PERRY DIME PUBLICATIONS · WHITE PAPER

The Last Will Be First and the First Will Be Last

Itô calculus, queueing theory, and a case for torque per capita as the North Star KPI for repairing the healthcare system

Jefferson Richards · November 23, 2026
Published by Perry Dime Publications, a division of Richards Plus, LLC · www.perrydime.com
Draws on: *Put Options on Loyalty: The Architecture of Certainty* (Perry Dime Publications, 2026); *Coupled Cost of Delay* (shop-floor queueing); a financial-operations queueing brief; a queueing-theory synthesis report; and the Itô calculus literature.

ABSTRACT

Every supply chain, trading desk, claims department and clinic is a system of queues. Work arrives, waits, is served, and leaves. Where the wait is longest, the system is constrained, and the people at that constraint decide how fast everything else can move. This paper proposes a single North Star KPI, **torque per capita**: the amount of constraint-clearing work each person in a queue can deliver per unit time, given the technology they have been handed.

We argue that the mathematics to model such systems continuously has matured. Itô calculus replaces the discrete Markov chains used in queueing and behavioral models with stochastic differential equations, which GPUs can simulate as full distributions rather than point estimates. We connect this to three existing bodies of work, a shop-floor model of coupled delay, a financial-operations queueing brief, and a behavioral-insurance framework, and apply the result to the healthcare coverage crisis during open enrollment. The proposal is employer-sponsored **captive insurance with medical stop-loss** as a governance instrument that aligns incentives for health. The title is the organizing principle: in a queue served by those who can state their needs, the people who learn to articulate them move to the front.

1. Thesis and North Star KPI
2. Why now
3. Itô calculus in brief
4. From queues to diffusions
5. Three arenas, one queue
6. Torque per capita, defined
7. Healthcare: captive insurance as governance
8. Narrative, orthogonality and the next cohort
9. Last first, first last
10. Risks and guardrails
11. Research agenda
Sources and notes

1 Thesis and North Star KPI

> *“But many that are first shall be last; and the last shall be first.” — Matthew 19:30 (see also Matthew 20:16)*

This paper reads that line as a statement about queues. It is also a statement about uncertainty. The author’s working framing is that the order of events in a complex system is governed by processes that are random in their detail and lawful in their aggregate. In mathematical terms, that is the territory of stochastic differential equations, and, once space is added, stochastic partial differential equations. We use this as a metaphor and a modeling program, not as a theological proof.

Torque per capita

Picture a machinist at a potter’s wheel, or at a CNC machine. Clay becomes a finished good only if torque is applied to the wheel. If the wheel does not turn, nothing reaches the next step in the supply chain. Torque per capita (TPC) generalizes this image. It measures how much work per unit time the people in each queue of a supply chain can apply, given the technology they are enabled with. It is a case for investing in the tools that increase that torque, because those tools free the bottlenecks that cap the whole chain.

WORKING DEFINITION

Torque per capita at a stage is the rate of value-moving work completed by one person at that stage, per unit time, at the quality the next stage requires. System torque is limited by the stage where it is lowest relative to demand.

2 Why now

Three developments make this the right moment to build the model.

1. **Itô calculus is getting attention.** Kiyosi Itô’s stochastic calculus has long been the working language of mathematical finance (Black–Scholes) and stochastic differential equations. It is now also cited for applications in machine learning. Material that was scarce several years ago is now easy to find.

2. **Probabilistic compute.** GPUs do not only handle zeros and ones. They operate on continuous-valued arrays, so they can simulate thousands of random paths of a stochastic process in parallel and return the full distribution of outcomes, the “gray scale” between black and white.
3. **Language models lower the cost of articulating a need.** People without a science, technology, engineering or mathematics background can now state a need precisely enough for a system to act on it. That changes who gets served.

The open *P versus NP* problem is relevant as context only. Queueing and scheduling problems are often computationally hard, and the interest in that question reflects a general push toward better tools for hard optimization. This paper makes no claim about its resolution.

