A Spirit Lives in the Tree: Belief as a Productive Input and the Life Cycle of Ritual

I develop a dynamic learning model of the institutional supply of motivational belief. An informed authority creates manufactured scarcity over access to divine favor and institutes ritual as the exclusive access channel. Workers believe successful ritual execution raises the perceived return to effort, allowing the authority to convert belief into rents. The belief is false in its object but real in its effects: it raises effort by inflating perceived productivity. Yet productive false belief is self-undermining. Stronger belief increases effort but also inflates expected success beyond what the true technology can sustain, making belief most fragile when it works best. Ritual complexity acts as ex-ante reputation insurance: it weakens current incentives by lowering perceived execution success, but slows belief erosion by making future failures attributable to imperfect performance rather than to the absence of divine favor itself. As technology makes effort productive on its own, belief-based incentives become redundant, generating secularization from within even under monopoly. The model explains why rituals emerge under weak technology, become elaborate to manage anticipated disappointment, and decline with technological progress.

Keywords: economics of religion; motivational beliefs; ritual complexity; belief-based rent extraction; misspecified Bayesian learning.

JEL Codes: Z12; D83; C73; D91; D42.

1. Introduction

Why do people work hard when they cannot be sure effort will pay off? Sometimes the answer is not better tools or higher wages, but a story. A farmer who believes that the right ritual wins divine favor works harder than one who sees the harvest as pure luck, even if the ritual itself does nothing. The belief is false in its object, but the extra effort and output are real in its effects. Belief works like a technology: it does not change what effort physically accomplishes, but it changes how much people are willing to supply by changing how productive they believe effort to be.

This connects to what anthropologists call the ritual economy: beliefs and values are not private opinions outside economic life, but forces that shape production, distribution, and use. Motivated by this idea, I build belief into the production environment. When technology

is too weak to motivate effort on its own, an authority can use belief to make people work harder. I argue that this is one economic origin of ritual. The key question is how an authority preserves a false but productive belief when outcomes keep producing disappointing evidence, and what happens when technology later becomes a substitute source of motivation.

Such a story is valuable only where outcomes are uncertain.[1] We surround risky activities with ritual but leave routine ones alone. Malinowski's Trobriand fishermen used elaborate magic on the dangerous open sea but none in the calm lagoon. The reason, in this model, is economic rather than mystical: when effort reliably pays off, people work hard on their own; when the link between effort and success is uncertain, they hold back, and a motivating story has room to work. Ritual fills the gap where technology runs out.

But a productive false belief is fragile. If the authority promises that ritual and effort bring success, every failed harvest becomes evidence against him, and belief begins to erode. The belief works by raising effort, but it also raises expectations beyond what the true technology can sustain. The central mechanism is therefore self-undermining: false belief raises perceived productivity, perceived productivity raises effort and output, but inflated expectations make failure more disappointing and accelerate belief erosion. The belief is therefore most fragile when it works best.

Ritual complexity manages this disappointment. By making the ritual harder to perform perfectly, the authority makes failure easier to attribute to imperfect ceremony rather than to the absence of divine favor itself. Complexity buys plausible deniability. But it is not free. A ritual that is too hard to perform seems less likely to deliver divine favor, weakening the very effort incentives the myth was meant to create and reducing the credit the authority receives after success. Ritual complexity therefore creates a tradeoff between motivating effort today and protecting belief tomorrow.

[1] Singh (2018) develops a cultural-evolutionary account of shamanism in which shamans specialize in influencing or diagnosing uncertain, important, and otherwise uncontrollable outcomes. Here, I model how such uncertainty can be converted into a belief-based incentive system whose failures are managed through ritual complexity.

As technology improves, effort becomes productive on its own, and the authority lets the myth fade—not because he is exposed or displaced by a rival faith, but because the myth no longer earns its keep. This is a theory of how myths die from within: progress quietly makes them redundant. It reframes the decline of magic and the rise of secular life not as enlightenment defeating superstition, but as technology removing the economic reason the superstition existed. The same authority who once had reason to preserve the myth now has reason to abandon it.

I formalize this argument as a dynamic learning game between a fully informed, long-lived shaman, representing a monopolistic religious institution or charismatic authority claiming privileged access to divine power, and short-lived Bayesian workers who learn under a culturally inherited, misspecified subjective model. The shaman creates manufactured scarcity over a fictitious productive input: access to divine favor. He does not monopolize a real technology. He monopolizes access to something workers believe is productive. The scarcity is manufactured, but its behavioral effects are real because workers change effort in response to it. Workers believe that successful ritual execution raises the marginal return to effort, so monopoly control over the ritual channel generates rents. In this sense, authority converts public belief into rent by monopolizing a procedure believed to grant access to success.[2]

Building on Bénabou and Tirole (2016)'s view of beliefs as economic assets, this paper shifts attention from the demand side to the supply side of motivational belief. Bénabou and Tirole emphasize why individuals may value optimistic beliefs: such beliefs can sustain effort, self-control, and motivation. Here, I take the demand for motivational belief as given and study its institutional supply. When technology is weak and outcomes are uncertain, people are receptive to narratives that make effort appear more effective than the true technology alone would imply. The shaman meets this demand by defining ritual as the exclusive access channel to divine favor, converting a demand for optimistic belief into a specific institutional product that he alone controls.

[2] Relatedly, Harvey (2025) reads ancient Mesopotamian and Israelite creation myths as works of political economy that naturalized divine service, portraying humans as created to labor for the gods while temple and palace institutions extracted resources to support ritual provision and religious personnel. I formalize a dynamic mechanism behind such belief-based extraction: ritual complexity preserves the system by shaping how failures are interpreted.

Two objects carry the economics. The myth is a motivation technology: a belief-based incentive device that raises effort by increasing workers' perceived return to effort. Ritual complexity is a technology of blame allocation: a reputation-insurance device that shapes how harvest outcomes are interpreted. By making perfect execution harder to verify, complexity allows failed harvests to be attributed to imperfect performance rather than to the absence of divine favor itself, protecting future rents at the cost of weaker current incentives. The optimal ritual therefore trades off current motivational force against future credibility.

Unlike costly-signaling or club-good accounts, ritual complexity here is not chosen by believers to signal commitment or deter free-riding.[3] It is chosen by the authority to manage interpretation. Complexity is costless to prescribe directly but costly indirectly: it lowers perceived execution success, effort, and current tribute. Its function is not to screen types but to create deniability, slow belief erosion, and extend surplus extraction.[4]

The model is subtle because the data from which workers learn are endogenously generated. Workers' effort determines the true likelihood of harvest success, while ritual complexity affects effort and reshapes how informative success or failure is about the myth. The shaman does not manipulate outcomes directly; he shapes the data-generating process through which workers learn about his legitimacy.

This same logic extends beyond religion: the mechanism, not only its secularizing endpoint, applies wherever an authority controls a prescribed path to success. A parent may

[3] See Iannaccone (1992) for the classic club-good account of religion, in which sacrifice, stigma, and behavioral restrictions screen low-commitment members and reduce free-riding when religious participation depends on hard-to-monitor effort and commitment. The mechanism here is different: the baseline model abstracts from rival religious suppliers and outside religious options, and ritual complexity has no direct resource cost, so the standard screening or costly-signaling logic is absent. Complexity is costly only indirectly by lowering perceived execution success and current effort and is chosen by an informed authority to trade off current losses against slower belief erosion. See also Iyer (2016) for a broader survey of the economics of religion.

[4] Relatedly, Coyne and Mathers (2011) provide a rational-choice account of rituals, emphasizing identity, signaling, and common knowledge as functional channels behind ritual emergence; path dependence, switching costs, and identity attachments as sources of ritual persistence; and new alternatives or changing identity payoffs as sources of ritual change. My mechanism is different: ritual dynamics arise from a single dynamic rent-extraction problem faced by an informed authority, who uses ritual complexity to preserve belief and the rents it supports by shaping how failures are interpreted.

say that "a spirit is watching" to motivate good behavior,[5] or a society may sustain an "American Dream" narrative[6] that keeps people striving even when their circumstances would otherwise discourage them, while blaming them when they fall short. What matters is not religion itself, but monopolistic control over a procedure whose true relevance is difficult for outsiders to verify.

The rest of the paper proceeds as follows. Section 2 sets up the production environment, workers' perceived technology, and Bayesian updating under misspecification. Section 3 characterizes the shaman's optimal ritual-complexity policy and the resulting belief dynamics. Section 4 studies technological progress and shows how belief-based motivation becomes redundant, generating secularization from within.

2. The Model Setup

2.1 Environment

Consider an economy with a continuum of workers indexed by $i \in [0,1]$ and a single shaman. Both discount future payoffs at rate $\delta \in (0,1)$ over an infinite horizon $t = 0,1,2,...$ The shaman is the sole religious authority, whom workers regard as the intermediary who controls access to the ritual procedure required to obtain divine favor. Workers believe that perfect performance of the prescribed ritual raises their return to effort and thereby the probability of a successful harvest outcome.

In fact however, the ritual has no divine productive effect and influences output only by inflating workers' perceived marginal returns to effort. [7] The shaman knows that the ritual

[5] A parent who cannot constantly monitor a child may create a myth "a spirit is watching" that induces good behavior. When contrary evidence appears, the parent can add qualifications that preserve belief but weaken the myth's clarity and motivational force.

[6] Similarly, secular mobility narratives such as the American Dream can motivate effort while interpreting failure as imperfect individual execution rather than as evidence against the promise itself.

[7] "Magic enables man to carry out with confidence his important tasks, to maintain his poise and his mental integrity in fits of anger, in the throes of hate, of unrequited love, of despair and anxiety. The function of magic is to ritualize man's optimism, to enhance his faith in the victory of hope over fear. Magic expresses the greater value for man of confidence over doubt, of steadfastness over vacillation, of optimism over pessimism." Following Malinowski (1948), magic can be understood as occupying the space where empirical knowledge and practical technique reach their limits. In the model, the shaman's claim to mediate divine favor

code is entirely under his control and has no true divine effect, whereas workers enter each period with prior belief $\rho_t \in (0,1)$ that the prescribed ritual is divinely authorized and genuinely raises the probability of a successful harvest if perfectly executed.

The shaman's surplus extraction rests entirely on an informational advantage generated by his institutional monopoly[8] over the "divine production process." [9] This institutional monopoly turns belief into an incentive system: by controlling the ritual channel through which divine favor is supposedly accessed, the shaman induces workers to exert more effort than the true technology alone would justify.

2.2 Workers

Inherited belief ρ_0

All workers share a common inherited prior $\rho_0 \in (0,1)$ that the myth M, the theological narrative underpinning the shaman's authority, is true. Let $\rho_t \in (0,1)$ denote the common public belief at the beginning of period t. Workers update this belief after observing harvest outcomes according to Bayes' rule. The current belief $\rho_t \in (0,1)$ is the sole state variable in the model.

Assumption: Bayesian Workers under a Misspecified Model

Workers are Bayesian within their subjective model. They correctly update beliefs about the shaman's authority after observing harvest outcomes, but their likelihood function is misspecified: they believe that divine favor can supplement the technological return to effort. They also interpret ritual complexity as a change in divine demands rather than as a strategic

performs the same function: it raises workers' perceived control over production precisely in the residual domain of uncertainty left by technology.

