The Grief Seed Economy: An Agent-Based Model of the Magical Girl System in *Puella Magi Madoka Magica*

Yuna Hoshinomiya
yuna@yuna0x0.com
Independent researcher

29 July 2026

Abstract

We model the system connecting magical girls (girls, hereafter), witches and Kyubey in *Puella Magi Madoka Magica* as a resource cycle and study its behaviour at the population level. Every rule we impose is taken from the setting: cleansing moves corruption into a grief seed rather than destroying it, a filled seed hatches into a witch unless Kyubey consumes it, a defeated witch leaves a seed only sometimes, a clouded gem saps its owner's combat power, and familiars mature into witches of their own. The parameters the setting fixes are set from it, and the rest are calibrated against three things the setting shows: a city holds a handful of magical girls, veterans run a few years, and most careers end in transformation. Four results follow. First, at the density the setting shows, combat with witches dominates the corruption budget and conflict between girls is minor, a factor of 0.15; the ratio rises with density and passes parity only near nine standing girls, far above anything the setting depicts. Second, the rate at which girls leave is set by the contract rate and is insensitive to the seed supply, which instead changes how they leave and how much Kyubey collects. Third, cumulative return saturates in the offer rate. The advantage aggressive recruitment buys is temporary, because it comes from a stock that can only be spent once, which is the structure of a Ponzi scheme. Fourth, a city closed to migration has only two outcomes: the girls starve for want of seeds, or familiars multiply and the city is lost; the first is also where Kyubey's return is highest. Each result is a property of the model under the stated assumptions, and we validate the implementation against an exact conservation identity.

Keywords agent-based simulation · common-pool resources · population dynamics · evolutionary game theory · fictional world modelling

1 Introduction

This paper studies the subsystem of *Madoka Magica* that connects magical girls (girls, hereafter), witches and Kyubey. We formalise it as a resource cycle whose rules are taken from the setting:

1. A magical girl's soul is held in a soul gem, which clouds through the use of magic, through injury, and through despair, and clouds slowly even when she does nothing.
2. Cleansing does not destroy that corruption. A grief seed absorbs it, and a seed that fills up hatches into a witch unless Kyubey consumes it first.
3. A defeated witch leaves a grief seed only sometimes.
4. Familiars leave no seed, but one that feeds on enough humans matures into a witch of its own; and when a witch is defeated the familiars around it die with it. The baseline compresses this chain into a single maturation rate, and Section 4.7 implements it in full.
5. A gem that saturates breaks, and what reforms is a witch; the girl is dead at that point, and the change cannot be undone.
6. A clouded gem costs its owner her edge in a fight.
7. Kyubey consumes spent seeds so that the witch inside cannot return. The energy it collects comes from a girl's shift from hope to despair, which places the collection at the transformation.

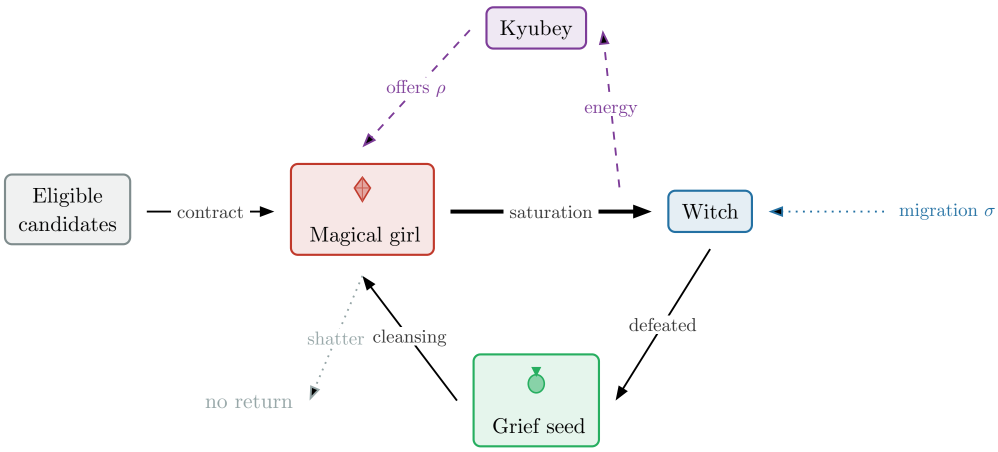

Figure 1: Schematic of the resource cycle. Solid arrows form the main loop; the purple dashed arrows are Kyubey's two roles, the contract and the harvest; the dotted arrows are the two boundary flows. One is witches migrating in from neighbouring cities. The other is girls lost to a shattered soul gem, which produces no witch and so returns no energy. The gem and seed icons are original shapes drawn for this figure.

Figure 1 draws these rules as a cycle. Written this way the model has the structure of a common-pool resource system, with one feature that the usual examples lack: the resource is not consumed and gone, it is transformed. Corruption leaves a girl only by moving into a seed, and a full seed becomes the next witch. The only exit is Kyubey. Systems of this general shape are well studied: fisheries bioeconomics [1], [2], [3] describes harvesting from a renewable stock, work on the commons [4], [5] describes what happens when access is shared, and evolutionary game theory [6], [7] describes contests over a divisible resource.

We then ask what the cycle does at the population level. How long do careers run? How many magical girls does a city support at once? Where does a girl's corruption budget go? What happens if a city is closed, so that its only witches are the ones its own girls become? Our interest is in behaviour that is hard to read off the rules directly.

The model contains no rule about territory, cooperation, or predation. Nothing in it distinguishes one girl from another except combat power, position, and a single aggression parameter. Several patterns nonetheless appear, and we spend most of the paper characterising them.

1.1 Contributions

- A spatial agent-based model of the cycle (Section 3) whose rules are drawn from the setting, with corruption conserved rather than discarded, probabilistic seed drops, familiar maturation, and a clouded gem that weakens its bearer.
- An exact conservation identity for corruption (Section 3.4), which decomposes a career into sources and sinks and doubles as a validation instrument.
- Four results (Section 4): combat dominates the corruption budget at the density the setting shows, while conflict dominates only far above it; the departure rate is set by recruitment; cumulative return saturates, so the system has the structure of a Ponzi scheme; and a closed city either starves or is overrun.
- Two policy studies that the setting itself raises: aggression under scarcity, and witch farming (Section 4.7), where familiars are modelled as individuals and the practice turns out to raise human losses without raising anyone's return.
- A discussion of what these imply for the system as a whole, and of how the model can serve as a rule core for narrative generation (Section 5).

2 Related work

Agent-based simulation. Deriving population-level regularities from local interaction rules is the standard method of agent-based social simulation, established in the Sugarscape tradition [8]. We follow that approach. More recently, language models have been used to drive agents that produce believable social behaviour [9]; we return to this in Section 5.