3 Itô calculus in brief

Itô calculus extends ordinary calculus to stochastic processes such as Brownian motion. Brownian paths are continuous but not differentiable anywhere and have infinite variation over any interval, so the standard Riemann integral does not apply. Itô’s construction defines the stochastic integral as a limit of Riemann sums, with the integrand evaluated at the *left* endpoint of each small interval:

$$Y_t = \int_0^t H_s \, dX_s \tag{1}$$

The integrand H must be *adapted*: its value at time t may depend only on information available up to time t. In finance, this is the rule “first decide what to do, then observe the price change.” It forbids clairvoyance, because a trading strategy cannot buy just before each uptick.

Two results matter here. First, a stochastic differential equation for a state variable X:

$$dX_t = \mu(X_t, t)\, dt + \sigma(X_t, t)\, dB_t \tag{2}$$

Second, **Itô’s lemma**, the change-of-variables rule. It differs from the ordinary chain rule because of a quadratic-variation term:

$$df(X_t,t) = \left(\partial f/\partial t + \mu\, \partial f/\partial x + \tfrac{1}{2}\, \sigma^2\, \partial^2 f/\partial x^2 \right) dt + \sigma\, \partial f/\partial x\, dB_t \tag{3}$$

The Stratonovich integral is an alternative that evaluates the integrand at the midpoint, obeys the ordinary chain rule, and can be converted to and from the Itô form.

WHY ADAPTEDNESS MATTERS BEYOND FINANCE

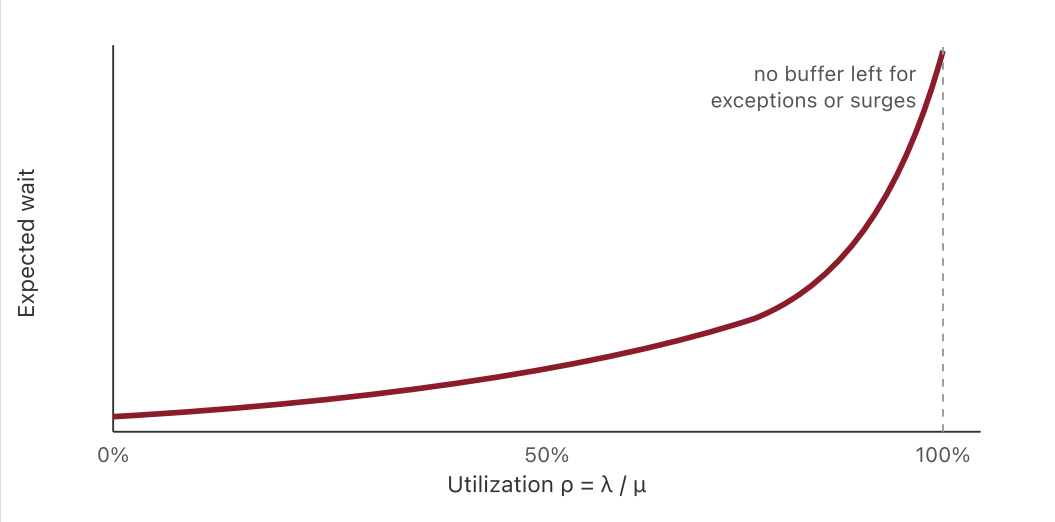
A manager, a clinician, or an insurer also acts on information available now and then observes what follows. Any model of behavior or flow that lets a decision-maker see the future overstates what is achievable. The adapted-integrand rule is a built-in honesty check for the models in this paper.

4 From queues to diffusions

A queue is the simplest system in which waiting has a price. The standard measures are arrival rate λ, service rate μ, utilization $\rho = \lambda/\mu$, and work in progress (WIP). Little's Law ties them together:

$$L = \lambda \times W \tag{4}$$

Average backlog equals throughput times time in the system. A team that completes 40 tasks a day, with each item spending two days in the process, holds about 80 items in progress. As utilization approaches 100%, waiting time rises sharply. In a basic single-server queue the expected wait grows in proportion to $1/(1 - \rho)$.



Illustrative shape of waiting time against utilization. The curve is schematic, not fitted to data.

