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

Arisu Hoshinomiya
arisu@yunya.pet
Independent researcher

28 July 2026

Abstract

We model the system connecting magical girls, witches and Kyubey in *Puella Magi Madoka Magica* as a closed resource cycle, and study its behaviour at the population level. Our simplification keeps five elements: soul gems accumulate corruption, corruption is cleansed using grief seeds left by defeated witches, a girl whose gem saturates becomes a witch, Kyubey collects the energy released by that transition, and new girls enter by contract. We implement this as a spatial agent-based model with explicit corruption accounting, endogenous recruitment, and hawk–dove contests over grief seeds. Four results follow. First, lifetime corruption is dominated by conflict between magical girls rather than by combat with witches: when witches are rare the ratio is about 5.9 to 1, and it falls below 1 when witches are plentiful. Second, the rate at which girls leave the population is set by the contract rate and is insensitive to the grief seed supply; a richer witch supply changes how girls leave rather than how often, and the share of departures ending in transformation falls from 88% to 31%. Third, cumulative energy return is bounded by the rate at which new candidates become available, so the contract rate that maximises return decreases as the evaluation horizon lengthens: the early return is paid out of a stock that is consumed once, which is the structure of a Ponzi scheme. Fourth, aggression over grief seeds reproduces the tragedy of the commons: collective survival is highest near an aggression level of 0.15, but under scarcity slightly more aggressive strategies invade a less aggressive population. Each result is a property of the model under the stated assumptions. We give the parameter ranges over which each holds and 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, witches, and Kyubey. We formalise it as a resource cycle with five elements, which serve as the assumptions of the model:

1. A magical girl's soul is held in a soul gem.
2. The gem accumulates corruption over time and through the use of magic.
3. Corruption can be cleansed by drawing on a grief seed, which a defeated witch leaves behind.
4. A gem that reaches saturation turns its bearer into a witch, and Kyubey collects the energy released by that transition.
5. New girls enter the system by contract with Kyubey.

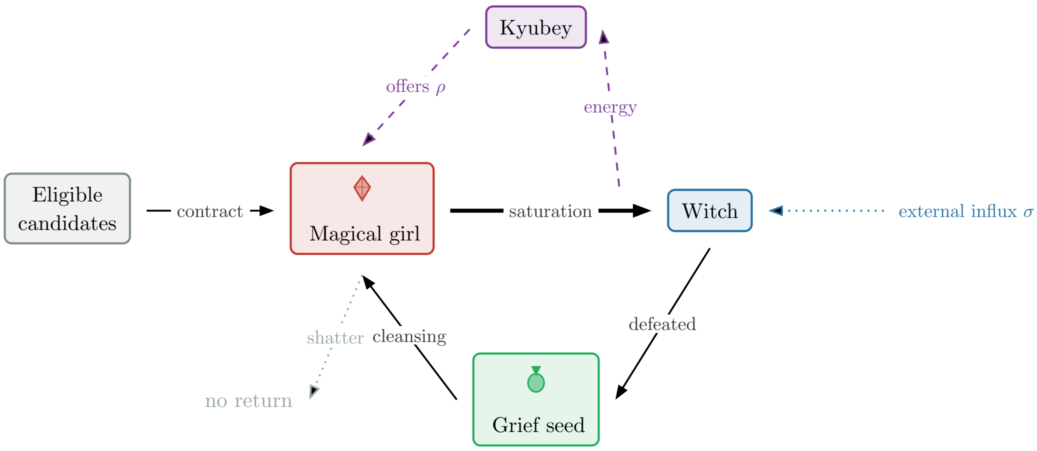

Figure 1: Schematic of the resource cycle. Solid arrows form the main loop; the purple dashed arrows are Kyubey's two roles; the dotted arrows are the two boundary flows. One is witches arriving from outside the loop. The other is girls lost to a shattered soul gem, which produces no witch and so returns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

These five assumptions define the model. Each affects only one stage of the daily cycle, so an alternative formulation can be substituted by changing the corresponding rule in Section 3.