[8] See Weingast (2015) on Adam Smith's analysis of the medieval Church as a monopoly provider of religious services whose authority over the masses generated rents and political power.

[9] In Malinowski (1948)'s account, magical knowledge is 'occult, taught through mysterious initiations, handed on in a hereditary or at least in very exclusive filiation,' even though many members of the community participate in the associated gardening, fishing, and other rituals under the magician's direction.

choice by the shaman. Thus, workers may update rationally from observed outcomes while still learning within an incorrect causal model. [10]

Remark This assumption separates learning from outcomes from learning from the ritual script itself: workers revise beliefs after observing harvest success or failure, but not merely from observing changes in ritual complexity. [11]

2.3 The Shaman's Instrument

Tribute rate τ

Let $\tau \in (0,1)$ define the tribute rate, defined as the share of harvest value claimed by the shaman each period. For simplicity, it is treated as exogenous. Since the mass of workers is normalized to one, aggregate expected output is equal to per-worker expected output.

Script complexity x_t

Let $x_t \geq 0$ A denote the complexity of the ritual script prescribed in period t. A higher x_t corresponds to a more elaborate ritual. Each period, the shaman chooses x_t to maximize the present discounted value of expected tribute revenue:

$$\Pi = \Sigma_{t=0}^{\infty} \delta^{t} \cdot \tau \cdot v \cdot E[Y_t],$$

Where $\delta \in (0,1)$ is a time-discounter factor, τ is the tribute rate, $v > 0$ is the value of a successful harvest, and since $Y_t \in \{0,1\}$ is a binary outcome, $E[Y_t]$ is equivalent to the probability of harvest success.

Remark This is not ex-post excuse-making. The shaman sets ritual complexity *before* the harvest, conditional on the current belief stock, shaping how informative the next signal will be rather than rationalizing the outcome after the fact. The behavioral restriction is that workers update from harvest outcomes but not from the script itself: they read variation

[10] Malinowski (1948) writes that magic "never originated" and "never has been made or invented." This is consistent with the assumption that workers view the shaman as a transmitter, rather than inventor, of ritual requirements: changes in the script are read as divine demands, not strategic manipulation.

[11] Workers read an elaborated script as a change in divine demands rather than as manipulation because their inherited causal model bundles three claims: divine favor exists, ritual is its channel, and the shaman transmits the spirit's demands. Workers can accept or reject this bundle, but the strategic hypothesis lies outside their model

in complexity as variation in divine demand, not institutional manipulation. This follows from their inherited subjective model, which is what renders their learning misspecified.

2.4 Technology

True Success Probability $p(e_t)$

The ritual economy produces a common aggregate binary harvest outcome $Y_t \in \{0,1\}$, where $Y_t = 1$ denotes a successful harvest and $Y_t = 0$ denotes failure. Conditional on worker effort $e_t \geq 0$, the true probability of success is

$$p_t \;=\; p(e_t),$$

Where p_t satisfies the following assumptions:

- $p(0) = 0$: without effort, there is no harvest.
- $p(e_t) \to \bar{p} \in [0,1]$, as $e_t \to \infty$: as effort becomes arbitrarily large, the probability approaches an upper bound $\bar{p}$.
- $p'(e_t) > 0$: more effort strictly raises the probability of success,
- $p''(e_t) < 0$: the return to effort is strictly diminishing.

Functional Form. For tractability, I parametrize $p(e_t) = z(1 - e^{-\alpha e_t})$, where $z \in (0,1)$ defines the technological limit and $\alpha > 0$ governs the speed of diminishing returns. This satisfies all four assumptions above. The parameter z has a natural interpretation: it is the maximum harvest probability attainable by pure technology alone, so $1 - z \in (0,1)$ is the residual gap that workers may believe divine favor can fill.

Assumption: Common aggregate shock

Workers are ex ante identical and face a common aggregate harvest shock. In a symmetric equilibrium, all workers choose the same effort e_t. The ritual economy then produces a single publicly observed binary harvest outcome $Y_t \in \{0,1\}$ common to all workers. Thus, the continuum of workers does not generate a cross-section of infinite signals. All workers observe the same aggregate harvest outcome and update the common public belief.

Ritual Execution Probability $s(x_t)$

Let $x_t \geq 0$ denote the procedural complexity of the ritual script prescribed in period t.[12] A higher x_t corresponds to a more demanding ritual. Each period, the shaman chooses x_t to maximize the present discounted value of expected tribute revenue. Workers assign probability

$$s_t = s(x_t)$$

to satisfactory performance of the ritual in period t. The function satisfies:

- $s(0) = 1$: a maximally simple ritual is executed perfectly with probability one.
- $s(x_t) \to 0$, as $x_t \to \infty$: a maximally complex ritual is executed perfectly with probability zero
- $s' < 0$: more complex rituals are harder to execute satisfactorily.

Functional Form. I set $s(x_t) = e^{-\gamma x_t},\ \gamma > 0$, which satisfies all assumptions above. Note that $s_t = 0$ corresponds to the limiting choice of $x_t = \infty$. One natural interpretation: if a ritual of complexity requires a sequence of x_t conditions each satisfied with probability $\theta \in (0,1)$, independently of one another, then $s(x_t) = \theta^{x_t} = e^{-\gamma x_t}$ with $\gamma = -\ln\theta > 0$. The function is declining and convex. Each additional condition reduces execution probability by $|s(x_t + 1) - s(x_t)| = \theta^{x_t}(1 - \theta)$, which shrinks to zero as $x_t \to \infty$. Each successive requirement has a smaller effect than the last: as perfect execution becomes less likely, there is simply less probability left to lose.

Assumption: Latent Ritual Execution

Let $R_t \in \{0,1\}$ denote whether to have satisfactory ritual execution and $Pr(\, R_t = 1 \mid x_t \,) = s(x_t)$ denote the probability of satisfactory ritual execution. For $x_t > 0$, R_t is latent and not

[12] To give content to "complexity," I rely on Wood's (1986) taxonomy of task complexity. Wood argues that tasks become more complex when they require more acts (*component complexity*), more precise coordination among those acts (*coordinative complexity*), or changing requirements over time (*dynamic complexity*). In the ritual setting, these dimensions correspond to additional ritual steps; stricter rules about their order, timing, performers, and combination; and state-contingent requirements, such as state-contingent requirements, such as newly announced divine demands in my setting. I summarize these procedural demands in a single variable, x_t since the model focuses on how ritual complexity affects execution probability and belief updating rather than on the internal sources of ritual complexity. High ritual complexity makes imperfect execution a reasonable expectation rather than an exceptional event, allowing workers to attribute failure to deficient performance of the script rather than to the absence of divine favor.

directly observed by workers. When $x_t = 0$, satisfactory execution is certain $s(0) = 1$, thus $Pr(\, R_t = 1 \mid 0\,) = s(0) = 1$.

Remark This assumption makes ritual complexity a source of deniability: after a failed harvest, workers cannot tell whether failure occurred because the myth was false or because the ritual was imperfectly executed.

Effort-Contingent Divine Supplement and Perceived Success Probability

Workers believe that divine favor, when accessed through satisfactory ritual performance, increases the return to their own effort. Let

$$q_t = q(e_t)$$

denote the effort-contingent component of the perceived divine supplement in period t. The function q satisfies:

- $q(0) = 0$: without effort, divine favor generates no perceived additional chance of harvest success.
- $q(x_t) \to \bar{q} \in [0,1]$, as $q \to \infty$: as effort becomes arbitrarily large, the perceived divine supplement approaches an upper bound $\bar{q}$.
- $q' > 0$: greater effort increases the perceived benefit from divine favor.
- $q'' < 0$: the perceived divine return to effort is diminishing.

Thus q_t captures the part of divine favor that workers believe operates through effort rather than replacing it.

In the perceived success probability,

$$\pi^A(e_t, x_t, \rho_t) = \underbrace{p_t(e_t)}_{\text{True return}} + \underbrace{\rho_t \, s(x_t) \, q(e_t)}_{\text{Perceived divine supplement}},$$

where ρ_t measures belief in the shaman's authority, $s(x_t)$ measures the perceived probability of satisfactory ritual execution, and $q(e_t)$ measures the perceived divine return attached to effort.

Remark The divine supplement $\rho_t \, s(x) \, q(e_t)$ has no true productive effect but increases perceived marginal returns to effort. This captures the idea that belief in divine favor raises effort by making workers perceive production as more responsive to their actions than it actually is.

Assumption on Divine Favor $q(e_t)$

To ensure that the perceived success probability is well-defined, assume that

$$p(e_t) + q(e_t) \leq 1, \text{ for all } e_t \geq 0.$$

For simplicity, we normalize the upper bound so that

$$\lim_{e_t \to \infty} [p(e_t) + q(e_t)] = 1, \text{ or } \bar{p} + \bar{q} = 1.$$

Thus, $q(e_t)$ can be interpreted as the divine component that is perceived to fill the residual gap left by technology. Technology alone generates success probability $p(e_t)$, while divine favor is believed to supplement production only in the part of success that technology cannot fully control.

Functional Form. I set $q(e_t) = (1-z)(1-e^{-\alpha e_t})$, so that divine favor responds to effort with the same diminishing returns as technology, governed by α. This implies $p(e_t) + q(e_t) = 1 - e^{-\alpha e_t} \in [0,1)$ holds for all $e_t \geq 0$, ensuring the perceived success probability $\pi^A(e_t, x_t, \rho_t)$ is well defined. The normalization $\bar{p} + \bar{q} = z + (1-z) = 1$ divides the maximum attainable success probability between technology (fraction z) and divine favor$(1-z)$.

2.5 Worker's Utility

Let

$$\tilde{v} \equiv (1-\tau)v$$

denote the net-of-tribute value of a successful harvest. Worker i's perceived flow utility in period t is

$$u_{it} = \tilde{v}\, \pi^A(e_t, x_t, \rho_t) - c(e_t),$$

where $c(0) = 0, c' > 0$. The first term $\tilde{v}\, \pi^A(e_t, x_t, \rho_t)$ is the worker's perceived expected net-of-tribute harvest value. The second term $c(e_t)$ is the direct cost of effort for harvest.

Functional Form. I set $c(e_t) = e_t$, so effort cost is linear. Combined with the strict concavity of $\pi^A(e_t, x_t, \rho_t)$ *in* e_t, this guarantees a unique interior optimum and yields the closed-form optimal interior effort rule, making its comparative statics fully transparent. [13]

[13] If instead $c(e_t) = e_t^2$, the worker's first-order condition can be rearranged as $(\alpha e_t)e^{\alpha e_t} = \frac{v z \alpha^2}{2}[z + \rho_t e^{-\gamma x_t}(1-z)]$, whose solution requires the Lambert W function with $e_t^* = \frac{1}{\alpha} W\big(\frac{v z \alpha^2}{2}[z +$

2.6 Timeline

Each period t proceeds as follows:

(i) The shaman observes the current public belief ρ_t and announces a ritual script of complexity $x_t \geq 0$.

(ii) After observing x_t, each worker chooses effort e_t^* to maximize its perceived expected utility:

$$e_t^* \in \arg\max_{e_t \geq 0} \{\tilde{v}\pi_t^A(e_t, x_t, \rho_t) - c(e_t)\}.$$

(iii) A common aggregate harvest outcome $Y_t \in \{0,1\}$ is realized and publicly observed by all workers.

