

A Particle-Type Framework for Social Dynamics: Typed Agents, Four Interaction Channels, and a Quantitative Design Study

A Particle-Type Framework for Social Dynamics: Typed Agents, Four Interaction Channels, and a Quantitative Design Study

Abstract

Social influence campaigns, platform integrity teams, and policymakers currently rely on influence models that treat agents as homogeneous, which prevents them from reasoning about how leaders, carriers, mediators, and unobserved populations contribute differently to behavioral cascades. The specific gap is that no existing sociophysics model assigns distinct interaction potentials per agent type and per interaction channel, so the predictive value of such a typing scheme is unknown. We formalize a "Standard Model of social particles": agents typed as baryons (stable leaders), leptons (mobile information carriers), bosons (mediators), and dark matter (unobserved populations), interacting through four channels (strong family bonds, electromagnetic like/dislike attraction, weak random transformative events, and gravity-like cultural pull), with an effective-potential description adapted from particle physics. Because no simulation has yet been run, we derive analytical design-study results under stated assumptions. Under a linear-threshold mean field with mean degree 10, lepton edge weight 0.08, threshold 0.2, and boson mediation amplification $1 + \mu f_M$ with $\mu = 3$, $f_M = 0.1$, the critical lepton seed fraction is 0.1923 versus 0.6667 for a homogeneous baseline, a factor-3.47 reduction. Weak-force events at rate 10^{-2} per day convert baryons with probability 0.3935 over 100 days, 4.13 times the lepton rate. An unobserved population fraction of 0.66 biases observed reach downward by a factor of 2.941. The typed model carries 60 extra coupling parameters, requiring a log-likelihood gain above 60 (AIC) or 207.2 (BIC at $n = 1000$) to justify itself. These thresholds define the falsification criteria for the planned agent-based test.

1. Introduction

Marketing teams that run influencer campaigns, trust-and-safety teams that track coordinated manipulation, and civil-society organizations that design behavior-change programs all face the same operational problem: they must decide where to intervene in a social system, and their current models tell them only how many agents to activate, not which kinds of agents. Influencer-selection models do distinguish node classes in practice — [4] models influencer marketing campaigns on a Twitter social graph and a synthetic Watts-Strogatz graph — but the underlying dynamics models in wide use still treat the transmitting population as interchangeable. The cost of this homogeneity is concrete: a campaign or counter-campaign designed on a homogeneous-agent model can be sized wrong by a large factor, and Section 4 derives a factor of 3.47 between the seed fraction a typed model and a homogeneous model require for the same cascade.

What would change if a typed framework held? A campaign designer could compute, before spending, how many mobile carriers versus mediators versus leaders must be activated. A platform integrity team could estimate how much observed diffusion underestimates true reach when part of the population is invisible to measurement — Section 4 derives an underestimate by a factor of 2.941 at a 0.66 unobserved fraction. A policy analyst could estimate how large a random shock (a "weak-force" event) must be to flip stable leaders, which [2] shows is a live concern: that work documents how social bots were used during the 2016 US Presidential Election to sway conversation around candidates, and reviews bot-creation mechanics and detection methods, so the question of how non-human and human agent classes differ in diffusion dynamics is already operational.

This paper does not claim the framework works. It builds the framework, derives its first quantitative consequences, and states exactly what must be measured for the framework to be accepted or rejected. It borrows its formal template from particle physics, where [8] presents the effective-potential formalism — a real scalar function of fields whose minima are the vacua — as the standard tool for analyzing transitions between states, and where [11] documents both the power of the Standard Model Lagrangian structure and its cost: at least 19 arbitrary, un-derived parameters supplied from measurement. Any social analog must confront the same parameter-count problem, and Section 4 quantifies it.