Bioeconomics. The Gordon–Schaefer model couples logistic stock growth to harvesting effort and identifies three reference points: maximum sustainable yield, maximum economic yield, and the open-access equilibrium at which economic rent is fully dissipated [1], [2], [3]. It is the closest standard model to the recruitment problem we set up, with one structural difference worth stating early. In a fishery the harvested stock and the renewable resource are the same population. In our model they are not: grief seeds are the renewable resource, while magical girls are both the harvested stock and the harvesting effort.

Contests. The hawk–dove game [6], [7] predicts a mixed equilibrium whenever the cost of escalation exceeds the value of the contested resource. We use a spatial version in which the model decides for itself how much a seed is worth: the less spare seed capacity the girls are carrying between them, the more a seed is worth taking, so the intensity of conflict tracks scarcity instead of being fixed in advance.

Formal descriptions of rule systems. Linear logic [10] suits resource-transforming rules because state change needs no frame axioms [11], and Ceptre [12] builds a rule specification language on it. Answer set programming has been used to describe generative design spaces declaratively [13], [14] and to analyse game systems bidirectionally [15]. These tools are strong at answering whether a state is legal and whether one state can be reached from another. Our questions are about how much and how many, which is why we use simulation instead.

3 Model

3.1 State

Time is discrete, one step per day. Space is a $G \times G$ toroidal grid of city districts, with $G = 6$. The state consists of a set of magical girls, a set of witches, a scalar pool P of eligible candidates, and Kyubey's cumulative energy return E.

A magical girl i has a position, a corruption level $c_i \in [0, 1]$, a base combat power w_i, a list of held grief seeds with their remaining capacities, an aggression parameter $a_i \in [0, 1]$, an age, and bookkeeping counters. Her effective power is $w_i(1 - \varphi c_i)$, so a clouded gem costs her the edge in a fight. A witch has a position and a power w.

3.2 Daily cycle

1. **Arrivals** the candidate pool grows by Λ up to a cap $P_{\max}$. Each standing witch matures a familiar into a new witch with probability f, and witches migrate in from neighbouring cities at rate σ, both drawn as Poisson counts.
2. **Drain and movement** each girl ages one day and accrues passive corruption α. With probability h she moves one district toward the nearest witch; otherwise she moves at random.
3. **Combat** in each district holding a witch, the girl present with the highest effective power engages it and accrues βw corruption. She wins with probability $w_i^{\text{eff}} / \big(w_i^{\text{eff}} + w\big)$, removing the witch and, with probability d, taking a grief seed of capacity κ. On a loss she accrues a further $0.8\beta w$, and her gem shatters with probability p_s.
4. **Contests** in each district holding two or more girls, a random pair plays hawk–dove. Girl i escalates with probability a_i. If both escalate, the loser pays C corruption and the winner $0.35C$, the loser's seed changes hands with probability 0.5, and the loser's gem shatters with probability $1.5p_s$. If only one escalates, she takes a seed from a non-escalating holder with probability $0.35V$. Here Σ_{cap} is the seed capacity the girls are still carrying unused between them, so the contested value $V = \big(1 + \Sigma_{\text{cap}}/M\big)^{-1}$ is highest when they are running short.
5. **Cleansing** while c_i exceeds a trigger c_{trig} and a seed is held, corruption is drawn down toward c_{targ} against the seed's remaining capacity. The corruption is not destroyed; it now sits in the seed. A seed worked to exhaustion is taken by Kyubey with probability q, and otherwise hatches into a witch.
6. **Transformation** a girl with $c_i \geq 1$ becomes a witch of power λw_i^{eff}, and Kyubey's return increases by one unit.
7. **Contracts** Kyubey makes Poisson(ρ) offers per day. Each is accepted with probability π, consumes one unit of the candidate pool, and returns η units of energy. A new girl's aggression is set in one of two ways: from a fixed value, or by imitating a girl whose career recently ended, with the longer-lived ones more likely to be copied. The second is success-biased transmission [16], where learners favour models who appear to have done well.

3.3 Parameters

Table 1 lists the parameters. Those in the upper block are fixed by the setting or calibrated against it; those in the lower block are ones the setting does not quantify, and we treat them as a design space to explore rather than as estimates, in the sense of [13]. Absolute magnitudes such as days and head counts should be read as internal units. The calibration and the sources behind it are recorded in the project's `docs/CANON.md`.

Symbol	Value	Meaning
α	0.010	passive corruption per day
κ	0.9	grief seed capacity, three cleansings
d	0.5	chance a defeated witch leaves a seed
q	0.5	chance Kyubey takes a spent seed before it hatches
φ	0.6	how far corruption saps combat power
f	0.005	familiar maturation per witch per day
σ	0.10	witches migrating in per day
ρ	0.008	contract offers per day
p_s	0.002	gem shatter probability per lost engagement
β	0.030	corruption per unit of witch power engaged
$c_{\mathrm{trig}}, c_{\mathrm{targ}}$	$0.55, 0.25$	cleansing trigger and target
λ	1.6	power multiplier on transformation
G	6	grid side, in districts
h	0.75	probability of moving toward a witch
δ	0.004	witch dormancy rate per day
$\Lambda, P_{\mathrm{max}}$	$0.05, 20$	candidate inflow and pool cap
π	0.55	offer acceptance probability
C	0.11	corruption cost of an escalated contest
a	0.25 or swept	aggression
η	see text	energy taken at the contract

Table 1: Model parameters. The upper block is set from the setting or calibrated against it; the lower block holds parameters the setting does not quantify.

3.4 Corruption accounting

Corruption within a girl is conserved. It is generated, moved into a seed, or carried to the end of her career, and there is no other channel. Every career therefore satisfies an exact identity. Consider a girl who ends her career after L days. Over that career she moves P units into seeds, takes C_{cb} corruption in combat with witches and C_{cf} in contests, and finishes with gem load c_{end}. She satisfies

$$\alpha L + C_{\mathrm{cb}} + C_{\mathrm{cf}} = c_{\mathrm{end}} + P. \tag{1}$$

Rearranged, career length is a corruption budget divided by the passive drain:

$$L = (c_{\mathrm{end}} + P - C_{\mathrm{cb}} - C_{\mathrm{cf}})/\alpha. \tag{2}$$

Two things follow before any simulation. A girl who never fights, never contests and never cleanses has a career of $1/\alpha = 100$ days, which serves as the model's natural unit of lifespan. Cleansing and conflict also enter Equation 2 with opposite signs, so what sets longevity is the balance between them rather than either one alone.

The identity is local to a girl. Across the system corruption is conserved too, but by a longer route: what leaves a girl enters a seed, and a full seed becomes a witch, whose defeat may return a seed to another girl. The only place corruption leaves the system is Kyubey's consumption of spent seeds.