(iv) Public belief is updated from ρ_t to ρ_{t+1} according to Bayes' rule.

2.7 Equilibrium Objects Under the Parametric Specification

Perceived Productivity Index $\Phi(\rho_t, x_t)$

Substituting the functional forms into the perceived success probability $\pi^A(e_t, x_t, \rho_t)$ and collecting the common term yields

$$\pi^A(e_t, x_t, \rho_t) = (1 - e^{-\alpha e_t})\Phi(\rho_t, x_t),$$

where $\Phi(\rho_t, x_t) \equiv z + \rho_t e^{-\gamma x_t}(1 - z) \in [z, 1]$ is perceived productivity index. It equals 1 when $\rho_t = 1$ and $x_t = 0$ (full belief and a simple ritual) and approaches to z as $\rho_t \to 0$ (divine favor either ceases to be believed) or $x_t \to \infty$ (divine favor becomes inaccessible through an impossibly demanding ritual).

The marginal perceived return to effort is

$$\frac{\partial \pi^A}{\partial e_t} = \alpha e^{-\alpha e_t}\Phi\big(\rho_t, x_t\big) > 0.$$

$\rho_t e^{-\gamma x_t}(1 - z)]\big)$. Note all qualitative results continue to hold, but the comparative statics lose their transparent closed-form expressions.

Since $\frac{\partial^2 \pi^A}{\partial {e_t}^2} = -\alpha^2 e^{-\alpha e_t} \Phi\big(\rho_t, x_t\big) < 0, t$he perceived return is strictly concave in effort. Two comparative statics follow from Φ. First, stronger belief shifts the perceived return upward, raising the marginal incentive to exert effort at every effort level; formally, $\frac{\partial\, \Phi(\rho_t, x_t)}{\partial \rho_t} > 0$. Second, greater ritual complexity lowers the perceived return, suppressing workers' effort incentives and reducing current output, which is the cost of complexity that the shaman must weigh against its benefit as a reputation insurance device; formally, $\frac{\partial\, \Phi(\rho_t, x_t)}{\partial x_t} < 0$.

Myth Premium

The myth premium measures the inflation of workers' perceived harvest success probability relative to the true probability:

$$m_t(x_t, \rho_t) = \frac{\pi^A(e_t, x_t, \rho_t) - p_t(e_t)}{p_t(e_t)} = e^{-\gamma x_t} \rho_t \frac{1-z}{z}.$$

If $m_t = 0.5$, workers acts as if their chances are 50% higher than they really are. When the myth premium m_t is large, workers are heavily leveraged on belief. Three forces govern the size of the premium. First, m_t is increasing in ρ_t: false belief amplifies effort most when workers already trust the shaman deeply; a more faithful congregation (with a higher ρ_t) is more inflated. Second, m_t is decreasing in x_t: a harder ritual is executed perfectly with lower probability, so divine favor is rarely accessed and the myth's perceived contribution shrinks. Third, m_t is increasing in $\frac{1-z}{z}$: as technology becomes more powerful, the residual gap the myth can claim credit for narrows, and the premium vanishes as $z \to 1$.

Workers' Optimal Effort e_t^*

The myth premium $m_t(x_t, \rho_t)$ means workers perceive a higher marginal return to effort than technology alone delivers, inducing them to exert strictly more effort than the true production function would justify. Each worker chooses e_t to maximize perceived utility: $u_{it} = \tilde{v}\, \pi^A(e_t, x_t, \rho_t) - c(e_t)$. The first-order condition gives:

$$e_t^* = \frac{1}{\alpha} \ln(\tilde{v} \alpha \Phi(\rho_t, x_t)).$$

An interior solution requires $\tilde{v}\alpha\Phi(\rho_t, x_t) > 1$; otherwise, perceived returns are too low to justify any effort. Letting $A_t \equiv \tilde{v}\alpha\Phi(\rho_t, x_t)$, three comparative statics follows.

DEFINITION 3.3.1 (Active State and Active Region)

Belief $\rho_t \in (0,1)$ *is an active state if there exists a finite* $x_t \in [0,\infty)$ *such that* $A_t \equiv \tilde{v}\alpha\Phi(\rho_t, x_t) > 1$. *The active region is the set of all active states:* $\mathcal{A} \equiv \{\rho_t \in (0,1): \exists\, x_t \in [0,\infty)$ such that $A_t \equiv \tilde{v}\,\alpha\,\Phi(\rho_t, x_t) > 1\}$.

LEMMA 2.7.1 (The Worker's Optimal Effort Rule)

The worker exerts positive effort if and only if $A_t \equiv \tilde{v}\alpha\Phi(\rho_t, x_t) > 1$ *The equilibrium effort rule is* $e_t^* = \frac{1}{\alpha} max\{ln(\tilde{v}\alpha\Phi(\rho_t, x_t)), 0\}$. *When* $A_t \equiv \tilde{v}\alpha\Phi(\rho_t, x_t) > 1$, (1) $\frac{\partial e_t^*}{\partial \rho_t} > 0$, *stronger belief raises* Φ *and thereby raises effort; (2)* $\frac{\partial e_t^*}{\partial x_t} < 0$, *higher ritual complexity lowers* Φ *and thereby lowers effort; and (3)* $\frac{\partial e_t^*}{\partial \tilde{v}} > 0$, *a higher net-of-tribute harvest value raises the return to effort and increases effort.*

True Expected Output E[Y_t]

Although workers exert effort in response to inflated perceived returns, the harvest outcome is governed by the true production function $p(e_t) = z(1 - e^{-\alpha e_t})$, not by the myth. Substituting the optimal effort rule $e_t^* = \frac{1}{\alpha}\ln(\tilde{v}\alpha\Phi(\rho_t, x_t))$,

$$\mathrm{E}[Y_t] = p_t^* = z(1 - e^{-\alpha e_t^*}) = z(1 - \frac{1}{\tilde{v}\alpha\Phi(\rho_t, x_t)}) = z(1 - \frac{1}{A_t}).$$

Two comparative statics follow. First, $\frac{\partial E[Y_t]}{\partial \tau} < 0$: a higher tribute reduces $\tilde{v}$, shrinking workers' incentives, suppressing effort, and lowering true output. Second, $\frac{\partial E[Y_t]}{\partial \rho_t} > 0$: stronger belief raises the perceived productivity Φ, induces more effort, and thereby raises true output even though the belief itself has no direct productive effect.

PROPOSITION 2.7.1 (Religious Belief Raises Output) *Stronger belief in the shaman raises true expected output, even though the belief is itself false:*

$$\frac{\partial E[Y_t]}{\partial \rho_t} > 0.$$

REMARK. This is the model's core placebo result: the myth has no direct productive content but raises real output entirely through the effort channel. Workers perceive a higher marginal return, exert more effort, and generate more output than the true technology alone would sustain—but only for as long as the belief holds.

Belief Updating

After observing the harvest outcome $Y_t \in \{0,1\}$, workers update ρ_t by Bayes' rule within their misspecified model. The two posteriors are

$$\rho_{t+1}^{+} = \frac{\rho_t \, \pi^A(e_t, x_t, \rho_t = 1)}{\rho_t \, \pi^A(e_t, x_t, \rho_t = 1) + (1 - \rho_t) \, p_t},$$

$$\rho_{t+1}^{-} = \frac{\rho_t[1 - \pi^A(e_t, x_t, \rho_t = 1)]}{\rho_t[1 - \pi^A(e_t, x_t, \rho_t = 1)] + (1 - \rho_t)(1 - p_t)},$$

where $\rho_{t+1}^{+} > \rho_t > \rho_{t+1}^{-}$. Note $\pi^A(e_t, x_t, \rho_t = 1)$ is the perceived success likelihood under the myth-true hypothesis. Belief rises after a successful harvest and falls after failure. The shaman uses x_t to manage how informative each outcome is about his authority. Two propositions characterize this mechanism.

PROPOSITION 2.7.2 (Complexity Creates Deniability After Failure)[14]

A more demanding script is unlikely to be executed perfectly, so workers attribute failure to imperfect ceremony rather than to the myth's falsity. A bad harvest is less diagnostic of the shaman's legitimacy when ritual requirements are stringent. Formally, for $e_t^* > 0$ and $\rho_t > 0$,

$$\frac{\partial \rho_{t+1}^{-}}{\partial x_t} > 0.$$

Proof

Recall the posterior belief after a harvest failure

$$\rho_{t+1}^{-} = \frac{\rho_t[1 - \pi^A(e_t, x_t, \rho_t = 1)]}{\rho_t[1 - \pi^A(e_t, x_t, \rho_t = 1)] + (1 - \rho_t)(1 - p_t)}.$$

Differentiate with respect to x_t,

$$\frac{\partial \rho_{t+1}^{-}}{\partial x_t} = \frac{\rho_t(1 - \rho_t)\left[-(1 - p_t)\dfrac{\partial \pi_A(e_t, x_t, \rho_t = 1)}{\partial x_t} + (1 - \pi_A)\dfrac{\partial p_t}{\partial x_t}\right]}{\left[\rho_t[1 - \pi^A(e_t, x_t, \rho_t = 1)] + (1 - \rho_t)(1 - p_t)\right]^2}.$$

The denominator is strictly positive, so the sign of $\frac{\partial \rho_{t+1}^{-}}{\partial x_t}$ is determined by the sign of

[14] Malinowski (1948) emphasizes that magical efficacy depends on exact spell, rite, and taboo observance; if any condition is neglected, failure follows.

$$\left[-(1-p_t)\frac{\partial \pi^A(e_t, x_t, \rho_t = 1)}{\partial x_t} + (1-\pi_A)\frac{\partial p_t}{\partial x}\right].$$

It can be shown that

$$\pi^A(e_t, x_t, \rho_t = 1) = (1 - \tfrac{1}{\tilde{v}\alpha\Phi(\rho_t=1, x_t)})\Phi(\rho_t = 1, x_t) = \Phi(\rho_t = 1, x_t) - \tfrac{1}{\tilde{v}\alpha},$$

$$p_t = z(1 - e^{-\alpha e_t}) = z(1 - \frac{1}{\tilde{v}\alpha\Phi(\rho_t, x_t)}) = z - \frac{z}{\tilde{v}\alpha\Phi(\rho_t, x_t)}.$$

It follows that

$$\frac{\partial \pi_A}{\partial x_t} = \frac{\partial \Phi\big(\rho_t = 1, x_t\big)}{\partial x_t} = \frac{\partial[z + e^{-\gamma x_t}(1-z)]}{\partial x_t} = -\gamma e^{-\gamma x_t}(1-z) < 0,$$

$$\frac{\partial p_t}{\partial x_t} = \frac{z}{\tilde{v}\alpha}\frac{1}{\big[\Phi\big(\rho_t, x_t\big)\big]^2}\frac{\partial\big[\Phi\big(\rho_t, x_t\big)\big]}{\partial x_t} = -\frac{z(1-z)\gamma\rho_t e^{-\gamma x_t}}{\tilde{v}\alpha\big[\Phi\big(\rho_t, x_t\big)\big]^2} < 0.$$

Substitute,

$$\left[-(1-p_t)\frac{\partial \pi^A(e_t, x_t, \rho_t = 1)}{\partial x_t} + (1-\pi_A)\frac{\partial p_t}{\partial x}\right]$$

$$= (1-p_t)\gamma e^{-\gamma x_t}(1-z) - (1-\pi_A)\frac{z(1-z)\gamma\rho_t e^{-\gamma x_t}}{\tilde{v}\alpha[\Phi(\rho_t, x_t)]^2}.$$

Collecting the common factor,

$$= \gamma e^{-\gamma x_t}(1-z)\left[1 - p_t - \frac{z\rho_t(1-\pi_A)}{\tilde{v}\alpha[\Phi(\rho_t, x_t)]^2}\right].$$

$\gamma e^{-\gamma x_t}(1-z) > 0$, so the sign reduces to $\big[1 - p_t - \frac{z\rho_t(1-\pi_A)}{\tilde{v}\alpha[\Phi(\rho_t, x_t)]^2}\big]$.