Figure 1 shows the cycle in schematic form. Written this way, the model has the structure of a common-pool resource system. Grief seeds are the renewable resource. Their supply depends partly on witches arriving from outside the modelled loop and partly on the same magical girls who consume them, since a girl who transforms becomes a source of seeds for the others. Kyubey obtains its energy from the turnover of that population. 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 girls does a city support at once? How is the seed supply divided? How do these depend on how often witches appear and how often contracts are offered? Our interest is in behaviour that is hard to read off the rules directly, and in particular in the interaction between the seed economy and the contests over it.

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) with explicit corruption accounting, endogenous recruitment, and hawk–dove contests over grief seeds.
- 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): conflict dominates the corruption budget; the departure rate is set by recruitment while the seed supply changes only the manner of departure; cumulative energy return is bounded by candidate inflow; and aggression reproduces the tragedy of the commons, with an interior optimum.
- 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 contested value rises as free seed capacity per girl falls, so the intensity of conflict is endogenous.

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 on hard constraints and reachability. Our questions are quantitative and demographic, 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 combat power w_i, a list of held grief seeds with their residual capacities, an aggression parameter $a_i \in [0, 1]$, an age, and bookkeeping counters. A witch has a position and a power w.

3.2 Daily cycle

1. **Candidates and witch arrivals** the candidate pool grows by Λ up to a cap $P_{\max}$. Witches arrive from outside the loop at rate σ, drawn as Poisson(σ). Each standing witch becomes dormant with probability δ.

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 combat power engages it and accrues βw corruption. She wins with probability $w_i/(w_i + w)$, removing the witch and gaining 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. A shattered gem removes the girl without producing a witch, so it returns no energy.

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 $0.35V$, where the contested value $V = \left(1 + \Sigma_{\text{cap}}/M\right)^{-1}$ rises as free seed capacity per girl falls. If neither escalates, nothing happens.

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 residual capacity.

6. **Transformation** a girl with $c_i \geq 1$ becomes a witch of power λw_i, and Kyubey's return increases by one unit.

7. **Contracts** Kyubey makes Poisson(ρ) offers per day. Each is accepted with probability π and consumes one unit of the candidate pool. New girls take their aggression either from a fixed value or by imitating recent departures, weighted by how long those girls survived.

3.3 Parameters

Table 1 lists the seventeen parameters. We treat them as a design space to explore rather than as quantities to estimate, in the sense of [13], and report results as functions of the two we vary most: the witch arrival rate σ and the offer rate ρ. Absolute magnitudes such as days and head counts should be read as internal units.

Symbol	Value	Meaning
α	0.0022	passive corruption per day
β	0.030	corruption per unit of witch power engaged
κ	1.5	grief seed capacity
$c_{\text{trig}}, c_{\text{targ}}$	$0.55, 0.25$	cleansing trigger and target
p_s	0.012	gem shatter probability per lost engagement
λ	1.6	power multiplier on transformation
G	6	grid side, in districts
h	0.75	probability of moving toward a witch
σ	swept	witch arrivals per day
δ	0.004	witch dormancy rate per day
$\Lambda, P_{\max}$	$0.25, 90$	candidate inflow and pool cap
ρ	swept	contract offers per day
π	0.55	offer acceptance probability
C	0.11	corruption cost of an escalated contest
a	0.25 or swept	aggression

Table 1: Model parameters.

3.4 Corruption accounting

Corruption within a girl is conserved. It is generated, cleansed away, 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 cleanses P units, takes C_{cb} corruption in combat with witches and C_{cf} in contests, and finishes with gem load c_{end}. She satisfies

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

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

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

Two things follow before any simulation. A girl who never fights, never contests and never cleanses has a career of $1/\alpha = 455$ days, which serves as the model's natural unit of lifespan. And cleansing and conflict enter Equation 2 with opposite signs, so what sets longevity is the balance between them rather than either one alone.

In steady state, Little's law [16] relates the standing population to the contract rate $R = \rho\pi$ and the mean career length, $M = R\,\overline{L}$, and Kyubey's return arrives at

$$\dot{E} = R f_w, \tag{3}$$

where f_w is the share of departures that occur by transformation rather than by shattering. Equation 3 is central to what follows: it says the seed economy reaches Kyubey's return only through f_w.