The contributions are:

1. A typed-agent formalism for social dynamics with four particle classes and four interaction channels, expressed as an effective potential in the form used in [8] (Section 3).
2. A mean-field derivation of the critical seed fraction for influence cascades, giving 0.1923 for the typed model against 0.6667 for a homogeneous baseline under stated assumptions (Section 4.1).
3. A derivation of weak-force perturbation sensitivity, giving baryon conversion probability 0.3935 over 100 days at event rate 10^{-2} per day, 4.13 times the lepton conversion probability, and a second-order reach estimate of 78.7 additional adopters (Section 4.2).
4. A derivation of the observational bias introduced by an unobserved "dark" population, a factor of 2.941 at $f_D = 0.66$ (Section 4.3).
5. A parameter-cost accounting: 64 typed coupling parameters against 4 homogeneous ones, with AIC and BIC break-even log-likelihood gains of 60 and 207.2 respectively, and a specification cost of 694.18 bits using the per-parameter entropy of 11.5697 bits (Section 4.4).
6. A spinodal calculation for the gravity channel: the external cultural-pull field needed to erase a metastable consensus is $h_{\text{spin}} = 0.3849$ for the chosen potential (Section 4.5).

2. Background and Related Work

The proposed framework sits at the intersection of multirelational social network analysis, agent-based diffusion modeling, and physics-inspired sociophysics. Each strand supplies one ingredient, and each also supplies a caution.

Multirelational structure. [1] proposes a multidimensional social network in which the relations extracted from users' collective behavior in an online sharing system are grouped into layers built on two relation types, direct and object-based; this layering is the empirical precedent for typing social interactions into distinct channels rather than collapsing them into one weighted graph. [6] introduces a multirelational social network covering direct and object-based relationships. Together these works establish that social data already separates into distinguishable relation types; what they do not do is assign different interaction dynamics per type, which is the step this paper takes.

Non-human and unobserved agents. [2] investigates the role of social bots in content spreading, reviewing how bots are created and operated, the challenges of detecting them via graph-centric, crowdsourcing, and feature-based supervised methods, and case studies including bot influence on political discussion during the 2016 US Presidential Election and social spam campaigns. This motivates two of the four particle classes: agents that diffuse without being human (a distinct type with distinct dynamics), and the population that platform measurement misses, which this paper models as dark matter. The detection problem documented in [2] is precisely the boundary-drawing problem for the dark sector.

Measuring dynamics. [3] defines differences between network snapshots by means of a Graph Differential Tuple, enabling distance measures and investigation of the speed of changes. Any test of the typed framework against empirical networks will need exactly such snapshot-difference machinery, because the four interaction channels predict different time signatures: strong bonds should change slowly, weak-force events should appear as discontinuous jumps. [3] supplies the measurement instrument for that comparison.

Typed diffusion models in practice. [4] models influencer marketing campaigns and validates them on a real-world Twitter social graph and a synthetic Watts-Strogatz small-world graph. This is the closest existing practice to typed-agent dynamics, and its finding that conclusions are graph-dependent is a direct warning for the framework proposed here: type-dependent results must be validated per network, not assumed universal.

Physics-inspired social modeling. [5] models the Information Society with Ising agents and a multi-agent system. The external field in [5] is analogous to the gravity channel in this paper — an unconscious, global pull.

Methodological framing. We acknowledge that all models are "wrong" in some sense and the goal is to be wrong in a meaningful way. This paper adopts that stance explicitly: the typed framework is a hypothesis to be tested, not a truth claim. The paper assumes that validating the typed model requires demonstrating agreement with observed network structure, functional behavior, and dynamic cascades, not merely structural similarity. Applied here, this means showing that the typed model reproduces empirical network structure (L1) is not enough; it must also match functional and dynamic behavior (L2-L3), which is why the test design in Section 7 checks cascade dynamics, not just clustering.

The physics template and its cost. [8] introduces the effective-potential formalism used in particle physics. Section 3 adopts exactly this restricted form for a social order parameter. [8] also connects first-order phase transitions to observable signatures (gravitational waves from early-Universe transitions), which is the structural analogy for this paper's claim that weak-force perturbations should produce detectable discontinuities in social time series. The Standard Model contains many arbitrary parameters. The 19-parameter lesson is quantified for the social case in Section 4.4. We assume a specification cost of about 11.57 bits per continuous parameter at tolerance 10^{-4} . Section 4.4 uses this figure to price the typed model's parameter specification. We adopt a Lakatosian framing of hard core and protective belt to articulate falsification conditions. Section 6 uses this framing to state which observation would expel the typed framework from consideration rather than merely adjusting its protective belt.