In steady state, Little's law [17] relates the standing population to the contract rate $R = \rho\pi$ and the mean career length, $M = R\,\overline{L}$. Note that the mean career lengths reported here cover completed careers only and are therefore biased downward by right-censoring, so substituting them into this relation understates M. Kyubey's return arrives at

$$\dot{E} = R(\eta + f_w), \tag{3}$$

where f_w is the share of departures that occur by transformation rather than by shattering. Here η is the energy taken at the contract. The setting describes the energy as coming from the shift between hope and despair, which places the collection at the transformation; whether a share is also taken at the contract itself is reported differently by different sources, and we cannot settle it. We therefore report η as a range that includes taking nothing at the contract: the main runs use $\eta = 0.2$, and Section 4 gives the return for $\eta \in [0, 0.5]$.

4 Results

Runs are 8000 days. Steady-state quantities average the last 2000 days; career statistics pool completed careers across replicates. Unless stated otherwise the parameters are those of Table 1, giving $R = 0.0044$ contracts per day.

4.1 Baseline

Figure 2 shows a run at the calibrated parameters. The city holds $M = 2.8$ girls and $W = 0.23$ standing witches. Median career length is 158 days and the ninetieth percentile is 877 days, so most girls last months while a few run for years. That heavy tail is one of the calibration targets rather than an independent check on it: Mami Tomoe, the veteran the setting introduces first, contracted years before the story begins (see `docs/CANON.md`, section 4). 87.1% of departures are transformations. A girl consumes 4.6 seeds over her career and defeats 0.29 witches a week. Over the run Kyubey took 135 spent seeds out of circulation while 115 hatched into new witches.

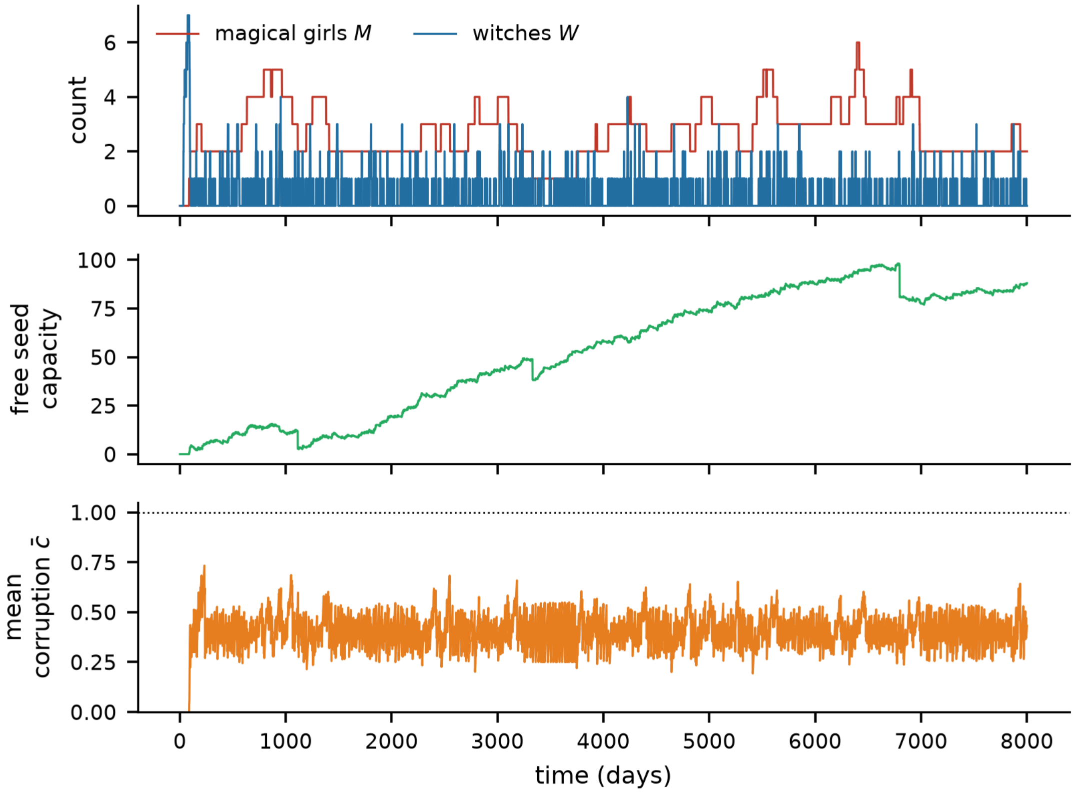

Figure 2: Baseline run at the calibrated parameters. Top: magical girls and standing witches. Middle: total unused grief seed capacity held by the population. Bottom: mean gem corruption, with the transformation threshold dotted.

4.2 Result 1: combat dominates the budget at the density the setting shows

Decomposing Equation 1 across the migration rate (Figure 3, Table 2) gives the first result. At the calibrated point, corruption taken fighting witches exceeds corruption taken from other girls by a factor of 6.7; the infighting share is 0.15 of the combat share. The setting does depict girls contesting hunting grounds: Kyoko Sakura arrives from Kazamino to take over Mitakihara after Mami Tomoe's death (see `docs/CANON.md`, section 4). But at the calibrated density those contests contribute little to the budget. The reason is density: with fewer than three girls in thirty-six districts, two of them rarely stand in the same place.

Sweeping the offer rate to raise the standing population (Figure 4) shows how crowded a city has to be before contests take over instead. The ratio climbs monotonically with density: it is 0.61 at seven standing girls and 1.08 near nine, so parity falls between the two, an order of magnitude above the one to five girls the setting depicts. Conflict dominance is therefore a property of a crowded city, not of this system as such.

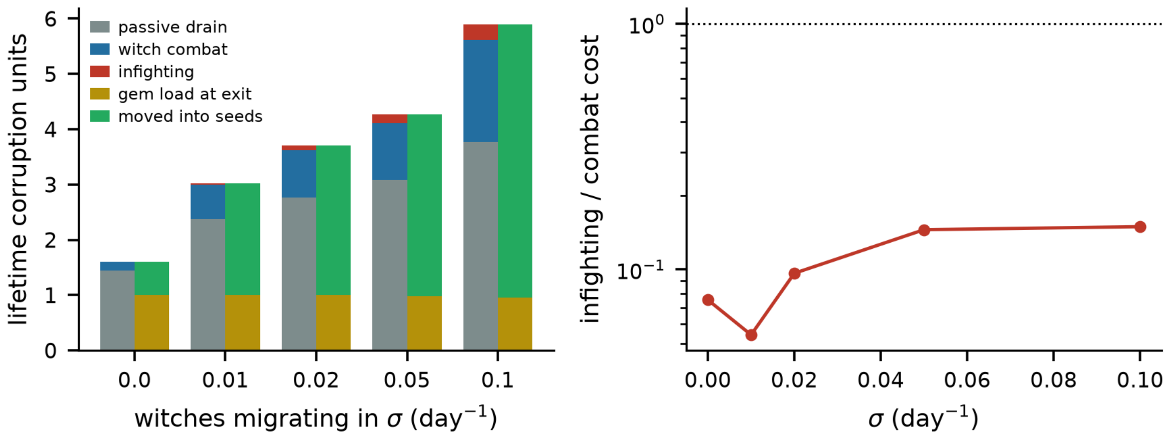

Figure 3: Corruption budget. Left: lifetime sources (left bar: passive drain, witch combat, conflict between girls) against sinks (right bar: gem load at the end of a career, corruption moved into seeds). Right: ratio of conflict cost to witch combat cost, with parity dotted.