4 Results

Runs are 8000 days. Steady-state quantities average the last 2000 days; career statistics pool completed careers across replicates. Unless stated otherwise $\rho = 0.03$, giving $R = 0.0165$ contracts per day.

4.1 Baseline

Figure 2 shows a run at $\sigma = 0.02$. The system settles at $M = 16.4$ girls and $W = 0.037$ standing witches. Witches are hunted down nearly as fast as they appear, so the visible witch population is far smaller than the population of girls hunting them. Median career length is 411 days and the mean is 673; 84.9% of departures are transformations; Kyubey collects 0.0126 units per day. Mean corruption holds near 0.42 without drift, which indicates the run has reached steady state.

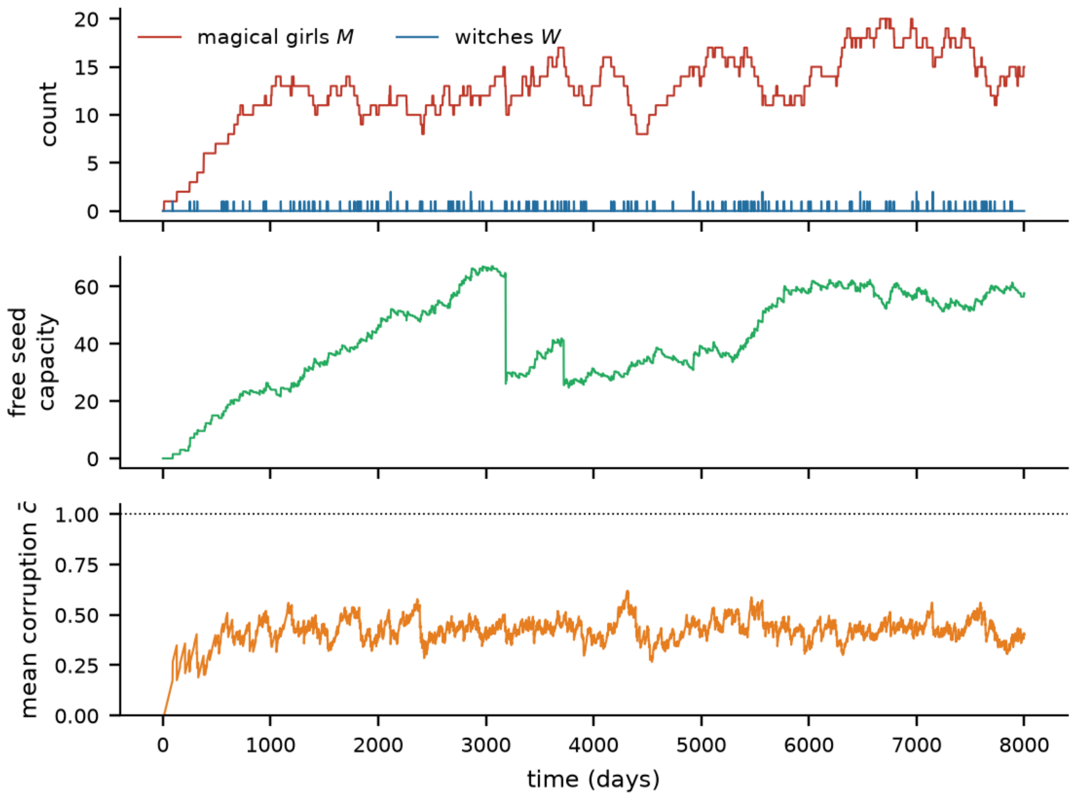

Figure 2: Baseline run at $\sigma = 0.02$, $\rho = 0.03$. Top: magical girls and standing witches. Middle: total unused grief seed capacity held by the population. Bottom: mean gem corruption, with the transformation threshold dotted.

4.2 Result 1: conflict dominates the corruption budget

Decomposing Equation 1 across witch arrival rates (Figure 3, Table 2) gives the first result. When witches are rare, corruption taken in contests with other girls exceeds corruption taken fighting witches by a factor of 5.3 to 5.9. Most of what ends a career comes from passive drain and from other magical girls.