3. Methods

3.1 Agent types

We assume a closed population of N agents partitioned into four classes:

- **Baryons** B : stable leaders. Low mobility, high connection stability, high resistance to conversion. Count $N_B = f_B N$.
- **Leptons** L : mobile information carriers. High mobility, low tie strength, low conversion resistance. Count $N_L = f_L N$.
- **Bosons** M : mediators. Agents whose presence amplifies transmission between other agents without themselves being primary sources. Count $N_M = f_M N$.
- **Dark matter** D : unobserved populations. Agents invisible to platform measurement but participating in dynamics. Count $N_D = f_D N$.

We assume throughout the design study $N = 1000$, $f_B = 0.04$, $f_L = 0.20$, $f_M = 0.10$, $f_D = 0.66$. These are illustrative design values, not measurements; Section 6 treats sensitivity to them.

3.2 Interaction channels

Four channels, each with its own potential, mirror the four physical forces:

- **Strong s :** family/close bonds. Short-range, high weight, slow to form or break.
- **Electromagnetic e :** like/dislike attraction, signed (attractive or repulsive), medium range.
- **Weak w :** random transformative events at low per-agent rate λ_w , changing an agent's state discontinuously.
- **Gravity g :** unconscious cultural pull, long-range, unscreened, modeled as an external field h on a social order parameter.

3.3 Effective potential

Following the tree-level form in [8], which restricts renormalizable potentials in $d = 4$ to quartic powers of the fields, we write a social effective potential for a scalar order parameter ϕ (population-level behavioral alignment, $\phi \in [-1, 1]$):

$$V(\phi) = \Lambda^4 + t\phi + \frac{1}{2}m^2\phi^2 + \frac{1}{6}\kappa\phi^3 + \frac{1}{24}\lambda\phi^4 - h\phi,$$

where t is a seed influence (analog of a source term), m^2 sets the stability of the current consensus, κ and λ are self-interaction strengths of the behavioral field, and h is the gravity-channel external field, directly analogous to the external field in the Ising treatment of [5]. Minima of V are social vacua (stable consensus states); a first-order transition between vacua, in the sense of [8], corresponds to a rapid population-wide behavioral flip. The weak channel enters not through V but as a Poisson process on agent states, as specified in Section 4.2.

3.4 Cascade dynamics

Transmission is a linear threshold process. Agent i adopts a behavior when the weighted influence from active neighbors reaches its threshold θ_i ; we take a homogeneous threshold $\theta = 0.2$ for the design study. Lepton edges carry weight $w_L = 0.08$. Boson mediation multiplies lepton influence by an amplification factor $(1 + \mu f_M)$ with mediation strength $\mu = 3$. Each agent has mean degree $\bar{k} = 10$. The homogeneous baseline assigns every edge weight $w_h = 0.03$ and no mediation. All values in this paragraph are stated design assumptions; the derivations in Section 4 show every arithmetic step from them.

3.5 Validation design (specified, not executed)

No simulation or empirical measurement is reported in this paper. The planned validation, specified here for completeness, follows the practice of [4]: run the typed dynamics on a real-world social graph (the Twitter social graph used in [4] is a candidate substrate) and on a synthetic Watts-Strogatz graph, compare against homogeneous-agent and random baselines on clustering, cascade size, and cascade timing, and use snapshot-difference measures in the sense of [3] to test the predicted time signatures of the four channels. Results of that validation are future work.

4. Analysis

Every input number in this section is a stated design assumption from Section 3 or a figure attributed to a cited work; every arithmetic step is shown.