σ	$\overline{L}$	passive	combat	conflict	into seeds	M	f_w
0.00	145	1.45	0.15	0.011	0.61	0.5	0.995
0.01	237	2.37	0.62	0.034	2.02	1.1	0.993
0.02	276	2.76	0.86	0.083	2.71	1.7	0.981
0.05	308	3.08	1.04	0.151	3.28	2.2	0.948
0.10	376	3.76	1.85	0.276	4.94	3.1	0.891

Table 2: Lifetime corruption budget in corruption units, mean career length in days, standing population, and transformation share, against the rate at which witches migrate in. Six replicates per row, $n = 165$ to 214 careers.

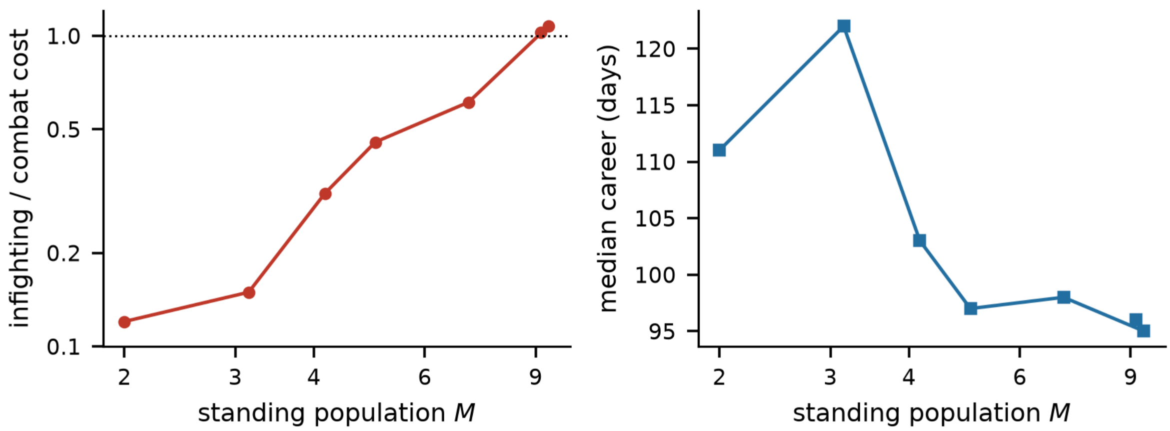

Figure 4: Left: the ratio of infighting to combat cost against standing population, with parity dotted. Right: median career against standing population. Density is varied through the offer rate.

4.3 Result 2: departures track recruitment, not the seed supply

Across the migration range the rate at which girls leave the population stays between 0.0034 and 0.0045 per day, tracking the contract rate $R = 0.0044$ as Equation 3 requires (Figure 5, centre). Enriching the seed supply changes not how many girls end their careers but how they end and how long they last: the transformation share falls from 99.5% to 89.1%, while the ninetieth percentile of career length rises from 256 to 1054 days. Kyubey's return falls with the transformation share, from 0.0050 to 0.0038 units per day.

The direction runs against intuition: the more witches a city has, the worse it is for Kyubey. Not because girls then survive indefinitely, but because a larger share of them are lost to a shattered gem or to a contest, and neither death leaves a witch behind, so Kyubey collects no despair from it. A richer seed supply does not keep everyone alive longer. It keeps a few alive much longer.

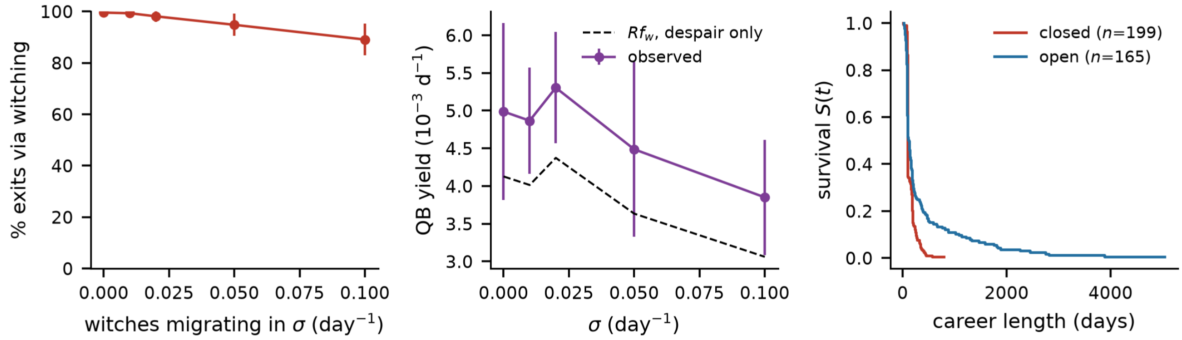

Figure 5: Left: transformation share against the migration rate. Centre: Kyubey's return, with the despair-only prediction Rf_w from Equation 3 overlaid; the gap is the energy taken at the contract. Right: career survival curves for a closed city and one at the highest migration rate.

The full (ρ, σ) surface (Figure 6) shows the return rising with the offer rate throughout and depending only weakly on migration. At the lowest offer rate the transformation share is at its most sensitive: in a closed city it is 1.00, and every girl ends as a witch.