The ratio depends strongly on witch density. It falls monotonically to 0.8 at $\sigma = 0.5$: when witches are common, seeds are common, the contested value V drops, and escalation stops being worth its cost. Conflict in this model is therefore a symptom of scarcity. Territorial exclusion appears where seeds are hard to come by and fades where they are not.

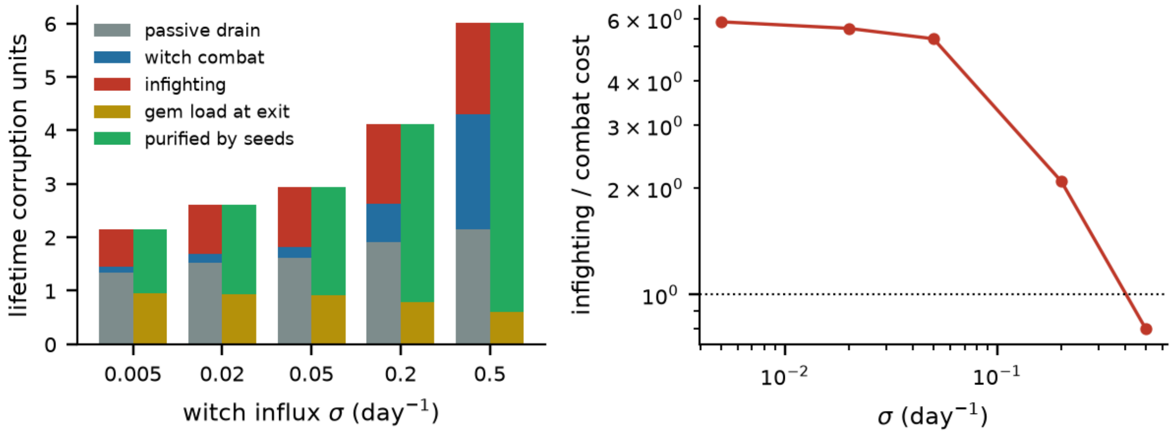

Figure 3: Corruption budget. Left: lifetime sources (left bar: passive drain, witch combat, conflict between girls) against sinks (right bar: gem load at the end of a career, corruption cleansed by seeds). Right: ratio of conflict cost to witch combat cost, with parity dotted.

σ	$\overline{L}$	passive	combat	conflict	cleansed	M	f_w
0.005	609	1.34	0.118	0.694	1.19	11.5	0.878
0.02	689	1.52	0.165	0.929	1.67	13.1	0.844
0.05	733	1.61	0.211	1.114	2.02	15.7	0.813
0.2	866	1.91	0.717	1.495	3.33	22.2	0.608
0.5	980	2.16	2.145	1.716	5.42	30.6	0.307

Table 2: Lifetime corruption budget in corruption units, mean career length in days, standing population, and transformation share. Four replicates per row, $n = 441$ to 501 careers.

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

Across a hundredfold range of witch arrival rates, the rate at which girls leave the population stays between 0.0138 and 0.0157 per day, tracking the contract rate as Equation 3 requires (Figure 4, centre, where the dashed line is Rf_w). Enriching the seed economy does not reduce how many girls end their careers. It changes how they end: the transformation share falls from 87.8% to 30.7% as σ goes from 0.005 to 0.5, because combat mortality arrives before saturation does. Kyubey's return falls with it, from 0.0126 to 0.0043 per day.

This is the only route by which the seed economy affects Kyubey, and its direction is worth noting. A world with many witches is a poor one for Kyubey, not because girls survive but because the way they are lost changes: a larger share end in a shattered gem or a contest, neither of which returns energy. Meanwhile the standing population nearly triples, from 11.5 to 30.6, and mean career length rises 61%, from 609 to 980 days, while the median rises only 15%, from 415 to 478 days. Abundance mostly enlarges the population and the upper tail of the career distribution (Figure 4, right; the 90th percentile moves from 1128 to 2612 days).