4.1 Critical seed fraction for cascades

Setup. In the mean field, a susceptible agent's influence from active leptons is

$$I_L = \bar{k} f_L w_L (1 + \mu f_M).$$

Step 1: mediation amplification. With $\mu = 3$ and $f_M = 0.10$:

$$1 + \mu f_M = 1 + 3 \times 0.10 = 1.30.$$

Step 2: influence per unit lepton fraction. With $\bar{k} = 10$ and $w_L = 0.08$:

$$\bar{k} w_L (1 + \mu f_M) = 10 \times 0.08 \times 1.30 = 1.04.$$

Step 3: critical fraction. The cascade condition is $I_L \geq \theta$ with $\theta = 0.2$, so

$$f_L^{\text{crit}} = \frac{\theta}{\bar{k} w_L (1 + \mu f_M)} = \frac{0.2}{1.04} = 0.19231.$$

Step 4: no-mediation comparison. With $f_M = 0$ the amplification is 1.00 and

$$f_{L, \text{no med}}^{\text{crit}} = \frac{0.2}{10 \times 0.08 \times 1.00} = \frac{0.2}{0.80} = 0.25.$$

Step 5: relative reduction from mediation.

$$\frac{0.25 - 0.19231}{0.25} = \frac{0.05769}{0.25} = 0.2308,$$

i.e. boson mediation lowers the required lepton seed fraction by 23.08%.

Step 6: homogeneous baseline. With all edges at $w_h = 0.03$ and no type structure, the critical active fraction is

$$f_h^{\text{crit}} = \frac{\theta}{\bar{k} w_h} = \frac{0.2}{10 \times 0.03} = \frac{0.2}{0.30} = 0.6667.$$

Step 7: typed-to-homogeneous ratio.

$$\frac{f_h^{\text{crit}}}{f_L^{\text{crit}}} = \frac{0.6667}{0.19231} = 3.467.$$

Under these assumptions the typed model predicts cascades from a seed fraction 3.47 times smaller than the homogeneous model. This is the paper's headline analytical contrast; it is a consequence of the assumed weights, not a measurement.

4.2 Weak-force perturbation sensitivity

Setup. Weak events arrive as a Poisson process with rate $\lambda_w = 10^{-2}$ events per agent per day (assumption). Per-class weak-channel coupling (exposure factor): $g_B = 0.5$ for baryons, $g_L = 0.1$ for leptons (assumptions). Effective conversion rate is $\lambda_w g$.

Step 1: baryon conversion probability over $T = 100$ days.

$$\lambda_w g_B T = 10^{-2} \times 0.5 \times 100 = 0.5,$$

$$p_B = 1 - e^{-0.5} = 1 - 0.60653 = 0.39347.$$

Step 2: lepton conversion probability over the same window.

$$\lambda_w g_L T = 10^{-2} \times 0.1 \times 100 = 0.1, \quad p_L = 1 - e^{-0.1} = 1 - 0.90484 = 0.09516.$$

Step 3: ratio of conversion probabilities.

$$\frac{p_B}{p_L} = \frac{0.39347}{0.09516} = 4.134.$$

Under these assumptions, a stable leader is 4.13 times as likely as a mobile carrier to be flipped by the same random-event stream, because the baryon's larger weak-channel coupling ($g_B = 0.5$ versus $g_L = 0.1$) gives a larger effective hazard over the window. This is the opposite of the intuition that leaders are safe from random shocks, and it is a direct consequence of the assumed resistance values.

Step 4: second-order reach of converted baryons. The baryon population is $N_B = f_B N = 0.04 \times 1000 = 40$. Converted baryons over the window:

$$N_B^{\text{conv}} = N_B p_B = 40 \times 0.39347 = 15.74.$$

We assume each converted baryon contacts its full neighborhood of $\bar{k} = 10$ agents and that each contacted agent adopts with probability $q = 0.5$ (an assumed susceptibility to strong-bond influence, chosen equal to g_B for symmetry of the design study). Second-order adopters:

$$N^{(2)} = N_B^{\text{conv}} \times \bar{k} \times q = 15.74 \times 10 \times 0.5 = 78.70.$$

So a 100-day weak-event stream at 10^{-2} events per agent per day is projected, under these assumptions, to produce about 78.7 second-order adopters through baryon conversion alone, on top of the 15.74 converted leaders themselves.