The step to continuous models

Most queueing and behavioral models are written as Markov chains: the system sits in one of a countable set of states and jumps between them with fixed probabilities. When volumes are large, a standard approximation replaces the jumpy count of items in the queue with a continuous diffusion that is kept from going negative by a reflecting boundary at zero:

$$dQ_t = (\lambda - \mu)\,dt + \sigma\,dB_t + dL_t \tag{5}$$

Here Q is the queue length, σ captures variability in arrivals and service, and L is the local-time term that pushes Q back whenever it hits zero. This is the heavy-traffic diffusion approximation from classical queueing theory. It turns a table of transition probabilities into a small number of parameters, a drift (λ − μ) and a volatility σ, that can be estimated from timestamps and simulated at scale. Itô's lemma then gives the dynamics of any smooth function of the state, such as expected cost of waiting.

5 Three arenas, one queue

The three source documents look at the same structure from three angles. In each, work or risk arrives, waits, and moves on, and a few constrained nodes govern the whole.

Arena	Source	What queues	Core mathematics
Shop floor	*Coupled Cost of Delay* (a hardware-integration colleague, de-identified)	Engineering tasks coupled by interfaces and dependencies	CCoD = $(I - W)^{-1} d$; spectral radius of W
Finance	Queueing brief for alternative-asset operations (a financial-operations colleague, de-identified)	Trades, reconciliations, NAV, KYC, approvals	λ, μ, ρ, WIP, Little's Law; constraint analysis
Insurance and behavior	*Put Options on Loyalty*	Individuals moving among behavioral states	Markov transition matrix; stationary π with πP = π

5.1 Shop floor: delay is a property of position

Reinertsen's Cost of Delay assigns a value-per-time to each work item and sequences work by CoD divided by duration (the CD3 rule). That treats each item's delay cost as independent. In hardware integration it is not: a change to one subsystem invalidates verification downstream. The coupled model puts the dependencies in a weighted matrix W and writes the total delay cost as a Leontief-style solution:

$$\mathrm{CCoD} = (\mathrm{I} - \mathrm{W})^{-1}\,\mathrm{d} = (\,\mathrm{I} + \mathrm{W} + \mathrm{W}^2 + \ldots\,)\,\mathrm{d} \tag{6}$$

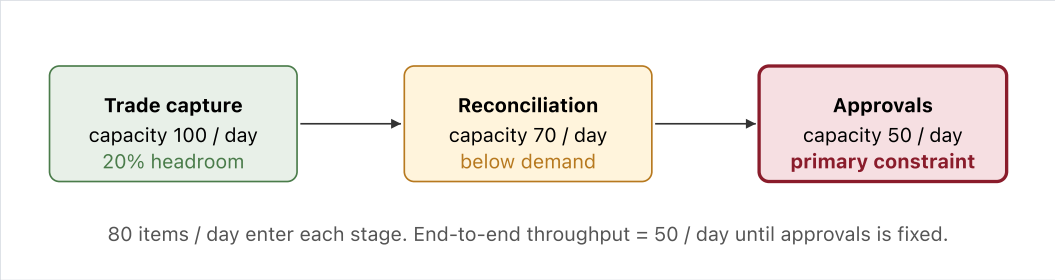
The series makes the blast radius explicit: first-order, second-order and higher-order paths each add cost. When feedback cycles push the largest eigenvalue of W toward 1, the series diverges. That is a formal account of the "integration wall," the sudden schedule collapse that real programs experience.

Where Itô enters. Equation (6) is deterministic and static, while real coupling weights and delay costs fluctuate and, as the model itself notes, grow as downstream items harden their assumptions. Replacing d and W with stochastic processes gives a stochastic extension, for example $\mathrm{d}d_t = -\kappa(d_t - \bar{d})\,\mathrm{d}t + \sigma\,\mathrm{d}B_t$ (a mean-reverting process), from which one can read off the *probability* that the spectral radius crosses 1 before a milestone, instead of a single yes-or-no answer.

The shop-floor communication problem. The same document stresses that coupling often lives only in experienced engineers' heads. Subject-matter experts and newcomers (the author's name for them is "fresh meat") must communicate effectively. The newcomer who can state what they need and where they are blocked gets unblocked first. This is the first appearance of the title's principle.