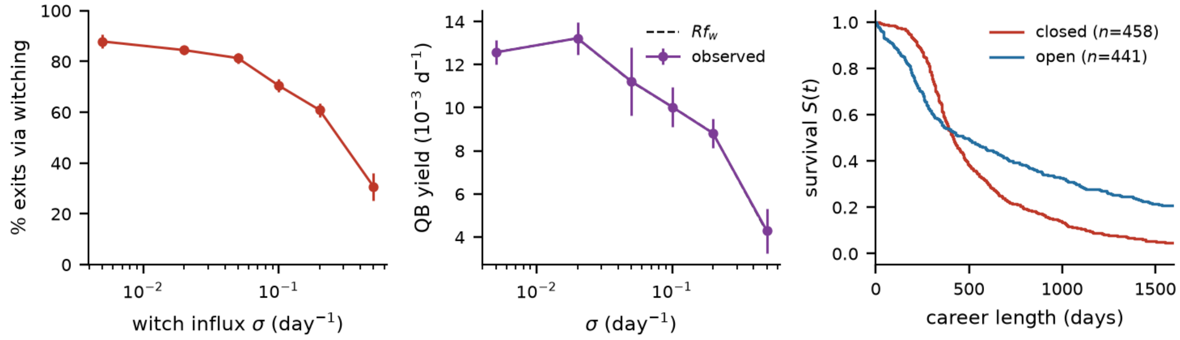

Figure 4: Left: transformation share against witch arrival rate. Centre: Kyubey's return, with the prediction Rf_w from Equation 3 overlaid. Right: career survival curves for a low arrival rate ("closed", $\sigma = 0.005$) and a high one ("open", $\sigma = 0.5$).

The full (ρ, σ) surface (Figure 5) shows the return rising with the offer rate and falling with the witch arrival rate throughout. At the corner $\sigma = 0.5$, $\rho = 0.006$ the transformation share reaches zero: every girl is lost to combat or to a contest, and no energy is returned at all.

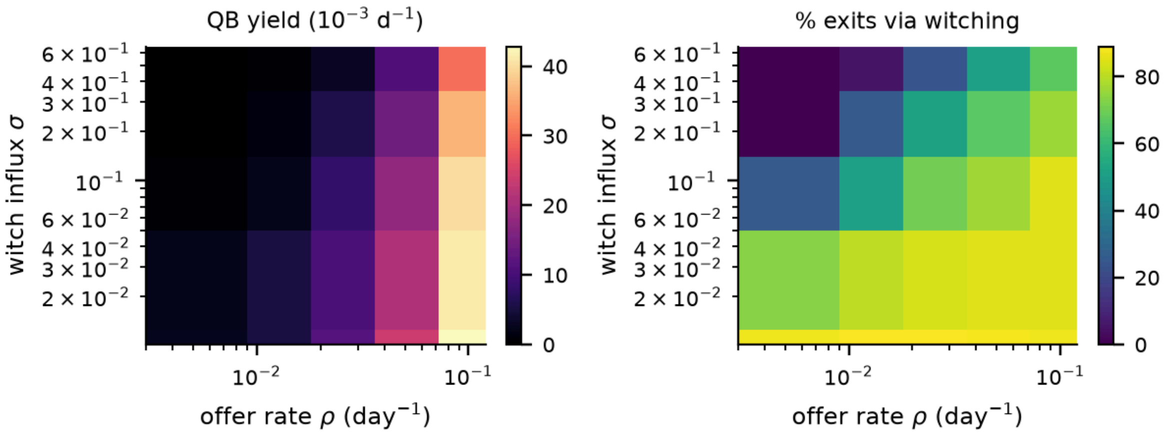

Figure 5: Kyubey's return (left) and transformation share (right) over offer rate and witch arrival rate. Two replicates per cell.

4.4 Result 3: cumulative return is bounded by candidate inflow

Figure 6 restricts the candidate pool, setting $\Lambda = 0.03$ and $P_{\max} = 60$, and sweeps the offer rate. Cumulative return saturates: past $\rho \approx 0.066$ the total after 8000 days sits between 235 and 244 units however aggressively Kyubey contracts. That level is set by a flow term $\Lambda f_w T \approx 204$ plus a one-time stock term $P_{\max} f_w \approx 51$. Their sum, 255, exceeds the observed total by about 17 units, which is the standing population at the end of the run: those girls have consumed candidates but have not yet converted.