Since $(1-p_t) = 1 - z(1 - \frac{1}{\tilde{v}\alpha\Phi(\rho_t, x_t)})$,

$$1 - p_t - \frac{z\rho_t(1-\pi_A)}{\tilde{v}\alpha[\Phi(\rho_t, x_t)]^2}$$

$$= 1 - z\left[1 - \frac{1}{\tilde{v}\alpha\Phi(\rho_t, x_t)}\right] - z\left[\frac{\rho_t(1-\pi_A)}{\tilde{v}\alpha[\Phi(\rho_t, x_t)]^2}\right]$$

$$= 1 - z\left[1 - \frac{1}{\tilde{v}\alpha\Phi(\rho_t, x_t)} - \frac{\rho_t(1-\pi_A)}{\tilde{v}\alpha[\Phi(\rho_t, x_t)]^2}\right]$$

Substitute $A_t = \tilde{v}\alpha\Phi(\rho_t, x_t)$,

$$= 1 - z\left[1 - \frac{1}{A_t} - \frac{\rho_t(1-\pi_A)}{A_t\Phi(\rho_t, x_t)}\right]$$

$$= 1 - z \left[1 - \frac{1}{A_t} \left[1 + \frac{\rho_t (1 - \pi_A)}{\Phi(\rho_t, x_t)} \right] \right]$$

$$= 1 - z + \frac{z}{A_t} \left[1 + \frac{\rho_t (1 - \pi_A)}{\Phi(\rho_t, x_t)} \right]$$

$$\geq 1 - z > 0.$$

Therefore,

$$\frac{\partial \rho^-_{t+1}}{\partial x_t} > 0.$$

■

PROPOSITION 2.7.3 (Complexity Dilutes Credit After Success)

Workers attribute a good harvest to their own effort rather than divine favor when the ritual is unlikely to have been executed perfectly. A good harvest is therefore less diagnostic of the shaman's legitimacy. Formally, for $e^*_t > 0$ and $\rho_t > 0$,

$$\frac{\partial \rho^+_{t+1}}{\partial x_t} < 0.$$

Proof

Recall the posterior belief after a successful harvest:

$$\rho^+_{t+1} = \frac{\rho_t \, \pi^A(e_t, x_t, \rho_t = 1)}{\rho_t \, \pi^A(e_t, x_t, \rho_t = 1) + (1 - \rho_t) \, p_t}.$$

Substituting $\pi^A(e_t, x_t, \rho_t) = (1 - e^{-\alpha e_t})\Phi(\rho_t, x_t)$, and $p_t = z(1 - e^{-\alpha e_t})$ and canceling the common factor $(1 - e^{-\alpha e_t})$,

$$\rho^+_{t+1} = \frac{\rho_t \, \Phi(\rho_t = 1, x_t)}{\rho_t \, \Phi(\rho_t = 1, x_t) + (1 - \rho_t) \, z}.$$

Differentiating with respect to x_t,

$$\frac{\partial \rho^+_{t+1}}{\partial x_t} = \frac{\rho_t \big(1 - \rho_t\big) z \dfrac{\partial \Phi(\rho_t = 1, x_t)}{\partial x_t}}{\big[\rho_t \, \Phi + \big(1 - \rho_t\big) z\big]^2}.$$

The denominator $[\rho_t \, \Phi + (1 - \rho_t)z]^2 > 0$, so the sign is determined by $\frac{\partial \Phi(\rho_t = 1, x_t)}{\partial x_t}$.

Since $\Phi(\rho_t = 1, x_t) = z + e^{-\gamma x_t}(1 - z)$,

$$\frac{\partial \Phi(\rho_t = 1, x_t)}{\partial x_t} = -\gamma(1-z)e^{-\gamma x_t} < 0.$$

Therefore,

$$\rho^+_{t+1} = \frac{\rho_t \, \pi^A(e_t, x_t, \rho_t = 1)}{\rho_t \, \pi^A(e_t, x_t, \rho_t = 1) + (1-\rho_t)\, p_t} < 0.$$

■

COROLLARY 2.7.1 (Ritual Complexity Smooths Belief Updating)

Greater ritual complexity compresses the spread of posterior beliefs as a belief smoother that slows the speed at which any single harvest outcome revises workers' faith, in either direction. Formally, *for* $e^*_t > 0$ *and* $\rho_t > 0$,

$$\frac{\partial(\rho^+_{t+1} - \rho^-_{t+1})}{\partial x_t} < 0.$$

REMARK. Complexity operates symmetrically: it dampens the upward revision after success and the downward revision after failure. Only the latter generates a net benefit for the shaman; dampening the downward revision slows blame after failure. The former is a cost: the shaman wants workers to attribute good harvests to his divine authority, but complexity undermines that attribution. This asymmetry defines the tradeoff that the optimal complexity policy in section 3.

Proof

By Proposition 2.7.3,

$$\frac{\partial \rho^+_{t+1}}{\partial x_t} < 0.$$

By Proposition 2.7.2,

$$\frac{\partial \rho^-_{t+1}}{\partial x_t} > 0.$$

Therefore,

$$\frac{\partial(\rho^+_{t+1} - \rho^-_{t+1})}{\partial x_t} = \frac{\partial \rho^+_{t+1}}{\partial x_t} - \frac{\partial \rho^-_{t+1}}{\partial x_t} < 0.$$

■

3. Equilibrium

3.1 Markov Perfect Equilibrium

A Markov Perfect Equilibrium consists of $(x_t^*,\ {e_t}^*, \rho_{t+1}^{\pm})$

(1) **Shaman's ritual policy** $x_t^* \equiv x^*: \rho_t \in [0,1] \to \mathbb{R}_+$ to solve the following Bellman equation:

$$V(\rho_t) = \max_{x_t}\{\tau \cdot v \cdot p_t^* + \delta[p_t^* V(\rho_{t+1}^+) + (1-p_t^*)V(\rho_{t+1}^-)]\}$$

(2) **Each attached worker's effort rule** ${e_t}^*(\rho_t, x_t^*, \tilde{v}\ ) = \frac{1}{\alpha}\ln(\tilde{v}\alpha\Phi(\rho_t, x_t^*))$ solving

$$\max_{e_t}\{\tilde{v}\,\pi^A(e_t, x_t^*) - c(e_t)\}$$

(3) **Workers' Beliefs in the Shaman** are updated using Bayes' rule:

$$\rho_{t+1}^+ = \frac{\rho_t\,\pi^A(e_t^*)}{\rho_t\,\pi^A(e_t^*) + (1-\rho_t)\,p_t^*} = \frac{\rho_t\,\Phi}{\rho_t\,\Phi + (1-\rho_t)z}$$

$$\rho_{t+1}^- = \frac{\rho_t[1-\pi^A(e_t^*)]}{\rho_t[1-\pi^A(e_t^*)] + (1-\rho_t)(1-p_t^*)}$$

3.2 Existence and Uniqueness of the Value Function

Let $B\big((0,1)\big)$ be the Banach space of bounded, continuous functions on (0,1). The Bellman operator $\mathcal{T}: \mathrm{B}\big((0,1)\big) \to \mathrm{B}\big((0,1)\big)$ is defined by the following:

$$(\mathcal{T}W)(\rho_t) \equiv \max_{x_t}\{\,\tau \cdot \tilde{v} \cdot\ p_t^* + \delta[p_t^* W(\rho_{t+1}^+) + (1-p_t^*)W(\rho_{t+1}^-)]\}$$

PROPOSITION 3.2.1(Existence and Uniqueness of the Value Function)

There exists a unique value function $V \in B\big((0,1)\big)$ satisfying $\mathcal{T}V = V$. The value function V is also weakly increasing in ρ_t.

REMARK. No matter how optimistic or pessimistic the shaman's initial guess is about V, iterating the Bellman equation converges to the unique true value.

Proof

[Step 1] Claim: $\mathcal{T}$ maps $B([\mathbf{0},\mathbf{1}])$ into itself.

For any $W \in B([\mathbf{0},\mathbf{1}])$, the inner objective is continuous in (ρ_t, x_t) and uniformly bounded by $\tau v + \|W\|_\infty$. Therefore, $\mathcal{T}W \in B([\mathbf{0},\mathbf{1}])$.

[Step 2] Claim: $\mathcal{T}$ is a contraction mapping.

Blackwell's conditions hold: (i) Monotonicity: if $W \leq W'$, then $\mathcal{T}W \leq \mathcal{T}W'$, since $p_t^*, (1-p_t^*) \geq 0$. (ii) Discounting: $\mathcal{T}(W+c)(\rho_t) = \mathcal{T}W(\rho_t) + \delta c$, where $\delta \in (0,1)$, since $p_t^*, +(1-p_t^*) = 1$. Together, these imply $\mathcal{T}$ is a contraction mapping with modulus δ. By the Banach Fixed-Point Theorem, there exists a unique value function $\mathrm{V} \in \mathrm{B}\big((0,1)\big)$ satisfying $\mathcal{T}V = V$.

[Step 3] Claim: V is increasing in ρ_t.

Fix any ritual complexity policy x_t. A higher prior ρ_t raises the myth premium, which increases optimal effort e_t^* and raises the probability of a successful harvest. Note under the same policy x_t and outcome path, the shaman with a higher prior ρ_t obtains higher revenue in every period. Since the optimal value function $V(\rho_t)$ is the supremum over all feasible policies, $V(\rho_t)$ is increasing in ρ_t.

■

COROLLARY 3.2.1(Bound on the Value Function)

The unique fixed point $V \in B\big((0,1)\big)$ satisfies for all $\rho \in [0,1]$ and a fixed $z \in (0,1)$,

$$0 \leq V(\rho; z) \leq \frac{\tau \tilde{v} z}{1-\delta}.$$

REMARK The shaman can always walk away and collect zero tribute, so his value is never negative. $\frac{\tau\tilde{v}}{1-\delta}$ is what the shaman could ever earn, in the best possible world with a perfectly captive congregation (full belief $\rho = 1$, simple ritual $x = 0$).

3.3 Optimal Ritual Complexity Policy

The shaman manages a belief stock over an infinite horizon, using complexity to slow the rate at which reality reveals itself. Ritual complexity is therefore a strategic tool for reputation insurance.

LEMMA 3.3.1 (Infinite Complexity Is Never Optimal in an Active State)

Let $\rho_t \in (0,1)$. If ρ_t is an active state, then the optimal complexity is $x_t^* < \infty$.

Proof.

Suppose $x_t^* = \infty$.