5.2 Finance: the constraint is not the busiest team

The financial-operations brief shows how an alternatives manager's operations form a chain of queues. Its worked illustration has 80 items a day entering each stage, with capacity of 100 per day for trade capture, 70 for reconciliation and 50 for approvals.



Adding capacity to trade capture does not raise throughput. Source numbers are illustrative figures from the queueing brief.

The brief's practical advice (map, measure, model, improve; use percentiles as well as averages; never trade accuracy or segregation of duties for speed) is the discipline torque per capita needs. Its dashboard of arrivals, throughput, aging, and utilization are the raw measurements for the drift and volatility terms in equation (5).

The supporting synthesis report adds the supply-chain context: performance measurement across transport, information, inventory and resources, the bullwhip effect, and treatment of Markov chains, Brownian motion and random walks as the mathematical base for queueing. It also observes that sales funnels are queueing systems. The point carries over to claims, prior authorizations, referrals and appointments.

5.3 Insurance: behavior as a chain, and its continuous extension

Put Options on Loyalty models a population in three behavioral states: S_1 mutual cooperation, S_2 neutral or untested, and S_3 negligent or risky. Insurance policy language and pricing change the transition probabilities. The insurer's aim is a governed matrix whose stationary distribution concentrates in S_1:

$$\pi \, P_{\text{governed}} = \pi \tag{7}$$

The book also gives a continuous-time responsibility measure, with intention accumulated against negligence and a baseline $\lambda > 0$ for the untested:

$$R(t) = \int_0^t I(\tau)\, d\tau \,/\, \left(\int_0^t N(\tau)\, d\tau + \lambda \right) \tag{8}$$

A new participant has $R = 0$, not infinite responsibility; they are untested. As negligence accumulates R falls, and as intention accumulates R rises.

The Itô extension. Equations (7) and (8) are deterministic summaries of a random process. Let $p_t \in [0,1]$ be the fraction of a covered group in cooperative behavior. A natural continuous model is a mean-reverting diffusion on a bounded interval:

$$dp_t = \theta(g)\,(\,p^* - p_t\,)\,dt + \sigma\,\sqrt{(\,p_t(1-p_t)\,)}\,dB_t \tag{9}$$

The reversion speed $\theta(g)$ depends on the governance design g (premium structure, exclusions, incentives); p^* is the target cooperation level; and the square-root term keeps p within $[0,1]$. The discrete stationary vector π becomes a stationary *density* over p, which can be estimated and simulated. Adding a spatial index (clinic, employer, region) gives a stochastic partial differential equation, which is the SPDE form named in the author's framing. Because the integrand is adapted, the model respects the fact that the insurer sees behavior only after it happens.

> **STATUS OF EQUATION (9)**
> Equation (9) is a proposed model form, not an estimated result. Its parameters would have to be fit to real, consented data before any claim is made about its predictive power.

6 Torque per capita, defined

With the queue as the common object, TPC can be written down. For stage k with n_k people, let $\mu_k(T_k)$ be the service rate achieved with technology T_k:

$$TPC_k = \mu_k(T_k) \,/\, n_k \tag{10}$$

$$\text{System throughput} = \min_k\,(\,n_k \cdot TPC_k\,) \le \min\,(\,\lambda,\, \min_k\,\mu_k\,) \tag{11}$$

This gives a clear investment rule: spend on the technology that raises TPC at the constraint, and weight the choice by coupling, because a node that many others depend on carries a higher coupled cost of delay (equation 6). It also explains why aggregate productivity statistics hide the problem. A stage with idle capacity looks efficient while the real limit sits elsewhere.

Setting	Torque is...	Typical constraint	Technology that raises TPC
Shop floor	Parts and integrations completed per engineer per week	Late-found cross-domain coupling	Dependency-graph tooling that makes coupling visible
Alternatives operations	Reconciled items, approvals or NAV steps per analyst per day	Specialist approval queues	Workflow timestamps, triage, exception automation
Healthcare coverage	Resolved claims, authorizations and care decisions per staff member per day	Prior authorization and claims adjudication	Plan-level data and captive-owned decision rights (see §7)

The healthcare row is a hypothesis to be tested with plan data. We have not measured it here.