The decomposition explains the saturation. A high offer rate converts the standing pool of eligible candidates into energy quickly, but that pool is consumed once; after it is gone, the return rate is limited by Λ, the rate at which new candidates become eligible. The offer rate that maximises cumulative return therefore falls as the evaluation horizon lengthens: $\rho^* = 0.140$ at $T = 500$ and $T = 2000$ days, dropping to 0.110 at $T = 8000$. All the advantage of a high offer rate is drawn from a stock that does not replenish, which 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.

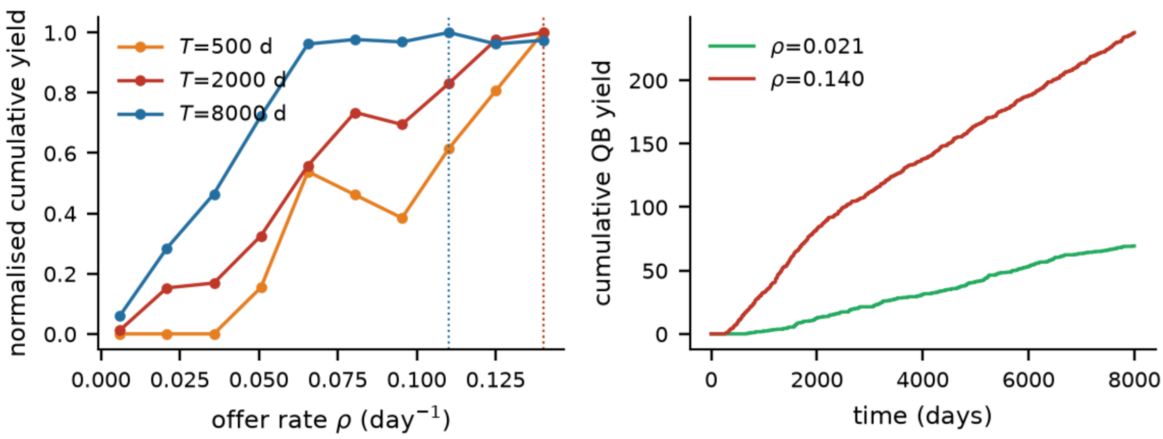

Figure 6: Left: normalised cumulative return against offer rate for three evaluation horizons, with the maximum of each marked. Right: return trajectories for a low and a high offer rate under a restricted candidate pool.

4.5 Result 4: aggression reproduces the tragedy of the commons

Fixing the aggression parameter so that populations are uniform (Figure 7, left) reveals a collective optimum at low but non-zero aggression. Median career length peaks at $a = 0.15$, at 524 days when seeds are scarce and 549 when they are plentiful, and falls to 140 and 132.5 days at $a = 1$. The standing population falls by a factor of 3.8 to 7.8 over the same range. Universal aggression is costly for everyone.

Two features are worth drawing out. At $a = 0$ the median career is 455 days in both regimes, which is exactly $1/\alpha$, the no-cleansing lifespan. Because only the highest-power girl present engages a witch, seed access concentrates on the strongest, and the median girl obtains no seed at all even when witches are plentiful. This is also why a small amount of aggression raises the median: taking seeds redistributes them away from whoever is monopolising the hunting.

The invasion analysis (Figure 7, right) separates individual from collective interest. Against a resident population at $a = 0.25$ under scarcity, a mutant at $a = 0.5$ reaches 1.035 times the resident median career, while $a = 0$ reaches 0.94 and $a = 1$ only 0.66. There is thus an interior best response above both the resident level and the collective optimum, and aggression drifts upward. When seeds are plentiful every mutant does worse than the resident, at 0.88, 0.70 and 0.56, and the pressure disappears. Socially transmitted aggression settles consistently with this, at $\bar{a} = 0.258$ under scarcity against 0.238 under abundance.