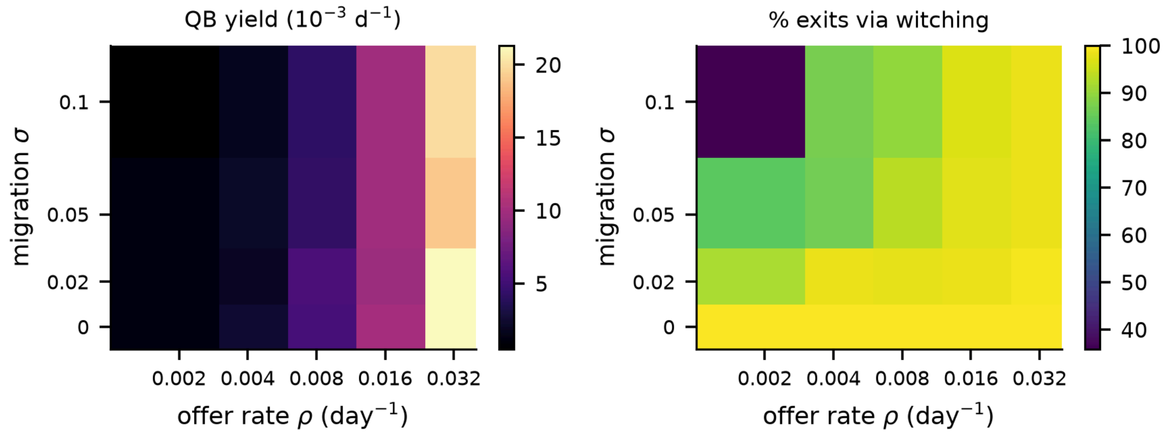

Figure 6: Kyubey's return (left) and transformation share (right) over offer rate and migration rate. Four replicates per cell.

4.4 Result 3: cumulative return saturates, and the advantage of haste is temporary

This set of runs restricts the candidate pool, setting $\Lambda = 0.004$ and $P_{\max} = 12$, and sweeps the offer rate over a sixty-four-fold range (Figure 7). Cumulative return saturates: past $\rho \approx$

0.016 the total after 8000 days sits between 45.0 and 46.9 units however aggressively Kyubey contracts, while at $\rho = 0.004$ it reaches only 13.8.

The reason is that a high offer rate converts the standing pool of eligible candidates into energy quickly, but that pool is consumed once; afterwards the return rate is limited by Λ, the rate at which new candidates become eligible. At $T = 500$ days the best rate tested is the highest one, 0.256, and by $T = 8000$ the choice hardly matters. This is the structure of a Ponzi scheme: the early return is paid out of a stock consumed once rather than out of a sustainable flow, and only a continuing supply of new candidates keeps the high rate going.

The same arithmetic can be read off the setting without simulating anything. A girl yields at most one seed, and only if she becomes a witch and someone defeats her; once the transformation share and the drop chance are counted, expected production is 0.44 seeds. A career needs about fourteen cleansings, which is 4.6 seeds. Expected production falls short of consumption by an order of magnitude, so the population cannot sustain itself and the gap must be filled by recruitment.

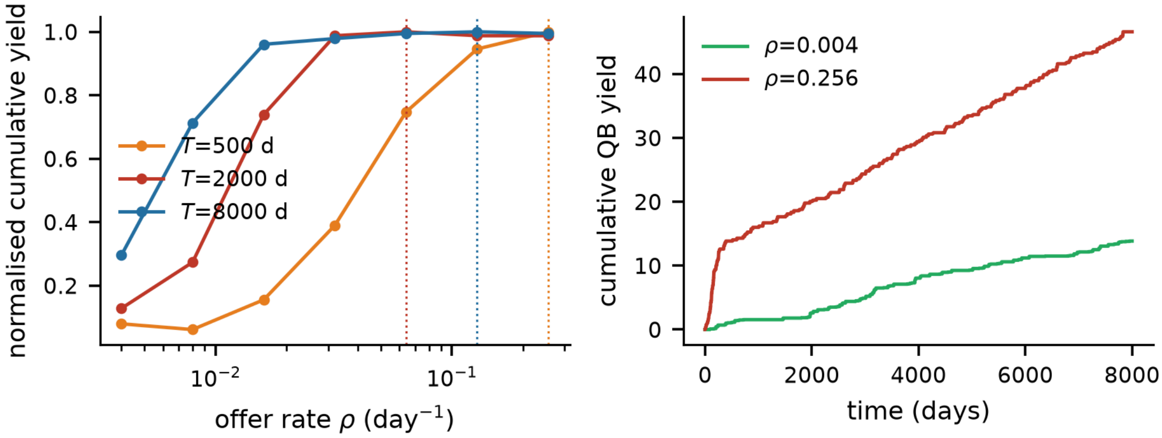

Figure 7: Left: normalised cumulative return against offer rate for three evaluation horizons, with the maximum of each marked. Right: return trajectories for the lowest and highest offer rates under a restricted candidate pool.

Because the setting settles neither whether energy is taken at the contract nor how much, we report the return against it. Over the baseline run, 27 transformations and 33 contracts give a total of 27.0, 30.3, 33.6 and 43.5 units at $\eta = 0$, 0.1, 0.2 and 0.5. The qualitative results do not depend on the choice: η adds a term proportional to recruitment, which strengthens rather than weakens the case for treating recruitment as Kyubey's control variable.

4.5 Result 4: a closed city starves or is overrun

If no witches migrate in, the only witches a city sees are the ones its own girls become and the ones its seeds hatch. Figure 8 sweeps the familiar maturation rate under that condition and finds no sustainable middle.

Below $f = 0.005$ the city starves. The standing population sits near half a girl, the median career is exactly 100 days, which is $1/\alpha$, and a career moves 0.4 to 0.6 units of corruption into seeds, so cleansing barely happens at all. Every departure is a transformation: $f_w = 1.00$. This side is not overrun, though: a girl becomes at most one witch, the contract rate is bounded, and witches are either defeated or go dormant, so the standing witch count settles near 0.40. The girls fare worst here while the city itself is never swamped. What separates the two outcomes is whether witch supply is linear or compounding: on this side each girl contributes one, on the other witches beget witches through their familiars.

The side that starves is also the one that suits Kyubey best. The return rate is 0.0055 units per day at $f = 0.001$, the highest anywhere in the closed-city sweep, and falls monotonically with the maturation rate to 0.0035 at $f = 0.02$. The regime worst for the girls is the most efficient one for Kyubey: when every departure is a transformation, no despair goes to waste.

At $f = 0.01$ and above the balance tips the other way. Witches accumulate faster than a handful of girls can clear them, reaching 45.9 standing witches at $f = 0.01$ and 173.5 at $f = 0.02$, and the city is lost in 33% and 100% of runs respectively.

The window in which a closed city looks anything like the setting is therefore narrow, and the setting itself tells us the city is not closed: Kyubey works worldwide, and girls travel between cities to take over vacant hunting grounds. What the closed case establishes is the direction of the imbalance. Left to itself, this economy does not reach a quiet equilibrium; it either fails to feed its girls or fails to contain its witches.