4.3 Observational bias from the dark sector

Setup. The dark fraction is $f_D = 0.66$ (Section 3.1 assumption): 66% of the participating population is invisible to platform measurement. We assume dark agents participate in dynamics identically to observed agents, so true reach R_{true} scales linearly with total population while observed reach R_{obs} scales with the observed fraction only.

Step 1: observed fraction.

$$f_{\text{obs}} = 1 - f_D = 1 - 0.66 = 0.34.$$

Step 2: bias factor. With uniform participation, $R_{\text{obs}} = f_{\text{obs}} R_{\text{true}}$, so

$$\frac{R_{\text{true}}}{R_{\text{obs}}} = \frac{1}{f_{\text{obs}}} = \frac{1}{0.34} = 2.9412.$$

An unobserved population fraction of 0.66 therefore biases observed reach downward by a factor of 2.941: a campaign that appears to reach 1,000 observed agents has, under this assumption, actually reached about 2,941. The bias is exactly reciprocal in the observed fraction, so at $f_D = 0.9$ it would be $1/0.1 = 10$; estimating f_D is therefore the dominant uncertainty, and it is the same boundary-drawing problem that [2] documents for bot detection via graph-centric, crowdsourcing, and feature-based supervised methods.

4.4 Parameter cost of the typed model

Setup. The typed model assigns one coupling per channel per ordered type pair: 4 channels $\times$ 4 $\times$ 4 type pairs = 64 parameters. The homogeneous model assigns one coupling per channel: 4 parameters. Extra parameters:

$$\Delta k = 64 - 4 = 60.$$

Step 1: AIC break-even. The Akaike Information Criterion is $AIC = 2k - 2 \ln \hat{L}$. Adding $\Delta k = 60$ parameters costs $2 \times 60 = 120$ in AIC, so the typed model must gain

$$\Delta \ln \hat{L}_{AIC} = \frac{120}{2} = 60$$

in log-likelihood over the homogeneous model to break even.

Step 2: BIC break-even. The Bayesian Information Criterion is $BIC = k \ln n - 2 \ln \hat{L}$. With $n = 1000$ (Section 3.1):

$$\Delta \ln \hat{L}_{BIC} = \frac{\Delta k \ln n}{2} = \frac{60 \times \ln 1000}{2} = \frac{60 \times 6.90776}{2} = \frac{414.47}{2} = 207.24.$$

Step 3: specification cost in bits. We assume a specification cost of about 11.57 bits per continuous parameter at tolerance 10^{-4} . Applying this per-parameter cost to the 60 extra typed parameters:

$$C_{\text{spec}} = 60 \times 11.5697 = 694.18 \text{ bits.}$$

The typed model must therefore buy, at minimum, a log-likelihood gain of 60 (AIC) or 207.2 (BIC at $n = 1000$), and its full specification carries a one-time description cost of about 694 bits beyond the homogeneous baseline. This is the social analog of the many un-derived parameters counted in the Standard Model, scaled to a model with 60 free couplings.

4.5 Spinodal field for the gravity channel

Setup. We take the potential of Section 3.3 with the simplest first-order-capable parameter choice: $m^2 = -1$, $\kappa = 0$, $\lambda = 6$, $t = 0$ (all design assumptions). Then

$$V(\phi) = \Lambda^4 - \frac{1}{2}\phi^2 + \frac{1}{4}\phi^4 - h\phi.$$

The metastable consensus sits near $\phi = 0$; the gravity field h tilts the potential. The spinodal field h_{spin} is the tilt at which the local minimum disappears, found where $V' = 0$ and $V'' = 0$ simultaneously.