Aggression here is therefore individually defensible under scarcity and collectively costly, which is the structure of the tragedy of the commons [4]. It arises without any rule mentioning cooperation, defection, or territory.

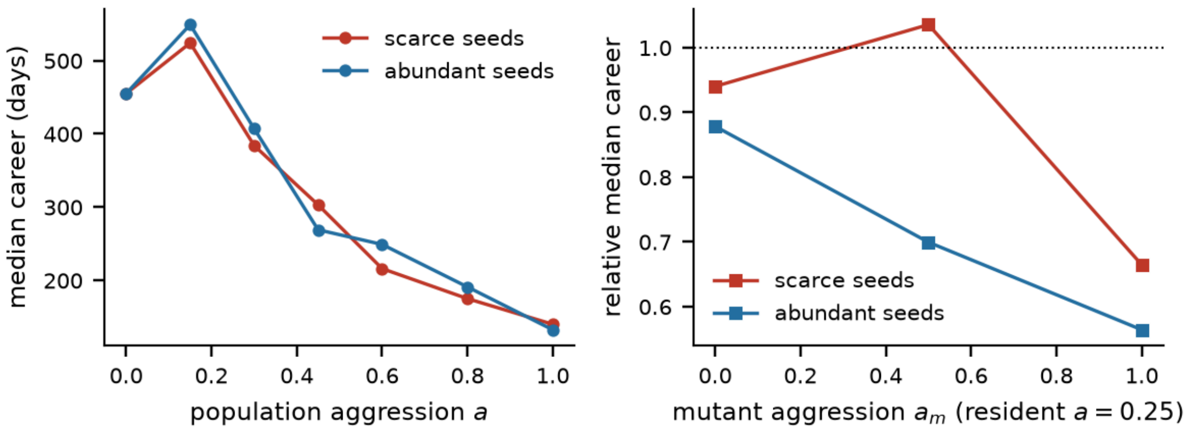

Figure 7: Left: median career length in uniform populations across the aggression parameter. Right: invasion fitness of mutants against a resident population at $a = 0.25$, as a ratio of median career lengths, with neutrality dotted.

4.6 Consistency with the reading behind the model

Three of the behaviours above agree with the reading of the setting that the model was built from: a city supporting only a handful of girls at a time, hunting grounds that are defended, and the strongest girl in an area tending to monopolise them. None of the three was written into the rules. In the model they follow from the resource structure and the engagement rule, which is a check that the model behaves plausibly over the parameter range used here.

Career length is deliberately excluded from that check. The passive drain α was set so that $1/\alpha$ falls a little over a year, and the median career duly lands near that value, so the two agreeing is calibration rather than a result.

5 Discussion

5.1 What the results say about the cycle

Taken together, the results describe a model in which the seed economy is a buffer rather than a solution. Seeds extend individual careers and enlarge the standing population, and Equation 2 shows exactly how: they add to the corruption budget. What they do not do is change the rate at which girls leave, which Equation 3 ties to the contract rate. In steady state, every contract eventually produces a departure, and a fixed share of departures returns energy.

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. The result is that Kyubey's optimum has an unusual shape: prefer few witches, so that girls saturate rather than shatter, and prefer a high offer rate only over short horizons, since the candidate pool is the binding constraint over long ones.

Result 4 adds a second-order effect that runs against the first. Conflict between girls consumes budget that would otherwise extend careers, so it shortens them, but it also redistributes seeds away from a monopolist and thereby raises the median. Neither the purely cooperative population nor the purely aggressive one does best. Near $a = 0.15$ redistribution has begun while escalation costs are still small, which is what makes that level the interior optimum.

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 on the same model (`src/madoka_abm/rl_kyubey.py`), using a macro-step of 100 days, six discrete offer rates, and the discount factor in place of the evaluation horizon. Training 2000 episodes at each of $\gamma = 0.9$, 0.99 and 0.999 produced no policy that beats the best fixed rate; the learned policies returned 5%, 14% and 26% less. Two features of the setup limit what the experiment can show. Cumulative return is nearly flat in the offer rate above $\rho \approx 0.096$, so the margin available to a state-dependent policy is smaller than the evaluation noise, and the learned tables are visibly under-fitted. We therefore record the question as open rather than settled.