As $x_t \to \infty$,

$$\Phi(\rho_t, \infty) = z + \rho_t \cdot 0 \cdot (1 - z) = z.$$

It follows that

$$e_t^{*,\infty} = \frac{1}{\alpha}\max\{\ln(\tilde{v}\alpha z), 0\}, \; p_t^{*,\infty} = z\big(1 - e^{-\alpha e_t^{*,\infty}}\big),$$

$$\pi_A(e_t^{*,\infty}, \infty, 1) = \big(1 - e^{-\alpha e_t^{*,\infty}}\big)[z + 0 \cdot (1 - z)] = p_t^{*,\infty}.$$

Bayes' rule then gives $\rho_{t+1} = \rho_t$ after every outcome. By induction, $\rho_{t+k} = \rho_t$ for all $k \geq 1$. Then The shaman collects $\tau v p_t^{*,\infty}$ each period, so $V(\rho_t) = \tau v p_t^{*,\infty} + \delta V(\rho_t) \;\Rightarrow\; V(\rho_t) = \frac{\tau v p_t^{*,\infty}}{1-\delta}$.

Because ρ_t is an active state, there exists $\widehat{x_t} < \infty$ with $\tilde{v}\alpha\Phi(\rho_t, \widehat{x_t}) > 1$. If $\tilde{v}\alpha z \geq 1$, then $\hat{p}_t^* > p_t^{*,\infty}$. If $\tilde{v}\alpha z < 1$, then $\hat{p}_t^* > p_t^{*,\infty} = 0$. In both cases, $\hat{p}_t^* > p_t^{*,\infty}$.

Consider the one-period deviation: play $\widehat{x_t}$ today, then revert to $x_T = \infty$ for all $T > t + 1$.

Since the relevant payoff $\hat{W}$ is $\tau v \hat{p}_t^* + \frac{\delta}{1-\delta}\tau v p_t^{*,\infty}$ under that one-time deviation,

$$\hat{W} - V = \left[\tau v \hat{p}_t^* + \frac{\delta}{1-\delta}\tau v p_t^{*,\infty} - \frac{1}{1-\delta}\tau v p_t^{*,\infty}\right] = \tau v \hat{p}_t^* - \tau v p_t^{*,\infty} = \tau v(\hat{p}_t^* - p_t^{*,\infty}) > 0,$$

, which contradicts the optimality of $x_t^* = \infty$.

■

COROLLARY 3.3.1 (Uniform Bounds on Compact Active-Belief Regions).

Let $K = [a, b]$ with $0 < a \leq b < 1$ be contained in the active region. Then there exist constants $\bar{x}_K < \infty$ and $\underline{e}_K > 0$ such that

$$\sup_{\rho \in [a,b]} x^*(\rho) \leq \bar{x}_K < \infty \;\text{ and }\; \inf_{\rho \in [a,b]} e^*(\rho, x^*(\rho)) \geq \underline{e}_K > 0.$$

Proof.

Define $s_t \equiv e^{-\gamma x_t} \in (0,1]$ for $x_t \in [0, \infty)$. As $x_t \to \infty$, $s_t \to 0$. Including this limiting case, the full action space $s_t \in [0,1]$ is a compact set.

Define $G([a, b]) \equiv \{(\rho_t, s_t) \in [a, b] \times [0,1]: \tau v p_t^*(\rho_t, s_t) + \delta \mathrm{E}[V(\rho_{t+1}) \mid \rho_t, s_t] = V(\rho_t)\}$. This is the set of all belief-complexity pairs that are optimal for the shaman.

Define $f(\rho_t, s_t) \equiv \tau v p_t^*(\rho_t, s_t) + \delta E[V(\rho_{t+1}) \mid \rho_t, s_t] - V(\rho_t)$. The functions $\Phi(\rho_t, s_t)$, e_t^*, p_t^*, $V(\rho_t)$are all continuous in $(\rho_t, s_t) \in [a,b] \times [0,1]$. Therefore, $f(\rho_t, s_t)$ is also continuous in $(\rho_t, s_t) \in [a,b] \times [0,1]$. Since $G([a,b]) = f^{-1}(\{0\})$ is the preimage of the closed set $\{0\}$ under the continuity of f, $G([a,b])$ is also closed. A closed subset of the compact set $[a,b] \times [0,1]$ is compact, so $G([a,b])$ is also compact. Every $\rho_t \in [a,b]$ lies in the active region by assumption. By Lemma just above, the optimal choice at every active belief satisfies $x_t^* < \infty$, so $s_t^*(\rho_t) = e^{-\gamma x_t^*(\rho_t)} > 0$. Therefore every element $\big(\rho_t, s_t^*(\rho_t)\big) \in G([a,b])$ satisfies $s_t^*(\rho_t) > 0$.

Suppose that $\inf_{(\rho_t, s_t) \in G([a,b])} s_t = 0$. Then there exists a sequence $\big\{({\rho_t}^{(n)}, {s_t}^{(n)})\big\}_{n=1}^{\infty} \in G([a,b])$ for all $n \geq 1$ with ${s_t}^{(n)} \to 0$. Since $G([a,b])$ is compact, this sequence must have a convergent subsequence with its limit$(\rho_t^*, s_t^*) \in G([a,b])$, where $s_t^* = 0$. This contradicts $s_t^*(\rho_t) = e^{-\gamma x_t^*(\rho_t)} > 0$. Therefore, $s_K \equiv \min\{\, s_t \colon (\rho_t, s_t) \in G([a,b])\} > 0$. Every optimal choice satisfies $e^{-\gamma x_t^*(\rho_t)} \geq s_K$. It follows that $x_t^*(\rho_t) \leq -\frac{1}{\gamma} \ln s_K \;\equiv\; \bar{x}_K < \infty \quad$ for all $\rho_t \in [a,b]$.

Since $\rho_t \in [a,b]$ is contained in the active region, every $\rho_t \in [a,b]$ satisfies $A_t = \tilde{v}\alpha\Phi(\rho_t, s_t) > 1$. Therefore, he effort rule gives $e_t^*(\rho_t, s_t) = \frac{1}{\alpha} \ln A_t > 0 \quad$ for every $(\rho_t, s_t) \in G([a,b])$.

Suppose that $\inf_{(\rho_t, s_t) \in G([a,b])} e_t^* = 0$. Then there exists a sequence $\big\{({\rho_t}^{(m)}, {s_t}^{(m)})\big\}_{m=1}^{\infty} \in G([a,b])$ for all $m \geq 1$ with ${s_t}^{(m)} \to 0$. Since $G([a,b])$ is compact, this sequence has a convergent subsequence with limit $({\rho_t}^{**}, {s_t}^{**}) \in G([a,b])$. There must be $({\rho_t}^{**}, {s_t}^{**})$ in $G([a,b])$ with the limiting effort of zero, which contradicts that $\rho_t \in [a,b]$ is contained in the effort active region. Therefore, $e_K \equiv \min_{(\rho_t, s_t) \in G([a,b])} e_t^*(\rho, s_t) > 0 \quad$ for all $\rho_t \in [a,b]$.

PROPOSITION 3.3.2 (Existence of an Interior Optimal Ritual Complexity)

Fix active $\rho \in \mathcal{A}$ with the floor condition $L_0(\rho) = \delta(\tilde{v}\alpha)^2\rho^2(1-\rho) > 1$ (as in Appendix Lemma 4.2.E.2). Then $x \in (0, \infty)$.

PROOF.

By Lemma 3.3.1, $x = \infty$ is never optimal. By Lemma 4.2.E.2, $x = 0$ is not optimal under the floor condition. The optimal interior complexity must lie in $x \in (0, \infty)$.

■

3.4 Equilibrium Belief Dynamics

The optimal interior complexity policy $x_t^*(\rho_t)$ partially insulates the shaman from harvest failures but cannot eliminate the tension between the myth and reality. This section characterizes the long-term dynamics of the belief process $\{\rho_t\}_{t\geq 0}$ under the equilibrium policy. We show first that belief erodes in expectation every period (Proposition 3.4.1), decompose the drift into three forces (Corollary 3.4.1), and then establish that collapse can be inevitable in the long run under certain cases (Proposition 3.4.2).

PROPOSITION 3.4.1 (Belief Erodes in Expectation Over Time)

Fix active belief $\rho_t \in (0,1)$. Suppose $z \in (0,1)$, and $x_t \in [0,\infty)$. Then, under the true data-generating process, the belief process $\{\rho_t\}_{t\geq 0}$ satisfies $E_t[\rho_{t+1}] < \rho_t$. The belief process $\{\rho_t\}_{t\geq 0}$ is a strict supermartingale. The downward drift equals

$$\Delta(\rho_t) \equiv \rho_t - E_t[\rho_{t+1}] = \frac{\rho_t^2(1-\rho_t)(\pi^A(e_t, x_t, \rho_t = 1) - p_t)^2}{\mathcal{D}^+\mathcal{D}^-} > 0,$$

where $\mathcal{D}^+ \equiv \rho_t\, \pi^A(e_t, x_t, \rho_t = 1) + (1-\rho_t)\, p_t$ and $\mathcal{D}^- \equiv \rho_t(1 - \pi^A(e_t, x_t, \rho_t = 1))) + (1-\rho_t)(1-p_t)$ are workers' perceived probabilities of success and failure, respectively.

Remark. The expectation is taken under the true data-generating process, not under workers' subjective predictive distribution. Workers themselves do not expect their belief to fall. Since they update by Bayes' rule within their own misspecified model, their posterior belief is a martingale under their subjective probability measure: the subjective probability of success is $\mathcal{D}^+$, the subjective probability of failure is $\mathcal{D}^-$, and $\mathcal{D}^+ \cdot \rho_{t+1}^+ + \mathcal{D}^- \cdot \rho_{t+1}^- = p_t$.

In contrast, under the true data-generating process, the distribution over harvest outcomes is $(p_t, 1-p_t)$, because the myth has no direct productive effect. When$e_t > 0$, $x_t < \infty$, and $z \in (0,1)$, workers assign too much probability to success relative to the true process, so $p_t < \mathcal{D}^+$. As a result, the belief-raising posterior ρ_{t+1}^+ is realized less often than workers expect, while the belief-lowering posterior ρ_{t+1}^- is realized more often. Therefore, belief erodes in expectation under the true process even though it is a martingale under workers' subjective model.

Proof

Define $\mathcal{D}^+ \equiv \rho_t\, \pi^A(e_t, x_t, \rho_t = 1) + (1 - \rho_t)\, p_t$ and $\mathcal{D}^- \equiv \rho_t[1 - \pi^A(e_t, x_t, \rho_t = 1)] + (1 - \rho_t)(1 - p_t)$, *where* $\mathcal{D}^+ + \mathcal{D}^- = 1$. Denote $\pi^A \equiv \pi^A(e_t, x_t, \rho_t = 1)$.