Step 1: derivatives.

$$V'(\phi) = -\phi + \phi^3 - h, \quad V''(\phi) = -1 + 3\phi^2.$$

Step 2: inflection point. $V'' = 0$ gives

$$3\phi_s^2 = 1 \quad \Rightarrow \quad \phi_s = \frac{1}{\sqrt{3}} = 0.577350.$$

Step 3: spinodal field. $V'(\phi_s) = 0$ gives $h_{\text{spin}} = -\phi_s + \phi_s^3$:

$$\phi_s^3 = \left(\frac{1}{\sqrt{3}} \right)^3 = \frac{1}{3\sqrt{3}} = 0.192450,$$

$$h_{\text{spin}} = -0.577350 + 0.192450 = -0.384900.$$

The magnitude of the external cultural-pull field needed to erase the metastable consensus is therefore

$$|h_{\text{spin}}| = 0.3849.$$

For $|h| < 0.3849$ the consensus survives as a local minimum; for $|h| \geq 0.3849$ it vanishes and the population rolls to the alternative vacuum. This is the gravity-channel analog of the external-field result in [5]; here the threshold field is computed explicitly rather than observed in simulation.

5. Results

All numbers below are computed in Section 4 from the stated design assumptions; none are measurements, and no simulation or empirical data are reported in this paper.

1. **Critical seed fraction (Section 4.1).** With mediation, $f_L^{\text{crit}} = 0.19231$; without mediation, 0.25; mediation lowers the requirement by 23.08%. The homogeneous baseline requires $f_h^{\text{crit}} = 0.6667$. The typed-to-homogeneous ratio is 3.467, i.e. the typed model reaches cascade threshold from a seed fraction 3.47 times smaller.
2. **Weak-force conversion (Section 4.2).** Over $T = 100$ days at $\lambda_w = 10^{-2}$ per day: baryon conversion probability $p_B = 0.39347$, lepton conversion probability $p_L = 0.09516$, ratio 4.134. Converted baryons: 15.74 of 40; projected second-order adopters: 78.70.
3. **Dark-sector bias (Section 4.3).** At $f_D = 0.66$, observed reach underestimates true reach by a factor of 2.9412; the factor is $1/(1 - f_D)$ and rises steeply with f_D .
4. **Parameter cost (Section 4.4).** $\Delta k = 60$ extra parameters; AIC break-even log-likelihood gain 60; BIC break-even gain 207.24 at $n = 1000$; specification cost 694.18 bits at 11.5697 bits per continuous parameter.
5. **Spinodal field (Section 4.5).** For $m^2 = -1$, $\kappa = 0$, $\lambda = 6$: $|h_{\text{spin}}| = 0.3849$.

Each result is conditional on the design assumptions of Section 3; the Discussion quantifies which assumptions the headline numbers are most sensitive to.

6. Discussion

The central result of this paper is the factor-3.47 gap between the seed fraction a typed model and a homogeneous model require for the same cascade (Section 4.1). It is an analytical consequence of assumed weights, not a finding about any real network, and the first question is whether the assumptions driving it are defensible.

Sensitivity of the headline number. The ratio $f_h^{\text{crit}}/f_L^{\text{crit}} = (\bar{k}w_L(1 + \mu f_M))/(\bar{k}w_h) = w_L(1 + \mu f_M)/w_h$ is independent of $\bar{k}$ and θ ; it depends only on the weight ratio $w_L/w_h = 0.08/0.03 = 2.667$ and the mediation factor 1.30. If the true lepton-to-homogeneous weight ratio were 1.5 instead of 2.667, the ratio would fall to $1.5 \times 1.30 = 1.95$; if $w_L = w_h$, the typed model would actually require a *larger* seed fraction than homogeneous (1.30 versus 1.00) because mediation amplification alone cannot compensate for equal weights. The headline contrast therefore rests almost entirely on the assumed weight separation between carrier edges and average edges, which is the single most important quantity for the planned validation to measure.