7 Healthcare: captive insurance as governance

The specific problem this paper aims at is the healthcare coverage crisis, and the timing is open enrollment. The proposal rests on one legal structure: a **group captive** or employer-owned captive insurer covering employee benefits, backstopped by **medical stop-loss** reinsurance.

How the mechanism is supposed to work

1. **Risk is retained where behavior can be influenced.** The employer, not a distant carrier, bears the first layer of claims. Profit lost to avoidable claims is felt directly.
2. **Insurance as governance.** Policy language, indemnities, exclusions and bylaws define expected conduct. Tight language narrows the variance of outcomes, which is the ε–δ argument in *Put Options on Loyalty*: for any tolerable variance ε there is a language boundary δ that keeps the group near the target state.
3. **Aligned incentives for health.** Where employees choose to participate, wellness and biometric programs, including wearable devices, give both sides data on risk. The employer's incentive to keep claims low and the member's incentive to stay healthy point the same way.
4. **Reinsurance market access.** The author's observation is that a conventional employer faces a small number of health insurers (about four, in the author's figure), while a captive can reach a much larger set of reinsurers (about fifty-six, in the same account) for stop-loss. More competing backstops should mean better pricing and terms. *These counts need independent verification before publication.*

5. **Data back to regulators.** Insurers model risk factors for each policy they write and report to the government. Captives writing many unique policies add to that model base.

In the model

The captive changes $\theta(g)$ in equation (9). Retained risk and shared savings raise the pull toward p^*, and stop-loss caps the volatility of the tail. Itô's lemma then gives the dynamics of expected claims cost as a function of p. The goal of the simulation is not a forecast of any person. It is to compare governance designs by the distribution of plan-level outcomes: median cost, tail cost, and the time-to-recover after a shock.

> COMPLIANCE AND FAIRNESS CONSTRAINTS
> In the United States, employer wellness programs, health-data collection and any premium differentials tied to health factors are bounded by HIPAA nondiscrimination rules, the ADA, GINA and ERISA, and state captive and stop-loss regulation varies. Participation has to be voluntary, data use disclosed, and the design reviewed by counsel and a licensed actuary. Nothing here is legal or actuarial advice.

8 Narrative, orthogonality and the next cohort

A technical fix will not work unless the public can discuss it. The author's concern is that a divided public conversation weakens the shared belief in majority rule with minority voice, the core assumption of Robert's Rules of Order.

Orthogonal arguments

Represent each argument as a vector in a joint semantic and sentiment space. Two speakers addressing the same topic can point in nearly perpendicular directions, so their cosine similarity is close to zero. They are not disagreeing along one axis; they are not on the same axis. Trigger phrases make this worse: a term like "climate change" can pull a listener into a pre-packaged narrative before the sentence ends.

$$\cos\theta = \langle a, b\rangle / (\|a\|\,\|b\|) \approx 0 \Longrightarrow \text{the arguments are orthogonal} \tag{12}$$

Legislative and parliamentary records make it possible to model how particular speakers argue and where they diverge. For the healthcare debate, the same tools could help find framings that sit in the overlap between audiences. This remains a research direction; no measurements are reported here.

Two cohorts on campus (a scenario, not a forecast)

The author expects the next generation of students to split into two groups. One is athletes, whose position is strengthened by name, image and likeness rules. The scenario extends that logic to biometric and movement data, whose value could be sold to recruiters or trainers. The other is "mathletes," students in science, technology, engineering and mathematics whose work is needed for infrastructure and technology. Campuses then start to resemble boot camps, where physical health matters because employers carry the cost. The arts and humanities remain the third, less funded group.

The point for this paper is the queue structure inside the scenario. Attention and funding flow to whoever generates measurable value, which pushes others to the back. Whether this scenario occurs is an empirical question. It is included to show where the model would apply.