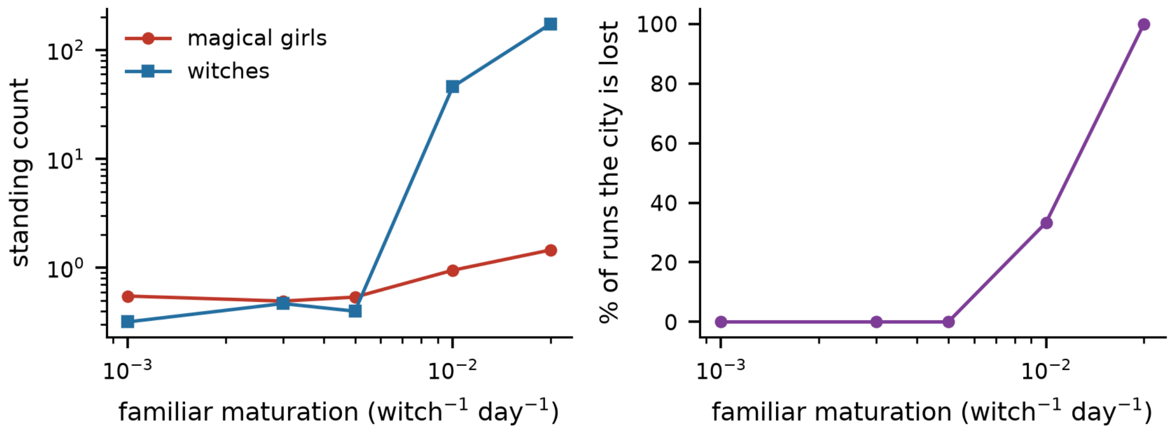

Figure 8: A city closed to migration. Left: standing counts of girls and witches against the familiar maturation rate. Right: the share of runs in which witches pass the cap and the city is counted as lost. Six replicates per point.

4.6 Aggression under scarcity

Fixing the aggression parameter so that populations are uniform (Figure 9, left) shows a cost to universal aggression where seeds are plentiful: the median career falls from 142 to 71 days as a goes from 0 to 1. Where seeds are scarce the same sweep is flat within noise, because contests are too rare to matter.

The invasion analysis (Figure 9, right) separates individual from collective interest and reverses with scarcity. Against a resident population at $a = 0.25$ under scarcity, a mutant at $a = 0.5$ reaches 2.07 times the resident median career and one at $a = 1$ reaches 1.44; under abundance the same mutants reach 0.95 and 0.52. Taking seeds pays when seeds are hard to come by and costs when they are not. Socially transmitted aggression settles at $\overline{a} = 0.241$ under scarcity and 0.258 under abundance, close to the resident value in both cases, because at these densities transmission has little to act on.

The setting depicts territorial disputes among a small number of girls, and the model places them accordingly: behaviour that pays the individual under scarcity, rather than a force shaping the population at large.

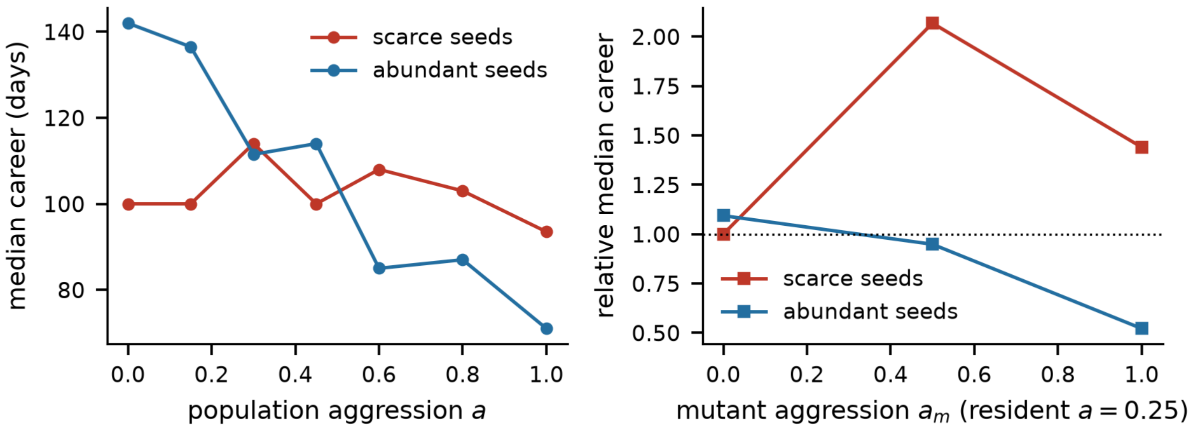

Figure 9: Left: median career length in uniform populations across the aggression parameter. Right: invasion fitness of mutants against a resident population at $a = 0.25$, as a ratio of median career lengths, with neutrality dotted.

4.7 Farming familiars

The baseline folds familiars into a maturation rate. The setting treats them as agents in their own right: a witch splits them off, they take humans until they mature into witches of their own, they leave no grief seed if killed first, and they die alongside their parent witch. A girl who meets one therefore chooses between culling it for nothing and leaving it to become a witch that may carry a seed. The setting names the second policy witch farming. Kyoko Sakura is the character who practises it, and Sayaka Miki objects to it on the grounds that leaving familiars alive gets people killed (see `docs/CANON.md`, section 2). None of its rates are quantified. We set the rate at which witches split off familiars so that a population half of which farms reproduces the maturation rate used in the baseline, then sweep the share of girls who farm.

The policy moves the supply it is meant to move. Realised maturation rises from 0.00222 to 0.00949 per witch per day across the sweep, a factor of 4.3, and the standing familiar count rises from 0.13 to 0.81.

Neither quantity the policy is meant to improve responds (Figure 10). The median career is 116, 110, 115, 116 and 106 days across the sweep, against a spread of 19 to 42 days between replicates, and Kyubey's return is 35.4, 37.3, 35.9, 33.6 and 34.0, against a spread of 6.0 to 8.2. Both are flat within noise. Repeating the sweep where witches are scarce ($\sigma = 0.02$), the regime in which the policy is argued for, gives returns of 40.9, 38.3 and 37.9, again flat within noise.

What does respond is the cost nobody in the loop is counting. Human losses rise from 16.3 to 24.6 per thousand days, an increase of one half, and the whole of the increase is in the familiar channel, which rises from 8.5 to 19.8 while the witch channel falls from 7.8 to 5.5. The share of runs in which the city is lost rises from 0 to 0.08.

The return does not move for the reason established in Section 4: it is limited by the rate at which candidates arrive rather than by the seed supply, so raising the seed supply buys nothing. The career does not move because the extra seeds arrive attached to extra witches, and the corruption spent fighting them offsets the corruption they cleanse. Within the model, farming is close to a pure transfer: the girls who practise it gain no measurable career, Kyubey gains no measurable return, and the humans pay.

This is an accounting result at the level of the population rather than a claim about any individual. A girl who is short of a seed at a particular moment faces a different question from the one measured here, and the model does not represent it.