Falsification. We state explicitly what would expel this framework rather than merely adjust an auxiliary assumption. The observation that would falsify the central claim is: on a real-world social graph, measuring per-type edge weights and finding that the carrier-type weight w_L does not exceed the population-average weight w_h by a factor of at least $1/(1 + \mu f_M)$ — under the measured mediation parameters — while cascade experiments simultaneously show no seed-fraction advantage for type-targeted seeding over random seeding at matched budget.

If type-targeted seeding performs no better than random seeding of the same number of agents, the typing scheme carries no predictive value and the 60 extra parameters buy nothing; that single experimental outcome would defeat the framework's reason to exist.

Limitations. The result rests on the following assumptions, with sensitivity:

- *Mean-field linear thresholds.* Real cascades on heavy-tailed graphs deviate from mean field; [4] found that conclusions drawn on a power-law Twitter graph did not hold on a Watts-Strogatz graph because the outdegree distributions differ, so the mean-field numbers here could shift substantially on any real graph.
- *Assumed weights and thresholds.* $w_L = 0.08$, $w_h = 0.03$, $\theta = 0.2$, $\mu = 3$ are design values; the headline ratio is proportional to w_L/w_h as shown above.
- *Uniform dark-sector participation.* Section 4.3 assumes dark agents participate identically to observed agents; if dark agents are systematically less active, the bias factor 2.941 is an overestimate.
- *Poisson independence of weak events.* Section 4.2 assumes events arrive independently at rate λ_w ; clustered shocks would change the conversion probabilities.
- *Asymptotic bit cost.* The per-parameter cost of 11.5697 bits is taken as an assumption; the 694.18-bit specification cost follows directly from this assumption.
- *Single scalar order parameter.* A one-dimensional ϕ cannot represent multiple competing behaviors; the spinodal calculation is illustrative of the method, not a prediction about any specific population.
- *No peer review, no simulation.* Every number is analytical under stated assumptions; none has been tested against data.

Open questions. What are the actual per-type edge weights on real graphs, and how do they vary across platforms? Does the weak-channel prediction of discontinuous jumps in social time series survive the snapshot-difference analysis of [3], which defines graph differences through a Graph Differential Tuple enabling distance measures and speed-of-change estimates? Can the four-channel time signatures (slow strong-bond change, discontinuous weak events) be separated empirically, or do they confound in observed data? How should the 64 couplings be parameterized to reduce the effective count — for example, by factorizing into channel and type terms — and would a factorized form retain the factor-3.47 advantage?

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. Three user groups are named in the Introduction. If the framework's central result holds after validation: a campaign designer changes the *composition* of a seed set, not just its size, targeting roughly $0.1923 \times N$ carriers rather than the $0.6667 \times N$ a homogeneous model prescribes; a platform integrity team changes its reach estimates by applying the factor $1/(1 - f_D)$, derived in Section 4.3, to correct observed diffusion for unobserved participants; a policy analyst changes the sizing of random-shock interventions using the conversion probabilities of Section 4.2. No deployment exists today because the framework is unvalidated.

(b) Readiness. Against the NASA Technology Readiness Level definitions, this work is purely analytical, has not been peer reviewed, and involves no experiments: it sits at TRL 1–2 (basic principles observed and reported; concept and/or application formulated). Analytical work of this kind cannot exceed TRL 3 (analytical and experimental proof of concept), and even TRL 3 would require the validation simulations of Section 3.5 to have been executed and checked against observations, which has not happened. No manufacturing or deployment readiness is claimed.

(c) Smallest next test. The smallest de-risking experiment is a simulation study on the two graph families used in [4] — the real-world Twitter social graph and a synthetic Watts-Strogatz small-world graph — measuring: (i) per-type edge weights w_L and w_h estimated from the graph; (ii) cascade seed-fraction thresholds for type-targeted versus random seeding at matched budget. The outcome that refutes the framework is the falsification condition of Section 6: type-targeted seeding shows no seed-fraction advantage over random seeding once measured weights are substituted into the Section 4.1 formula. A secondary check uses Graph Differential Tuple methods in the sense of [3] to test whether weak-channel events appear as discontinuous jumps in the graph-difference time series.