Under the true probability measure, $Y_t = 1$ with probability p_t and $Y_t = 0$ with probability $1 - p_t$. By iterated expectations,

$$E_t[\rho_{t+1}] = E_t\Big[E[\rho_{t+1}|Y_t]\Big] = p_t \cdot \rho^+_{t+1} + (1 - p_t) \cdot \rho^-_{t+1},$$

where the Bayesian posteriors are

$$\rho^+_{t+1} = \frac{\rho_t\, \pi^A}{\mathcal{D}^+},$$

$$\rho^-_{t+1} = \frac{\rho_t(1 - \pi^A)}{\mathcal{D}^-}.$$

Substituting,

$$E_t[\rho_{t+1}] = p_t \cdot \frac{\rho_t\, \pi^A}{\mathcal{D}^+} + (1 - p_t) \cdot \frac{\rho_t(1 - \pi^A)}{\mathcal{D}^-}.$$

Factoring ρ_t out,

$$E_t[\rho_{t+1}] = \rho_t \left\{\frac{p_t \pi^A}{\mathcal{D}^+} + \frac{(1 - p_t)(1 - \pi^A)}{\mathcal{D}^-}\right\}.$$

It is therefore sufficient to show that $\mathcal{D}^+\mathcal{D}^- > p_t\pi^A\mathcal{D}^- + (1 - p_t)(1 - \pi^A)\mathcal{D}^+$.

First, expand $\mathcal{D}^+\mathcal{D}^-$:

$$\begin{aligned}
&\mathcal{D}^+\mathcal{D}^- \\
&\quad = [\rho_t\, \pi^A + (1 - \rho_t)\, p_t][\rho_t(1 - \pi^A) + (1 - \rho_t)(1 - p_t)] \\
&\quad = {\rho_t}^2\pi^A(1 - \pi^A) + \rho_t(1 - \rho_t)\pi^A(1 - p_t) + p_t(1 - \rho_t)\rho_t(1 - \pi^A) + (1 - \rho_t)^2 p_t(1 - p_t) \\
&\quad = {\rho_t}^2\pi^A(1 - \pi^A) + \rho_t(1 - \rho_t)(\pi^A(1 - p_t) + p_t(1 - \pi^A)) + (1 - \rho_t)^2 p_t(1 - p_t).
\end{aligned}$$

Second, expand $p_t\pi^A\mathcal{D}^- + (1 - p_t)(1 - \pi^A)\mathcal{D}^+$:

$$\begin{aligned}
&p_t\pi^A\mathcal{D}^- + (1 - p_t)(1 - \pi^A)\mathcal{D}^+ \\
&\quad = p_t\pi^A\Big(\rho_t(1 - \pi^A) + (1 - \rho_t)(1 - p_t)\Big) + (1 - p_t)(1 - \pi^A)(\rho_t\, \pi^A + (1 - \rho_t)\, p_t) \\
&\quad = \rho_t p_t\pi^A(1 - \pi^A) + (1 - \rho_t)p_t\pi^A(1 - p_t) + \rho_t(1 - p_t)(1 - \pi^A)\pi^A \\
&\quad + (1 - \rho_t)p_t(1 - p_t)(1 - \pi^A) \\
&\quad = \rho_t\pi^A(1 - \pi^A) + (1 - \rho_t)p_t(1 - p_t).
\end{aligned}$$

Subtracting,

$$\mathcal{D}^+\mathcal{D}^- - [p_t\pi^A\mathcal{D}^- + (1-p_t)(1-\pi^A)\mathcal{D}^+]$$
$$= {\rho_t}^2\pi^A(1-\pi^A) + \rho_t(1-\rho_t)(\pi^A(1-p_t) + p_t(1-\pi^A)) + (1-\rho_t)^2 p_t(1-p_t)$$
$$- \rho_t\pi^A(1-\pi^A) - (1-\rho_t)p_t(1-p_t)$$
$$= \pi^A(1-\pi^A)(-\rho_t(1-\rho_t)) - (1-\rho_t)\rho_t p_t(1-p_t) + \rho_t(1-\rho_t)(\pi^A(1-p_t) + p_t(1$$
$$- \pi^A))$$
$$= \rho_t(1-\rho_t)\{-\pi^A(1-\pi^A) - p_t(1-p_t) + \pi^A(1-p_t) + p_t(1-\pi^A)\}$$
$$= \rho_t(1-\rho_t)\{(\pi^A)^2 - 2\pi^A p_t + {p_t}^2\}$$
$$= \rho_t(1-\rho_t)(\pi^A - p_t)^2.$$

Since $E_t[\rho_{t+1}] = \rho_t \left\{\frac{p_t\pi^A\mathcal{D}^- + \mathcal{D}^+(1-p_t)(1-\pi^A)}{\mathcal{D}^+\mathcal{D}^-}\right\}$

$$\Delta(\rho_t) \equiv \rho_t - E_t[\rho_{t+1}] = \rho_t\left\{\frac{\mathcal{D}^+\mathcal{D}^- - [p_t\pi^A\mathcal{D}^- + \mathcal{D}^+(1-p_t)(1-\pi^A)]}{\mathcal{D}^+\mathcal{D}^-}\right\}$$
$$= \rho_t\left\{\frac{\rho_t(1-\rho_t)(\pi^A - p_t)^2}{\mathcal{D}^+\mathcal{D}^-}\right\} = \frac{{\rho_t}^2(1-\rho_t)(\pi^A - p_t)^2}{\mathcal{D}^+\mathcal{D}^-}.$$

Since ${\rho_t}^2(1-\rho_t) > 0$ for $\rho_t \in (0,1), (\pi^A - p_t) > 0$ for $z \neq 1, and\ \mathcal{D}^+, \mathcal{D}^- > 0\ for\ e_t > 0,$

$$\Delta(\rho_t) \equiv \rho_t - E_t[\rho_{t+1}] = \frac{{\rho_t}^2(1-\rho_t)(\pi^A - p_t)^2}{\mathcal{D}^+\mathcal{D}^-} > 0.$$

■

Because the shaman's ritual inflates the perceived probability of success above its true level, harvest failures occur more frequently than workers anticipate. Each failure revises belief downward by more than each success revises it upward, producing a persistent downward drift in ρ_t. The magnitude of this drift is governed by three forces.

The Credibility Gap

The drift is increasing in $\pi^A(e_t, x_t, \rho_t = 1) - p_t$. The term $\pi^A(e_t, x_t, \rho_t = 1)$ is the myth's maximum potential promise: the highest success probability the myth could ever claim at given effort and ritual complexity. The gap $\pi^A(e_t, x_t, \rho_t = 1) - p_t$ therefore measures the full extent to which the myth overstates what technology delivers. The larger this overstatement, the more often failures occur relative to workers' expectations, and the faster belief erodes. Greater ritual complexity compresses the gap between promise and reality, slowing the rate at which workers learn the myth is false.

REMARK (Effort Amplifies the Credibility Gap) The credible gap $\pi^A(e_t, x_t, \rho_t = 1) - p_t = (1 - e^{-\alpha e_t}) \cdot e^{-\gamma x_t}(1 - z)$ is increasing in e_t. The myth is therefore most vulnerable precisely when it is working best. Strong belief drives high effort, which inflates the overstatement and makes each failure more damaging to the shaman's credibility. The myth's own success as a motivational device accelerates its eventual unraveling.

The Fragility of Certainty

The drift is decreasing in the variance of the workers' perceived outcome distribution $(\mathcal{D}^+, \mathcal{D}^-)$. When $\mathcal{D}^+$ is close to 1, workers expect success almost surely; a failure is a large surprise that produces a large downward revision. When $(\mathcal{D}^+, \mathcal{D}^-) = (1/2, 1/2)$, outcomes are maximally uncertain; no single harvest outcome becomes decisive. The per-period impact of each realization is therefore largest when workers hold the most confident priors and smallest when they hold the most uncertain ones. Greater ritual complexity pushes workers' perceived outcome distribution toward $(\mathcal{D}^+, \mathcal{D}^-) = (1/2, 1/2)$ reducing the downward revision triggered by each disappointing harvest.

The Belief-Stock Effect

The factor $\rho_t^2(1 - \rho_t)$ captures the sensitivity of the belief stock to new information. When the belief stock is very low (close to zero), workers are already deeply skeptical, so additional bad news has little faith left to erode. When the belief stock is very high (close to 1), the myth appears nearly unquestionable, and beliefs move only slightly at the extreme margin. Erosion is strongest at intermediate belief stocks, where workers retain enough openness to doubt that each disappointing harvest prompts sizable revisions.

Remark (Dogmatism, Unfalsifiability at Full Belief, and the Optimal Simplest Ritual) When $\rho_t = 1$, the belief state is absorbing regardless of the harvest outcome. If harvest is successful, faith is simply correct and evidence confirms it. If the harvest fails, the prior assigns no weight to the myth being false, so Bayes' rule has nothing to update. From the inside, the alternative world in which the myth is false does not exist, so a population at $\rho_t = 1$ behaves exactly as it would if the myth were true, not because the myth works, but because complete faith closes off the very possibility that it might be false. [15]

[15] The mode's logic suggests an interpretation of extreme information control in some closed regimes. In this interpretation, censorship functions not just by hiding bad signals but by keeping alternative social models

This is a degenerate limit in which the shaman can effectively choose $x_t = 0$, maximizing effort with the simplest possible ritual and no erosion of faith.

COROLLARY 3.4.1 (The Cost of Inflated Expectations).

For a given $\rho_t \in (0,1), z \neq 1$ and $e_t > 0$ the speed of belief erosion is jointly governed by three forces: (i) the credibility gap channel, whereby a larger gap between the myth's maximum potential promise and true technology (the full extent of the lie) accelerates erosion; (ii) the fragility-of-certainty channel, whereby failures are most damaging when success was expected almost surely; and (iii) the belief-stock effect, whereby intermediate conviction makes workers most responsive to disappointing outcomes. Moreover,

$$\frac{\partial \Delta(\rho_t)}{\partial x_t} < 0.$$

PROPOSITION 3.4.2 (Long-Run Collapse vs. Self-limiting Learning).

Let $\rho_0 \in (0,1)$, $z \in (0,1)$. The long-run behavior of the equilibrium belief process $\{\rho_t\}_{t=0}^{\infty}$ is determined by the parameters $\tilde{v}\alpha$ and z as follows:

(1) **(No-activity Regime)** If $\tilde{v}\alpha \leq 1$, then $e_t^* = 0$ and $\rho_t = \rho_0$ for all $t \geq 1$.

(2) **(Self-limiting Learning Regime)** If $1 < \tilde{v}\alpha < \frac{1}{z}$, then define

$$\underline{\rho} \equiv \frac{\frac{1}{\tilde{v}\alpha} - z}{1 - z} \in (0,1),$$

(Case 1: $\rho_0 \leq \underline{\rho}$)

$$e_t^* = 0 \text{ and } \rho_t = \rho_0 \text{ for all } t \geq 1.$$

(Case 2: $\rho_0 > \underline{\rho}$)

$\rho_t > \underline{\rho}$, $e_t^* > 0$, $x_t^* < \infty$, $\{e_t^*\} \to 0$, $\{p_t^*\} \to 0$, $\{x_t^*\} \to 0$, $\{\pi_t^A - p_t^*\} \to 0$.

(3) **(Persistent Learning and Collapse Regime)** If $\tilde{v}\alpha \geq \frac{1}{z}$, then

$$\rho_t \to 0.$$

RERMARK. $\underline{\rho}$ is the cutoff where the myth is just credible enough to sustain effort. Above it, workers believe the spirit might help, so they exert effort and the shaman keeps the

from becoming thinkable at all—effectively holding the belief state at $\rho_t = 1$. Once ρ_t falls just below 1, people begin to update their beliefs downward on average, and these small revisions, following the downward drift, can gradually accumulate into a spiral of disbelief.

system running. Below it, even with the simplest ritual executed perfectly, workers see that, even if the spirit helps, the return is too low to justify effort, and thus effort collapses. A false belief fully unravels only if production under positive effort continues to generate informative evidence as belief erodes. When belief erosion itself stops the evidence-generating activity, learning becomes self-limiting and a strictly positive residual belief persists. The learning stops early when technology is high

Proof.