9 Last first, first last

Return to the title. In queueing theory, the order of service is a design choice. First-come-first-served is fair in a narrow sense, but it is not optimal when items differ in the cost of waiting. The weighted-shortest-job-first and CD3 rules put items with the highest delay cost per unit of work at the front, and that is often not the item that arrived first.

Now add the human factor. A priority rule can only use what the system is told. The person who can state the cost of their delay and name the dependency they are blocking moves forward. The person who cannot is served late, however long they have waited. In that sense those who can articulate their needs are satisfied first, and the title applies in both directions: the newcomer who speaks clearly can go ahead of the veteran who does not, and the veteran's long wait becomes the last place in line.

Language models change who can do this. A person with no technical training can now describe a need in terms a scheduling or triage system accepts. That is a gain in torque per capita at the point where the person meets the system, which is often the longest queue in healthcare: the wait to be understood.

10 Risks and guardrails

- **Surveillance, consent, and unwitting agents.** Millions of people act as unwitting agents of behavioral modeling. Each phone is a sensor. Each public post, photo (now convertible to text), purchase and location trace is a data point. The people producing them usually believe almost nobody is looking, but because the material is public and crawlable, anyone who collects it at scale can read it. Surveillance programs enacted after 2001, such as the Patriot Act, are the best-known example of the same pattern of piecing scattered signals together. Credit scores are another: the author argues that lobbying for access to them reflects how valuable behavioral signals are.

  This matters for the models here because Itô-style models need many observed paths. Equation (9) describes a population-level diffusion, but the way to fit it is to watch the *idiosyncratic* part of each person's behavior: the small, individual deviations from the group mean, which is the $\sigma\, dB_t$ term. Unwitting agents supply exactly that noise, for free, and without agreeing to be a data point in a model of their own future behavior. The adapted-integrand rule from §3 limits what a model may use (only information available up to time t), but it says nothing about whether the information was obtained fairly.

  The same capability, applied to health data without consent, would be a serious harm. Any use of the models here should rest on opt-in participation, data minimization, clear limits on secondary use, and a disclosed boundary between data a person knowingly contributes (for example, a wearable joined voluntarily in a captive plan) and data scraped from them without their knowledge. Where the second kind is the only available source, the model should not be used to make individual decisions about coverage, pricing or employment. Aggregate use is the minimum standard, and individual use needs the person's agreement.
- **"Coercion" is a strong word for a reason.** The book uses "stochastic coercion" for engineering behavior toward a cooperative stationary state. A governed matrix that concentrates mass on S_1 also removes options. Designs need an exit, an appeals process, and transparency about what is measured.
- **Equity.** Employer-based captives suit stable, larger employers. Gig workers, small firms and the unemployed need a path, or the system leaves out the people who need coverage most.
- **Model risk.** Diffusion approximations assume large volumes and light tails. Rare catastrophic claims need explicit jump terms or reinsurance structure, not a larger σ.
- **Claims in this paper.** Statistics quoted from the author's narrative, such as the insurer and reinsurer counts, and the campus scenario, are unverified here.

Idiosyncratic versus adverse: a liability hypothesis

The unwitting-agent point has a second consequence, which concerns how harm is classified. Pharmacology and tort law distinguish two kinds of bad outcomes, and the difference decides who is responsible.

- **Idiosyncratic reaction.** A rare response that is already known to occur in a small share of people, though it cannot be predicted for any one person. The population-level frequency is documented and can be warned about.
- **Unexpected adverse reaction.** A harmful response that nobody had documented or could reasonably have anticipated.

In the model, the first lives in a *known* tail of the distribution of responses. The second lies outside anything the model had observed. Foreseeability is central to negligence: a harm that is rare but known can ground a duty to warn, to screen, or to monitor, while a truly unforeseeable harm generally cannot.

Influence as a nervous-system exposure

The author's argument is that phones couple influencers directly into users' nervous systems. Content arrives continuously, through the same device that carries a person's messages, location and health data, and it is tuned to each user's observed responses. If exposure of that kind is treated like a dose, then most people respond with ordinary engagement, a few show strong responses, and a rare few respond in extreme ways. Those extreme responses are, in the pharmacological sense, idiosyncratic. The existence of such a tail is known from the attention-targeting data itself, which is the same unwitting-agent data described above, because those who collect it can see who reacts strongly and how often.