(d) Binding constraints. The constraints that would bind first, in order: (1) *data* — per-type edge weights require labeling agents by type, which is exactly the detection problem [2] documents as unsolved for bot versus human classification via graph-centric, crowdsourcing, and feature-based supervised methods; (2) *measurement* — the dark fraction f_D is not directly observable and must be estimated, with the bias factor of Section 4.3 amplifying any estimation error; (3) *compute* — the 64-coupling model must be fit on graphs of the scale used in [4], which is feasible but requires the factorized parameterization discussed in Section 6 to avoid overfitting at $n = 1000$ -scale datasets, where the BIC penalty of 207.24 log-likelihood units applies. Cost figures, market sizes, and deployment timelines cannot be projected from the material in this paper and are left as open questions.

(e) Order-of-magnitude bounds. The only quantified bound available is the specification cost derived in Section 4.4: 694.18 bits to write down the typed model's extra couplings at tolerance 10^{-4} . No other cost or performance projection is supported by the derivations here.

(f) Risks and misuse. The same machinery that helps a campaign designer target carriers and mediators helps a manipulator do so more cheaply; the factor-3.47 seed reduction is a capability statement, not a safety statement. The dark-sector correction of Section 4.3 could also be misused to inflate claimed reach in reporting. Both risks argue for the validation being conducted with the detection-tool landscape of [2] in view, and for publishing the falsification criteria alongside any operational tooling.

8. Conclusion

This paper built a typed-agent framework for social dynamics — four particle classes, four interaction channels, an effective potential in the restricted quartic form of [8] — and derived its first quantitative consequences under fully stated assumptions. The principal analytical results are: a critical seed fraction of 0.19231 for the typed model against 0.6667 for the homogeneous baseline, a factor of 3.47; weak-event conversion probabilities of 0.39347 for baryons versus 0.09516 for leptons over 100 days, a ratio of 4.13, with 78.70 projected second-order adopters; a dark-sector reach bias factor of 2.941 at $f_D = 0.66$; a parameter cost of 60 extra couplings with AIC and BIC break-even log-likelihood gains of 60 and 207.24 and a specification cost of 694.18 bits; and a spinodal field $|h_{\text{spin}}| = 0.3849$ for the gravity channel. Following the advice of [7] to be wrong in a meaningful way, the paper states its own falsification condition: if type-targeted seeding shows no advantage over random seeding on real graphs once measured weights are used, the framework's 60 extra parameters buy nothing and it should be discarded. The next step is the simulation study specified in Section 7(c), on the graph families of [4], with snapshot-difference analysis in the sense of [3]; until that is done, every number here is a design-study projection, not a finding about the world.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.42-codeagent). Language models were tools in that process; they are not authors. The named author is accountable for the content.

- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by openai/gpt-oss-120b, zai-org/glm-5.3-flash, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 16 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] Multidimensional Social Network in the Social Recommender System. arXiv:1303.0093v1. <https://arxiv.org/abs/1303.0093v1> [2] Measuring social spam and the effect of bots on information diffusion in social media. arXiv:1708.08134v1. <https://arxiv.org/abs/1708.08134v1> [3] Quantifying Social Network Dynamics. arXiv:1303.5009v1. <https://arxiv.org/abs/1303.5009v1> [4] Modeling Influencer Marketing Campaigns in Social Networks. arXiv:2106.01750v3. <https://arxiv.org/abs/2106.01750v3> [5] Information Society: Modeling A Complex System With Scarce Data. arXiv:1201.1547v1. <https://arxiv.org/abs/1201.1547v1> [6] Social Recommendations within the Multimedia Sharing Systems. arXiv:1303.0284v1. <https://arxiv.org/abs/1303.0284v1> [7] A primer on data-driven modeling of complex social systems. arXiv:2210.08636v1. <https://arxiv.org/abs/2210.08636v1> [8] Cosmological phase transitions: from perturbative particle physics to gravitational waves. arXiv:2305.02357v3. <https://arxiv.org/abs/2305.02357v3> [11] QNFO: Standard Model