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, with fifteen to thirty policies learning against each other and therefore against a moving target. Memory makes reciprocity and reputation possible, so Result 4 would likely change, which is the reason to pay that cost.

Witches make no decisions in the model. They are a resource stream, so there is no objective for them to optimise, and giving them policies would add parameters without a question to answer.

Two methodological extensions are also open. One is to train a surrogate model for the parameter sweeps so that the resolution of Figure 5 is no longer limited by simulation cost. The other is to invert the problem and infer parameters from a desired set of population-level features. In both cases the bottleneck is generating training data rather than fitting the model. Compute estimates and entry points for each are recorded in the project's `docs/ROADMAP.md`.

6 Limitations

The model covers the magical girl, witch and Kyubey subsystem, and the five assumptions in the introduction are simplifications adopted for tractability. Parameters 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 term rather than modelled separately. Combat resolution allows one engagement per witch per day and awards it to the highest-power girl present; this is what produces the monopoly effect in Result 4, and a different engagement rule would weaken it. Aggression is a scalar with no memory, so reciprocity and reputation are excluded, and equilibria under repeated interaction would likely differ. Mean career lengths are computed over completed careers and are therefore biased downward by right-censoring, which is why M exceeds $R\overline{L}$ by 14% to 91% in our runs, the gap widening with σ as the surviving cohort grows. Because movement and engagements happen once per day, α fixes not only the unit of lifespan but also how many encounters a career contains. Scaling α with every per-day rate leaves the median career at $1/\alpha$, but the dimensionless outcomes move: over a fourfold range of α the transformation share runs from 0.78 to 0.88 and the conflict-to-combat ratio from 8.3 to 4.3. The direction of every relationship we report survives that rescaling; the levels do not, and belong to this time unit. Finally, Kyubey is parametric rather than optimising; our attempt to learn a state-dependent policy (Section 5) did not beat the best fixed rate, and Figure 6 suggests its objective would in any case depend sharply on the horizon.

7 Conclusion

Simulating the grief seed cycle brings several quantitative features of it into view. Conflict between magical girls, not combat with witches, consumes most of a career's corruption budget when witches are scarce, and that conflict subsides when they are not. The seed supply changes how careers end without changing how often they end, so enriching it lowers Kyubey's return by substituting one mode of loss for another. Kyubey's cumulative return is bounded by how quickly new candidates become available, which makes its preferred contract rate depend on the horizon over which it is judged. And aggression over seeds sits in the familiar shape of the tragedy of the commons, defensible individually under scarcity and costly collectively. 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] 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.

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 a summary at https://madoka-magica-abm.pages.dev. This is an independent study written for enjoyment. It has not been 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 210 lines of code). The experiment driver (`madoka-abm experiments`) regenerates every figure and the results file in under three minutes on one core. Replicates use deterministic seeds $0, 1, 2, ...$

Identity check. Equation 1 is exact, which makes it a strong test of the implementation. Across settings spanning $\rho \in [0.02, 0.08]$, $\kappa \in [1.5, 3.0]$ and $\sigma \in [0.007, 0.027]$, career length predicted by Equation 2 matched realised career length to a relative error of at most 8×10^{-15} over all 5444 recorded careers, which is floating-point precision. 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 for $\sigma \leq 0.2$, changing by 2% to 12% between windows against a replicate spread of 7% to 19%, so the 8000-day horizon is long enough for the reported statistics. The transient is longer when witches are plentiful: at $\sigma = 0.5$ the population is still climbing through the first 6000 days, so an early window such as $[2500, 3000]$ sits inside the transient, at 16.7 against 30.6 late. Extending those runs to 24000 days moves the population between 23.9 and 31.9 across 2000-day windows without a trend, so the 8000-day horizon captures the level but leaves a slow fluctuation of roughly $\pm 15\%$; the $\sigma = 0.5$ column should be read with that spread in mind. An earlier sweep at 2500 days was too short and produced systematically low populations; we note this because it is the failure mode most likely to affect a shorter reproduction.