SERAPHIS PUBLICATIONS

# The Mathematics of Seraphis

Authorised by technical director Stephen Ward

April 16, 2026

## 1 DISCLAIMER

This is a mathematical document assuming undergraduate-level knowledge, however the tone is more informal and takes an engineering rather than a formalist approach. Making our framework accessible is important, however the highest priority in the early stages of any startup is operational. The public could anticipate, without promises, a layman's version in mid-to-late 2026.

## 2 THE MARKOV CHAIN

Our model has one axiom: that help-seeking is universal. This means that if it always exists somewhere, it only transfers between targets. Friends, the Internet, authorities, and others. Let those be states of a Markov chain. Since we aim to cooperate with authorities rather than replace them, we will treat it as an absorbing Markov chain. That is,

$$p_{\text{authority}\to\text{authority}} = 1$$

This is not strictly true in real life. A child abuse investigation might close after reaching authorities or end with improper action. However, our goal is to get the case to authorities. What happens afterwards is not part of our work, and so this is simply setting the scope of the model rather than a simplification.

This is not as difficult to measure as it sounds for three reasons. Most importantly, the process of reaching out for help under pressure is not exclusive to child abuse, and so can be measured without the ethical issues characteristic of abuse research. We can remain anchored to the reality of abuse due to large amounts of secondary data on abuse rates and reporting rates. Knowing the inputs and outputs of the Markov chain creates a strong boundary condition. And the third reason is that in the digital age, an increasing amount of help-seeking is visible. This does not just describe online help-seeking directly, but also other forms which are digitally logged or trackable.

This makes the Seraphis model of child abuse prevention arguably more feasible from a data collection perspective than other proposed forms of monitoring child abuse risk. Issues exist, and certainly measurement is not solved, however they are within the realm of engineering and privacy compliance rather than being unmeasurable. We have enough to create a minimum viable product (MVP) measurement, and enough to iterate on it to improve accuracy afterwards. This is how engineering begins.

A pilot method of modelling pathways is through a modified version of the Iowa Gambling Task (IGT). This isolates the ventromedial prefrontal cortex, responsible for decision-making under pressure, without retraumatisation or ethical risk. Although the IGT has not been used for this exact purpose before, the degrees of freedom allow the decks to represent different types of help-seeking. You could model, for example, low-risk choices with occasional large fines which make the deck net-negative, analogous to the grooming relationship of abuse; against winning decks perceived as very risky, analogous to the relationship with authority.

The limitations are clear - this does not give exact pathways, but categories of pathways, while missing some environmental factors. This is how pilot programmes are. It is, however, enough to sketch the space and effectively target some interventions. Mathematicians reading should be reminded of the large operational pressure to begin programmes and prove that you aren't just a sophisticated model. A moral pressure also exists, that if you have the capacity to help people, you are not expected to wait to meet some imagined threshold of model fidelity. We may opt to name our modified version with the word "Gambling" switched to "Choice", due to the undesirable optics of requesting access to children to introduce them to gambling, even if it is an Iowa Gambling Task.

## 3 Convergence

We then measure the rate of convergence. Let $n$ be the number of states. Recall that the steady state of an absorbing Markov chain is the absorbing state,

$$\overline{\pi} := \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}$$

We define the convergence rate to mean the speed at which a vector with all entries $1/n$ converges to $\overline{\pi}$. Exactly,

$$\frac{\partial}{\partial t} ||M^t - \overline{\pi}||$$

This is written for continuous-time Markov chains simply because $\partial$ notation is cleaner, but the generalisation to discrete-time chains is trivial. The rate of convergence is proportionate to the spectral gap, and since the largest eigenvalue of any stochastic matrix is 1 this means

$$\propto 1 - \lambda_2$$

Roughly, the probability of not being absorbed by a time $t$ is

$$e^{-(1-\lambda_2)t}$$

A rate of convergence can be yielded from this relatively easily. If one can also estimate a rate of decay (time before reaching out and number of failed attempts are two of the best-studied numbers in the field, you can interpolate), a rate at which a child stops reaching out for help as it hasn't arrived, you can create what we call the Seraphis ratio. In the case where

$$\frac{\text{rate of convergence}}{\text{rate of decay}} \geq 1$$

a child is able to receive help faster than they stop reaching out on average. This is confounded by the stochasticity of the system, but it is a meaningful goal.

Two advantages are worth naming, even if they aren't strictly mathematical. Firstly, pushing a number above 1 across the country is a goal the public can rally behind, it's a clear metric with an achievable goal. A metric which gets attention and resources beats a metric relegated to a blackboard. Secondly, 1 could represent a tipping point in several ways. A potential perpetrator knowing that it is easier for them to be caught than for them to hide abuse sees a strong deterrent. Systems begin to reinforce safety over time at this threshold as it is the point where safety begins to dominate over silence. Even if 1 does not represent the immediate end of child abuse, it could plausibly represent its point of terminal decline.

## 4 Targeting

With data, the problem becomes how to target transition probabilities. Some are more sensitive than others and therefore higher-yield to target. Our method is the Chambers algorithm, named out of respect for a member's ex-physics teacher Mr Samuel Chambers.

We take the Gershgorin disks of the Markov matrix. For a real stochastic matrix these are

$$G = \bigcup_i G_i = \bigcup_i [M_{ii} \pm \sum_{j \neq i} M_{ij}]$$

Gershgorin disks have the property that they collectively contain all eigenvalues of the matrix,

$$\forall \lambda, \lambda \in \bigcup_i G_i$$

If we aim to maximise the spectral gap $1 - \lambda_2$, we would like to minimise $\lambda_2$. We can do this indirectly by minimising the space that other eigenvalues could be. Therefore, a major strategy involves finding the $G_i$ containing the largest absolute value

$$j : \max \left| \bigcup_i G_i \right| \in G_j$$

and then minimising it. Depending on whether the largest value is in the positive or negative direction, the diagonal value $M_{jj}$ should be decreased or increased. This makes the matrix more diagonal over time by targeting the diagonal $M_{jj}$ elements.

Readers might say this is unnecessarily convoluted an algorithm, and ask why we don't simply use a computer to test which is the most sensitive. Our algorithm is not perfect and stands to be updated in the future. However, it is chosen here because it is resistant to inaccuracies. Our measured figures could have ±5% from reality, and a simple algorithm would have a chaotic effect where noise wildly changes the result. This level of risk is not acceptable when working in something as sensitive as child protection.

The Chambers algorithm minimises differences because Gershgorin disks create destructive interference between the random variables. Several random variables, each with roughly the same range which is equal in both the positive and negative direction, will on average have a lower total inaccuracy. Our judgement is that in the early stages, the accuracy range will be large enough that minimising the randomness with the Chambers algorithm will beat a technically more efficient algorithm that behaves chaotically under noise.

## 5 Interventions & Conclusion

Interventions can then be implemented and measured for accuracy. If measurement processes remain as streamlined as the IGT pilot as they become more accurate, we could conduct them regularly before and after interventions to measure effectiveness and progress over time. Otherwise, we intend to conduct them at least annually, but preferably several times per year, in select regions.