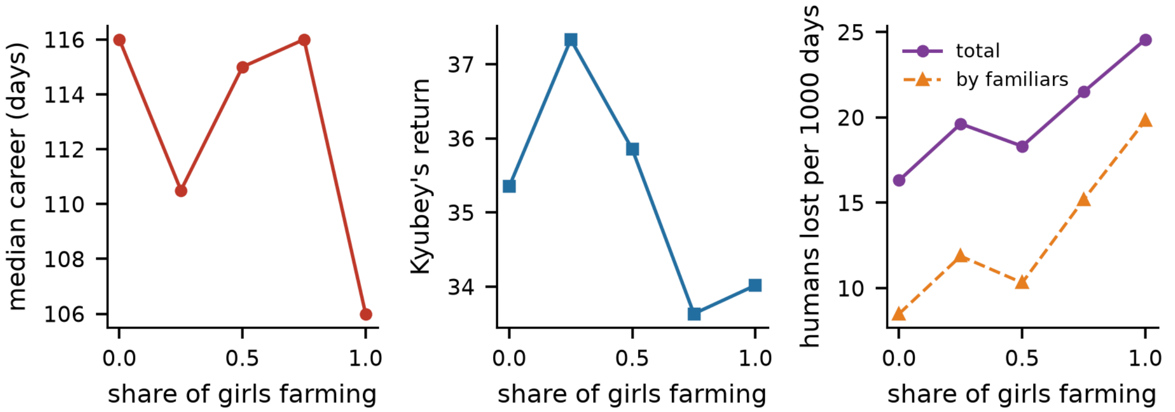

Figure 10: Familiars as individuals, across the share of girls who leave them to mature. Left: median career length. Middle: Kyubey's cumulative return. Right: human losses per thousand days, in total and through the familiar channel alone.

5 Discussion

5.1 What the results say about the cycle

Taken together, the results describe a system that cannot pay for itself. Cleansing does not remove corruption; it moves it into a seed, and a full seed becomes the next witch. The only way corruption leaves is through Kyubey, taking spent seeds out of circulation, and the only way the population is replenished is through contracts. Between those two flows the girls are intermediaries.

That is why the closed city has no settled state. With few witches the girls cannot cleanse and die on the passive drain alone; with many they are overwhelmed. The setting avoids both by leaving the city open, and the model shows how narrow the alternative is.

The comparison with fisheries is instructive at exactly the point where the analogy breaks. In an open-access fishery, effort rises until rent is dissipated and the stock is overexploited [1]. Here effort and stock are the same population, so raising effort does not merely deplete the resource, it depletes the harvesters. Kyubey's optimum has an unusual shape: few witches, so that girls saturate rather than shatter, and a high offer rate only over short horizons, since the candidate pool is the binding constraint over long ones. The starving side of Result 4 is that preference taken to its limit.

5.2 Using the model as a rule core

Simulation and language-model agents fail in opposite directions. A rule engine keeps invariants but produces numbers rather than description. Language-model agents produce fluent description [9] but do not reliably respect invariants across a long run. A practical arrangement puts the invariants in the simulation and the surface in the generative layer: the model adjudicates what is legal and computes consequences, and the language model proposes structured intents and renders the resulting state changes as prose.

The simulation layer has a second use beyond consistency. It supplies statistics for a generated narrative to respect: how long careers run, how often contests escalate, how sparse the population is, and where a girl's corruption budget goes. These are quantities a generative layer has no reliable way to estimate on its own.

5.3 Extensions

All three parties in the model could be made to learn, but the expected return differs. Two things separate them: whether the party makes decisions in the model at all, and what learning would cost relative to the result it could change.

Kyubey is the cheapest to add and the most directly useful. It is a single decision-maker with an explicit objective (cumulative energy return) acting once per macro-step, so a policy can be learned without changing anything else in the model. Result 3 compares only fixed offer rates, yet the level of the candidate pool is observable state, so a state-dependent policy might beat any fixed rate by throttling recruitment when the pool runs low and raising it once the pool recovers. If so, that Ponzi structure becomes a control problem rather than merely an observation. We tested this with tabular Q-learning (`src/madoka_abm/rl_kyubey.py`) at three discount factors. No learned policy beat the best fixed rate, and the shortfall was between 2.4% and 4.9%, which places the learned policies on the plateau the fixed rates already reach. This is the same ceiling as Result 3 and as the farming sweep above: what limits the return is the rate at which candidates arrive. `docs/ROADMAP.md` records what a stronger attempt would change.

Familiars are individuals in the module above, but their state is only a district and a feeding count, and they make no choices of their own. Giving them ground they defend, or letting a girl mark a maturing familiar as hers, would turn farming from a per-girl policy into a contested resource in its own right.

Magical girls cost considerably more but could change more. Aggression is currently a number without memory. If each girl instead learned when to escalate from local state, such as seeds held, own corruption, and the number of girls in the district, the model would become a standard sequential social dilemma. Memory makes reciprocity and reputation possible [18], so the aggression results would likely change, which is the reason to pay that cost.

Witches make no decisions in the model. They are a stage the corruption passes through, so there is no objective for them to optimise.

6 Limitations

The model covers the magical girl, witch and Kyubey subsystem, and the rules in the introduction are simplifications adopted for tractability even where they are taken from the setting. Parameters the setting does not quantify are chosen to cover distinct regimes rather than estimated, so absolute values are meaningful only inside the model.

Emotional state is folded into the passive drain rather than modelled separately, although the setting names despair as a source of corruption in its own right. The baseline folds familiars into a maturation rate and only Section 4.7 restores them as agents, so outside that section no result includes familiars having to be found and killed. Tuned to the same maturation rate, the two agree at the calibrated point: 0.34 standing witches against 0.33, a return of 33.0 against 35.9 with a replicate spread of 5.7 to 6.0, and the same transformation share of 0.92. The approximation does break down where girls are scarce, because culling depends on how many of them are around to do it and a fixed rate cannot respond to that. At $\sigma = 0.02$ the baseline holds 0.86 standing witches against the agent version's 8.39. That direction brings the overrun branch of Result 4 forward rather than pushing it back.

Combat resolution allows one engagement per witch per day and awards it to the girl with the highest effective power; a different rule would spread seed access more evenly across girls, or concentrate it in fewer of them.

Because movement and engagements happen once per day, α fixes not only how long a career runs but also how many witches it contains. Halving or doubling the length of a step therefore changes the outcome even if every other parameter is rescaled to match. Mean career lengths are computed over completed careers and are therefore biased downward by right-censoring. Populations are small at these parameters, which makes the steady state noisy: comparing the windows $[4000, 6000]$ and $[6000, 8000]$ days, standing populations move by -16% to $+25\%$ against a replicate spread of 10% to 32%.

Finally, the energy Kyubey takes at the contract is confirmed by the setting but not quantified, so η is a modelling assumption reported as a range.

