270

Treatments

Hour	Protocol 3	Optimized cystine (small)	Optimized cystine (full)	Optimized glutamate (full)
2-3	1/2 (0.5)	0.3623	0.0554	0.2358
3-4	1/24 (0.0417)	0.0093	0.0428	0.2498
4-5	1/24	0.0430	0.0725	0.0199
5-6	1/24	0.0194	0.0562	0.0170
6-7	1/24	0.0267	0.0485	0.0674

7-8	1/48 (0.0208)	0.0279	0.0646	0.0219
8-9	1/48	0.0228	0.0354	0.0126
9-10	1/48	0.0233	0.0088	0.0254
10-11	1/48	0.0287	0.0476	0.0417
11-12	1/48	0.0254	0.0409	0.0094
12-13	1/48	0.0162	0.0362	0.0139
13-14	1/48	0.0315	0.0525	0.0001
14-15	1/48	0.0273	0.0658	0.0205
15-16	1/48	0.0346	0.0382	0.0030
16-17	1/48	0.0263	0.0509	0.0037
17-18	1/48	0.0242	0.0449	0.1918
18-19	1/48	0.0227	0.0491	0.0046
19-20	1/48	0.0273	0.0475	0.0037
20-21	1/48	0.0306	0.0536	0.0063
21-22	1/48	0.0106	0.0429	0.0063
22-23	1/48	0.1600	0.0457	0.0229