The hypothesis follows. If the rare, extreme reactions are known at the population level, then events in which a crowd acts on online influence may be better analyzed under a *known-risk* framework than under an *unforeseeable-harm* one. The author raises the January 6, 2021 riot at the U.S. Capitol as an example where this question could be asked: whether adjudicating such events draws on tort concepts of known, rare reactions and duty to warn, similar to medical malpractice, in addition to criminal law about individual acts.

LIMITS OF THIS ARGUMENT

This is a legal hypothesis and not settled law. Medical malpractice doctrine applies to licensed professionals in a treatment relationship, and courts have not generally treated influencers or platforms that way; extending it would need a duty of care, causation and damages to be shown, and would face free-speech and Section 230 defenses. Criminal liability for individuals who broke the law is a separate question, and nothing here reduces personal accountability for chosen acts. The paper proposes the idiosyncratic-versus-unexpected distinction as a way to *model and classify* risk, and asks counsel to test where, if anywhere, it carries legal weight.

Why it matters for healthcare. The same classification governs how a captive insurer writes exclusions and prices tails. Idiosyncratic events are insurable and can be named in policy language and stop-loss terms, while unexpected ones become the exclusions and surprises that policy wording tries to narrow (the ε–δ argument in §7). A model that records rare-but-known reactions as such is also one that cannot later say the risk was unforeseeable. That raises the stakes on the consent boundary above: whoever holds the data that makes a tail *known* also takes on the duty that comes with knowing it.

11 Research agenda

1. **Instrument one queue.** Pick a single healthcare workflow, such as prior authorization at a self-funded employer, and capture timestamps for arrival, wait, service and rework. Estimate λ, μ, σ and ρ.
2. **Fit the diffusion.** Fit equation (5) to that data, validate against held-out weeks, and compare with a plain discrete-event simulation.
3. **Make coupling visible.** Build the dependency matrix W among clinical, administrative and payment steps, and compute the probability that its spectral radius approaches 1 under load.
4. **Estimate TPC.** Compute equation (10) at each stage before and after a technology change, and test whether system throughput rises only when the constraint stage improves.
5. **Simulate governance designs.** Using equation (9) on a GPU, compare captive structures by the full outcome distribution, with the compliance limits in §7 built into the scenarios.
6. **Test the narrative layer.** Measure semantic orthogonality in published health-policy statements and test which framings overlap across audiences.

Sources and notes

1. Richards, J. *Put Options on Loyalty: The Architecture of Certainty*. Perry Dime Publications, a division of Richards Plus, LLC, 2026. Source for the responsibility index, the ε–δ policy-language argument, the three-state Markov model and the Integral Mass captive case study.
2. *Coupled Cost of Delay: A Graph-Theoretic Extension of Product Development Economics for Complex Hardware Integration*. April 2, 2026 (author de-identified). Cites Reinertsen, *Principles of Product Development Flow*.
3. *Queueing theory for alternative asset management: a visual guide*. (author de-identified). Source for the three-stage constraint illustration and the operating-dashboard measures.
4. *Synthesis Report: Queuing Theory & Operations Research Foundations*. OCR synthesis of academic reference texts on stochastic processes, supply chain performance and project risk.
5. "Itô calculus." *Wikipedia*. en.wikipedia.org/wiki/Itô_calculus. Source for the Itô integral, adaptedness, the Itô–Stratonovich distinction and the finance interpretation. Itô's lemma (equation 3) and the heavy-traffic reflected diffusion (equation 5) are standard results stated from general knowledge of stochastic calculus and queueing theory; equations (9) to (11) are proposals made in this paper.
6. Matthew 19:30 and 20:16 (New Testament).

Perry Dime Publications, a division of Richards Plus, LLC · www.perrydime.com

© 2026 Perry Dime Publications. All rights reserved.

Prepared November 23, 2026 from a dictated narrative and the source documents listed above. Colleagues who contributed source documents are not named. This paper is a research proposal. It is not medical, legal, tax, insurance or investment advice.