By Proposition 3.4.1 $\{\rho_t\}$ is a bounded nonnegative supermartingale. By the martingale convergence theorem, $\{\rho_t\}$ therefore converges almost surely,

$$\rho_t \to \rho_\infty \in [0,1].$$

By Proposition 3.4.1 and including the case $e_t = 0$,

$$\Delta(\rho_t) \equiv \rho_t - E_t[\rho_{t+1}] = \frac{{\rho_t}^2(1-\rho_t)(\pi^A - p_t)^2}{\mathcal{D}^+\mathcal{D}^-} \geq 0.$$

Take expectation,

$$E(\Delta(\rho_t)) = E(\rho_t) - E(E_t[\rho_{t+1}]) = E(\rho_t) - E(\rho_{t+1}).$$

Summing from $t = 0$ to T,

$$\sum_{t=0}^{T} E\left[\Delta(\rho_t)\right] = E[\rho_0] - E[\rho_{T+1}] = \rho_0 - E[\rho_{T+1}].$$

Since $E[\rho_{T+1}] \geq 0$,

$$\sum_{t=0}^{T} E\left[\Delta(\rho_t)\right] = \rho_0 - E[\rho_{T+1}] \leq \rho_0.$$

Since $\Delta(\rho_t)$ is non-negative and $\sum_{t=0}^{T} E\left[\Delta(\rho_t)\right]$ is uniformly bounded by $\rho_0, \sum_{t=0}^{\infty} E\left[\Delta(\rho_t)\right]$ converges to a finite number. Therefore, it follows that

$$E\left[\sum_{t=0}^{\infty} \Delta\left(\rho_t\right)\right] = \sum_{t=0}^{\infty} E\left[\Delta(\rho_t)\right] \leq \rho_0 < \infty.$$

Since $0 \leq E\left[\sum_{t=0}^{\infty} \Delta\left(\rho_t\right)\right] < \infty$,

$$\Delta(\rho_t) \to 0.$$

[No Activity Regime]

Suppose $\tilde{v}\alpha \leq 1$. Since $\rho_t \leq 1$ and $e^{-\gamma x} \leq 1$, $\Phi(\rho_t, x_t) = z + \rho_t e^{-\gamma x_t}(1-z) \leq 1$. Then

$A_t = v\alpha\Phi(\rho_t, x_t) \leq 1$, and $e_t = 0$. Consequently, $p_t = \pi_t^A = 0$; failure occurs with probability 1, and there is no learning $\rho_t = \rho_{t+1}$. Therefore, $\rho_t = \rho_0$, for every $t \geq 1$.

[Persistent Learning and Collapse Regime]

Suppose $\tilde{v}\alpha \geq \frac{1}{z}$. For any $\rho_t \in [0,1]$ and $x_t < \infty$, $\Phi(\rho_t, x_t) = z + \rho_t e^{-\gamma x_t}(1-z) > z$. Since $A_t = \tilde{v}\alpha\Phi(\rho_t, x_t) > \tilde{v}\alpha z \geq 1$, optimal effort remains positive; $e_t > 0$. By Lemma 3.3.1, $x_t < \infty$ for the optimal effort $e_t > 0$.

Suppose $\rho_\infty = \rho^* \in (0,1)$. Then, there exists $\varepsilon > 0$ such that $[\rho^* - \varepsilon, \rho^* + \varepsilon] \subset (0,1)$. Apply Corollary 3.3.1. Then complexity is uniformly bounded above, and effort is uniformly bounded away from zero. It follows $0 < p_t^* < \pi_t^A$, ${\rho_t}^2(1-\rho_t) > 0$, and thereby $\Delta(\rho_t) > 0$ on $[\rho^* - \varepsilon, \rho^* + \varepsilon]$, which contradicts $\Delta(\rho_t) \to 0$. Therefore, $\rho_\infty \notin (0,1)$.

Since $\rho_0 \in (0,1), \{\rho_t\}$ is a strict supermartingale under positive effort, and effort is always positive in this regime, ρ_∞ couldn't be 1.

[Self-limiting Learning Regime]

Suppose $1 < \tilde{v}\alpha < \frac{1}{z}$. The largest possible perceived productivity is at $x_t = 0$, so $\Phi(\rho_t, 0) = z + \rho_t(1-z)$. Positive effort can be sustained at belief ρ_t if and only if $\tilde{v}\alpha[z + \rho_t(1-z)] > 1$. The cutoff $\underline{\rho} = \frac{\frac{1}{\tilde{v}\alpha} - z}{1-z} \in (0,1)$ is well-defined since $1 < \tilde{v}\alpha < \frac{1}{z}$.

Now consider two cases: (i) $\underline{\rho} \leq \rho_0$ and (ii) $\underline{\rho} > \rho_0$.

If $\underline{\rho} \geq \rho_0$, then no finite ritual can induce positive effort. Therefore, $e_t = 0$ for all $t \geq 1$. It follows that $p_t = \pi_t^A = 0$ and thereby $\Delta(\rho_t) = 0$, which implies $\rho_t = \rho_0$ for all $t \geq 1$.

Now suppose $\underline{\rho} < \rho_0$. Along the active part of the equilibrium path $\underline{\rho} < \rho_t$, $e_t > 0$ and there exists a compact set $K \subset (\rho_t - \varepsilon, \rho_t + \varepsilon)$ for small enough $\varepsilon > 0$ and $\rho_t \in (\underline{\rho}, 1)$. Therefore, as long as such compact set exists, the same learning force used in the persistent-collapse proof applies: beliefs drift downward until $\rho_t \leq \rho_0$. In the self-limiting learning regime, $\rho_t \to \underline{\rho}$, $e_t \to 0$, $p_t \to 0$, $A_t - p_t \to 0$.

■

Corollary 3.4.2 (Technology and the Stopping Point of Learning)

Suppose the economy is in the self-limiting learning regime $1 < \tilde{v}\alpha < \frac{1}{z}$. *Let* $\underline{\rho} \equiv \frac{\frac{1}{\tilde{v}\alpha} - z}{1-z} \in (0,1)$ *denote the limiting belief cutoff. Differentiating with respect to* z *leads to*

$$\frac{\partial \underline{\rho}}{\partial z} = \frac{\frac{1}{\tilde{v}\alpha} - 1}{(1-z)^2} < 0.$$

REMARK. When technology is poor, effort depends heavily on faith, so a modest decline in belief is enough to make workers stop exerting effort. Once effort stops, harvest outcomes no longer test the myth, and learning stops with a positive residual belief. Stronger technology has the opposite effect. Because production remains worthwhile even when faith is low, workers continue exerting effort, outcomes remain informative, and belief can erode further. Thus, weak technology preserves myth by ending experimentation early, while strong technology exposes myth to longer empirical testing.

4. Ritual Fills the Gap Where Technology Ends

PROPOSITION 4.1 (Self-Sufficiency Ceiling)

Fix any active belief $\rho \in \mathcal{A}$. *Set*

$$\tilde{\delta} = \frac{1}{1 + 4(2 + \frac{4\rho(\tilde{v}\alpha)^2}{(1-\ \rho)})} \leq \frac{1}{9}.$$

Then there exists $\bar{z} < 1$ *such that for* $\delta > \tilde{\delta}$, *and for all* $z \in (\bar{z}, 1)$, *the unique optimal complexity choice at belief* ρ *is* $x^* = 0$ *or* $s = 1$.

Remark. Once technology is strong enough that harvests can be sustained without meaningful help from ritual, the shaman optimally chooses zero complexity.

Proof

See appendix.

PROPOSITIN 4.2 (Technological Floor)

Fix active belief $\rho \in \mathcal{A}$. *Suppose the floor condition is satisfied;* $(\tilde{v}\alpha)^2\rho^2(1-\rho) > 1$. *If the shaman is sufficiently patient,* $\delta > \overline{\delta}(\rho) = \frac{1}{(\tilde{v}\alpha)^2\rho^2(1-\rho)}$, *as in Lemma 4.2.E.2, then there exists*

a threshold $\underline{z} > 0$ *such that for all* $z \in (0, \underline{z})$*, the optimal ritual complexity is strictly positive* $x^* > 0$.

PROOF.

By Lemma 4.2.A,

the normalized value $V_0(\rho_t)$ satisfies:

$$V_0(\rho) = \max_{s \in [0,1]} \{\tau \tilde{v} a_0(\rho, s) + \delta V^0(R_0^-(\rho, s))\}.$$

By Lemma 4.2.C, at $z = 0, x_t = 0$,

$$MC(\rho) = \frac{\gamma\tau}{\alpha\rho(1-\tau)} > 0.$$

By Lemma 4.2.D.2,

$$MB(\rho) \geq MC(\rho) \cdot L_0(\rho),$$

where $L_0(\rho) = (\tilde{v}\alpha)^2\rho^2(1-\rho)$.

For $\delta > \overline{\delta_1}$,

by Lemma 4.2.E.2, there exists ρ such that at $x = 0$,

$$L_0(\rho) > 1, \text{ and } MB(\rho) > MC(\rho).$$

By Lemma 4.2.B, as $z \to 0$,

$$\left\| \frac{V(\cdot; z)}{z} - V_0 \right\|_\infty \to 0,$$

, so there exists a small enough $\underline{z}$ such that for $z \in (0, \underline{z})$,

$$MB(\rho) > MC(\rho).$$

Therefore, for $z \in (0, \underline{z})$,

$$x_t^* > 0.$$

■

Theorem 4.1 (Ritual Fills the Gap Where Technology Ends)

Fix an active belief $\rho \in A$. Suppose $(\tilde{v}\alpha)^2\rho^2(1-\rho) > 1$. Then there exist the floor patience threshold $\overline{\delta}(\rho) = \frac{1}{(\tilde{v}\alpha)^2\rho^2(1-\rho)} \in (0,1)$ and the ceiling patience threshold $\widetilde{\delta}(\rho) = \frac{1}{1+4(2+\frac{4\rho(\tilde{v}\alpha)^2}{(1-\rho)})} \leq \frac{1}{9}$ such that

- If $\delta > \overline{\delta}(\rho)$, there exists $\underline{z}(\rho,\delta) = \frac{1}{(\tilde{v}\alpha)^2\rho^2(1-\rho)} > 0$ such that the optimal x_t is positive for all $z \in (0, \underline{z}(\rho))$.
- If $\delta > \widetilde{\delta}(\rho)$, there exists $\overline{z}(\rho) < 1$ the optimal x_t is zero for all $z \in (\overline{z}(\rho),1)$.

Remark (Technology, Ritual, and Secularization from Within).

Theorem 4 has a natural interpretation in terms of the gap between what technology can deliver and what people wish to believe. Let z denote the highest true success probability the economy can attain when workers exert maximum effort using only secular technology (no ritual). The residual $1 - z$ is then the unexplained space: the share of outcomes that technology and effort cannot account for, which ritual and divine favor can be plausibly imagined to fill.