7 Conclusion

Simulating the grief seed cycle as the setting describes it brings several quantitative features of it into view. At the density the setting shows, a girl's corruption budget is spent fighting witches, not fighting other girls; conflict dominance belongs to a city an order of magnitude more crowded. The rate at which girls are lost is set by how often Kyubey contracts, not by how many witches there are; what the witch supply changes is how they are lost and how long the survivors last. Cumulative return saturates in the offer rate: the advantage recruiting hard buys comes from a stock that is spent once, so it does not last. And a city closed to migration has no settled state at all: either its girls starve for want of seeds, or its witches multiply beyond what a handful of girls can contain, and Kyubey collects most in the first case. None of these were built into the rules, and each holds over a stated range of the parameter space.

Bibliography

[1] H. S. Gordon, "The Economic Theory of a Common-Property Resource: The Fishery," *Journal of Political Economy*, vol. 62, no. 2, pp. 124–142, 1954, doi: 10.1086/257497.

[2] M. B. Schaefer, "Some Considerations of Population Dynamics and Economics in Relation to the Management of the Commercial Marine Fisheries," *Journal of the Fisheries Research Board of Canada*, vol. 14, no. 5, pp. 669–681, 1957, doi: 10.1139/f57-025.

[3] C. W. Clark, *Mathematical Bioeconomics: The Optimal Management of Renewable Resources*, 2nd ed. New York: Wiley, 1990.

[4] G. Hardin, "The Tragedy of the Commons," *Science*, vol. 162, no. 3859, pp. 1243–1248, 1968, doi: 10.1126/science.162.3859.1243.

[5] E. Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action.* Cambridge University Press, 1990. doi: 10.1017/CBO9780511807763.

[6] J. Maynard Smith and G. R. Price, "The Logic of Animal Conflict," *Nature*, vol. 246, no. 5427, pp. 15–18, 1973, doi: 10.1038/246015a0.

[7] J. Maynard Smith, *Evolution and the Theory of Games.* Cambridge University Press, 1982. doi: 10.1017/CBO9780511806292.

[8] J. M. Epstein and R. Axtell, *Growing Artificial Societies: Social Science from the Bottom Up.* MIT Press, 1996. doi: 10.7551/mitpress/3374.001.0001.

[9] J. S. Park, J. C. O'Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein, "Generative Agents: Interactive Simulacra of Human Behavior," in *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST)*, 2023. doi: 10.1145/3586183.3606763.

[10] J.-Y. Girard, "Linear logic," *Theoretical Computer Science*, vol. 50, no. 1, pp. 1–101, 1987, doi: 10.1016/0304-3975(87)90045-4.

[11] J. McCarthy and P. J. Hayes, "Some Philosophical Problems from the Standpoint of Artificial Intelligence," in *Machine Intelligence 4*, Edinburgh University Press, 1969, pp. 463–502.

[12] C. Martens, "Ceptre: A Language for Modeling Generative Interactive Systems," in *Proceedings of the Eleventh AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 2015, pp. 51–57. doi: 10.1609/aiide.v11i1.12784.

[13] A. M. Smith and M. Mateas, "Answer Set Programming for Procedural Content Generation: A Design Space Approach," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 3, no. 3, pp. 187–200, 2011, doi: 10.1109/TCIAIG.2011.2158545.

[14] M. J. Nelson and A. M. Smith, "ASP with Applications to Mazes and Levels," in *Procedural Content Generation in Games*, in Computational Synthesis and Creative Systems., Springer, 2016, pp. 143–157. doi: 10.1007/978-3-319-42716-4_8.

[15] A. Summerville, C. Martens, B. Samuel, J. Osborn, N. Wardrip-Fruin, and M. Mateas, "Gemini: Bidirectional Generation and Analysis of Games via ASP," in *Proceedings of the Fourteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 2018, pp. 123–129. doi: 10.1609/aiide.v14i1.13013.

[16] R. Boyd and P. J. Richerson, *Culture and the Evolutionary Process.* Chicago: University of Chicago Press, 1985.

[17] J. D. C. Little, "A Proof for the Queuing Formula: L = \lambda W", *Operations Research*, vol. 9, no. 3, pp. 383–387, 1961, doi: 10.1287/opre.9.3.383.

[18] M. A. Nowak, "Five Rules for the Evolution of Cooperation," *Science*, vol. 314, no. 5805, pp. 1560–1563, 2006, doi: 10.1126/science.1133755.

APPENDIX A Reproduction and validation

Availability. The model, the experiment driver, the data behind every figure and both editions of this paper are at https://github.com/yuna0x0/madoka-magica-abm, with the project page at https://madoka-magica-abm.pages.dev. This is an independent study written for fun and not peer reviewed, and corrections are welcome through the repository.

The model is implemented in Python with NumPy as its only dependency (`src/madoka_abm/world.py`, about 390 lines of code). The experiment driver (`madoka-abm experiments`) regenerates every figure and the results file in about a minute on one core. Replicates use deterministic seeds $0, 1, 2, \ldots$

How this paper was produced. The research direction, the rules the model must obey, the principle that the model must agree with the setting, and every editorial decision are the author's. Implementation and drafting relied on a large language model (Claude Opus 5, Anthropic): most of the simulation code, the experiment runs, the source research behind `docs/CANON.md`, and the first drafts of both editions were produced with it and then reviewed, rejected or revised by the author.

It helps to read the paper in two parts. What can be checked mechanically is checked thoroughly: `madoka-abm-tools audit` verifies that all 65 values quoted here still match `results/`, that the two editions agree in structure and citations, and that the 60 values repeated in the READMEs and on the project page stay in step, while the identity in Equation 1 is an independent test of the implementation. What cannot be checked mechanically, such as the modelling choices, the reading of the setting, and whether the literature is invoked aptly, has been reviewed by one non-specialist reader. The quantitative results are therefore reproducible and auditable; the interpretive ones are just one reading, and corrections are welcome.

Sources for the rules. Every rule in the introduction is drawn from the setting, and the sources for each one, together with the quantities the setting does not fix, are tabulated in the project's `docs/CANON.md`.

Identity check. Equation 1 is exact, which makes it a strong test of the implementation. It is run over 27 parameter settings spanning $\rho \in [0.004, 0.016]$, $\kappa \in [0.6, 1.5]$ and $\sigma \in [0.02, 0.10]$, two seeds each. Career length predicted by Equation 2 matched realised career length to a relative error of at most 3.4×10^{-15} over all 2035 recorded careers, which is floating-point precision. `madoka-abm experiments` performs the check and records it under `results/`. A bookkeeping error in any corruption stage, such as a cost applied without being logged or a cleansing not debited against seed capacity, breaks the identity immediately.

Stationarity. Standing populations over $[4000, 6000]$ and $[6000, 8000]$ days agree within replicate noise at the calibrated migration rate and above. At the small populations these parameters produce, replicate spread is 10% to 32% of the mean, so single runs should not be read closely; the reported figures average six replicates.