For $z \in (0, \underline{z})$, technology is too weak to motivate sufficient effort on its own, and ritual becomes economically necessary. To sustain sufficient effort in future periods and to absorb blame for frequent failures that low technology makes inevitable, Theorem 4 implies that the optimal ritual complexity is strictly positive, $x > 0$. Once $z \in (\overline{z}, 1)$, Theorem 4 shows that optimal ritual complexity falls to zero; the economy no longer needs the myth to motivate effort.

This mirrors Malinowski (1948)' s observation that Trobriand fishermen used elaborate rites for risky ocean fishing but not for safe lagoon fishing: ritual elaborates where technique ends, and magic fills the gap where technology and knowledge end. In the model, the frontier $\overline{z}$ is the quantitative counterpart of this boundary. The shaman does not abandon complex ritual because he is "caught" by belief-erosion dynamics as in Section 3; Theorem 4's result applies for all active belief levels even close to 1. He abandons it because technology has made the myth redundant as a motivational device. Secularization here is a story of crowding out: as the technology frontier rises above the critical threshold, the residual gap the myth can plausibly claim credit for shrinks to nothing, and ritual authority dissolves from within.

Conclusion

This paper offers a theory of ritual life cycles grounded in a single economic problem: why a monopolistic authority may have an incentive to create and sustain a false but productive belief, and why it may eventually stop maintaining it. The belief is false in its object but productive in its effects: it does not change the true technology, but it inflates perceived returns to effort and thereby changes behavior.

The model explains three stages through this single mechanism. Ritual originates as a belief-based incentive device when technology is too weak to motivate effort on its own. Complexity expands as reputation insurance: by making perfect execution harder to achieve, it allows failure to be attributed to imperfect ceremony rather than to the falsity of the myth, slowing the rate at which reality reveals itself. Technological progress can make effort productive on its own, rendering belief-based motivation redundant. The authority then stops investing in the myth's reputation and allows belief to erode; not because the myth is exposed from outside, but because it no longer earns its keep.

The central result is secularization from within under monopoly. The same authority that built the myth eventually has reason to let it fade. Secularization therefore need not require religious competition or external exposure. It can arise through quiet crowding out: as technology replaces the incentive function that belief once performed, the authority's reason to manage belief disappears with it.

Although the model is framed around religious ritual, the mechanism is more portable. More generally, an authority defines a path to success, monopolizes access to that path, and interprets failure as imperfect compliance rather than as evidence that the path itself is irrelevant. What matters is not religion itself, but monopolistic control over a prescribed procedure for achieving success. A ritual, in the sense of this model, is any procedure whose requirements are defined and managed by a concentrated authority.

The deeper point is that a monopolist does not merely allocate access to an existing path; it can define what counts as the relevant path in the first place and collect rents from those who seek access to it. Similar structures appear in organizations, bureaucracies, professions, and political institutions wherever advancement, legitimacy, or belonging depends

on procedures set by a concentrated few whose true relevance to success is difficult for outsiders to verify.[16]

One natural direction for future work is to ask what happens when the monopoly over interpretation breaks down. In such an environment, belief-based rents may not need to disappear; they may instead be reorganized across multiple authorities, each claiming the capacity to interpret the same world and to define the proper path to success. The Protestant Reformation is suggestive in this respect: viewed through the lens of the model, it illustrates a tree splitting rather than dying.[17] A formal extension could allow multiple belief suppliers to compete over ritual complexity, doctrinal interpretation, and access rules, and then study how competition reshapes the dynamics of belief, effort, and rent extraction.

[16] This broader interpretation is related to Meyer and Rowan's (1977) account of formal organizational structures as myths and ceremonies. In their account, organizations adopt formal rules, procedures, and vocabularies to secure legitimacy and resources, even when these structures are only loosely connected to technical efficiency. The mechanism here is different but complementary. In Meyer and Rowan, institutionalized structures are taken-for-granted myths that organizations ceremonially adopt; here, an informed authority strategically manages a procedure to shape how failures are interpreted, preserve belief in the procedure's relevance, and sustain rents.

[17] Doctrines such as *sola scriptura* and the priesthood of all believers weakened the claim that access to divine truth and favor had to pass through a single institutional hierarchy, giving rise to multiple competing authorities rather than the disappearance of religious authority itself.

REFERENCES

Bénabou, Roland, and Jean Tirole. "Mindful Economics: The Production, Consumption, and Value of Beliefs." *Journal of Economic Perspectives* 30, no. 3 (Summer 2016): 141–64.

Coyne, Christopher J., and Rachel L. Mathers. "Rituals: An Economic Interpretation." *Journal of Economic Behavior and Organization* 78 (2011): 74–84.

Iannaccone, Laurence R. "Sacrifice and Stigma: Reducing Free-riding in Cults, Communes, and Other Collectives." *Journal of Political Economy* 100 (1992): 271–91.

Iyer, Sriya. "The New Economics of Religion." *Journal of Economic Literature* 54 (2016): 395–441.

Malinowski, Bronislaw. *Magic, Science and Religion and Other Essays*. Selected, and with an introduction by Robert Redfield. Glencoe, IL: The Free Press; Boston, MA: Beacon Press, 1948.

Meyer, John W., and Brian Rowan. "Institutionalized Organizations: Formal Structure as Myth and Ceremony." *American Journal of Sociology* 83 (1977): 340–63.

Harvey, Eric. *Reading Creation Myths Economically in Ancient Mesopotamia and Israel*. Cambridge: Cambridge University Press, 2025.

Singh, Manvir. "The Cultural Evolution of Shamanism." *Behavioral and Brain Sciences* 41 (2018): e66.

Weingast, Barry R. Adam Smith's Industrial Organization of Religion: Explaining the Medieval Church's Monopoly and Its Breakdown in the Reformation. Working paper, Stanford University, 2015.

Wood, Robert E. "Task Complexity: Definition of the Construct." Organizational Behavior and Human Decision Processes 37 (1986): 60–82.

Appendix Proof

A.1 Proof for Proposition 4.1

LEMMA 4.1 (Uniform Lipschitz Bound on p^*)

For all $\rho \in (0,1)$, $s \in (0,1]$, *and* $z \in (0,1)$, *the true success probability* $p^*(\rho, s)$ *satisfies:*

$$\frac{\partial p^*}{\partial \rho} \leq c_1(1-z), \qquad \frac{\partial p^*}{\partial s} \leq c_1(1-z),$$

where $c_1 = \frac{1}{\tilde{v}\alpha z}$. *Therefore,*

$$Lip_\rho(p^*) \leq c_1(1-z), \qquad Lip_s(p^*) \leq c_1(1-z).$$

PROOF

Differentiate p^* in ρ,

$$\frac{\partial p^*}{\partial \rho} = \frac{zs(1-z)}{\tilde{v}\alpha\Phi^2}, \qquad \frac{\partial p^*}{\partial s} = \frac{z\rho(1-z)}{\tilde{v}\alpha\Phi^2}.$$

Since $\rho \leq 1$, and $s \leq 1$,

$$zs \leq z, \quad z\rho \leq z.$$

Since $\Phi(\rho, s) = z + \rho s(1-z) \geq z$,

$$\Phi^2 \geq z^2, \qquad \frac{1}{z^2} \geq \frac{1}{\Phi^2}.$$

Therefore,

$$\frac{\partial p^*}{\partial \rho} = \frac{zs(1-z)}{\tilde{v}\alpha\Phi^2} \leq \frac{z(1-z)}{\tilde{v}\alpha z^2} = \frac{1}{\tilde{v}\alpha z}(1-z),$$
$$\frac{\partial p^*}{\partial s} = \frac{z\rho(1-z)}{\tilde{v}\alpha\Phi^2} \leq \frac{z(1-z)}{\tilde{v}\alpha z^2} = \frac{1}{\tilde{v}\alpha z}(1-z).$$

■

COROLLARY 4.1.1 (Cost Bound of Complexity)

For any $\rho \in (0,1), s \in (0,1), z \in (\frac{1}{2}, 1)$,

$$|p^*(\rho, 1) - p^*(\rho, s)| \leq c_1(1-z)(1-s).$$

COROLLARY 4.1.2 (Lipschitz in ρ)

For any $\rho_1, \rho_2 \in (0,1)$,

$$|p^*(\rho_1, s) - p^*(\rho_2, s)| \leq c_1(1-z)|\rho_1 - \rho_2|.$$

LEMMA 4.2 (Posterior Shift Bounds)

Fix $\rho \in (0,1), s \in (0,1]$, *and* $z \in (\frac{1}{2}, 1)$,

$$|\rho^+(s) - \rho^+(1)| \leq M^+(1-z)(1-s),$$
$$|\rho^-(s) - \rho^-(1)| \leq M^-(1-z)(1-s),$$

where $M^+ = 2$, *and* $M^- \equiv \frac{4\rho(\tilde{\nu}\alpha)^2}{(1-\rho)}$. *On any compact subset of active state set* $\rho \in [\underline{\rho}, \overline{\rho}] \subset \mathcal{A}$, M^- *can be bounded above by a constant depending only on* $\overline{\rho}$.

PROOF

After a successful harvest,

$$\rho^+(\rho, s) = \frac{\rho \cdot \Phi(1,s)}{\rho \cdot \Phi(1,s) + (1-\rho) \cdot z},$$

where $\Phi(1,s) = z + s(1-z)$.

Differentiate in s,

$$\frac{\partial \rho^+}{\partial s} = \frac{\rho(1-z) \cdot (1-\rho)z}{(\rho\Phi(1,s) + (1-\rho)z)^2}.$$

Since $z \leq \rho\Phi(1,s) + (1-\rho)z$,

$$\frac{\partial \rho^+}{\partial s} = \frac{\rho(1-z) \cdot (1-\rho)z}{(\rho\Phi(1,s) + (1-\rho)z)^2} \leq \frac{\rho(1-z) \cdot (1-\rho)z}{z^2} = \frac{\rho(1-z) \cdot (1-\rho)}{z}.$$

Since $\rho(1-\rho) \leq 1$,

$$\frac{\partial \rho^+}{\partial s} \leq \frac{1-z}{z} \equiv M^+(1-z),$$

, where $M^+ \equiv \frac{1}{z} \leq 2$ for $z \in (\frac{1}{2}, 1)$.

Integrating over $[s, 1]$,

$$|\rho^+(s) - \rho^+(1)| = \int_s^1 \frac{\partial \rho^+}{\partial \xi}\, d\xi \leq \int_s^1 M^+(1-z)\, d\xi = M^+(1-z)(1-s).$$

Similarly, after a failed harvest,

$$\rho^{-}(\rho, s) = \frac{\rho \cdot [1 - \Phi(1, s)]}{\rho \cdot [1 - \Phi(1, s)] + (1 - \rho) \cdot (1 - z)}.$$

Differentiate in s,

$$\frac{\partial \rho^{-}}{\partial s} = -\frac{\rho(1 - \ \rho)(1 - z)\left[(1 - p) - \frac{z\rho(1 - \pi^{A})}{\tilde{v}\alpha[\Phi(\rho, s)]^{2}}\right]}{[\rho(1 - \pi^{A}) + (1 - \rho)(1 - p)]^{2}} < 0,$$

where $(1 - p) - \frac{z\rho(1-\pi^{A})}{\tilde{v}\alpha[\Phi(\rho,s)]^{2}}) < (1 - p) < 1$ from Proposition 2.7.2.

Since $\rho(1 - \pi^{A}) + (1 - \rho)(1 - p) \geq \frac{z(1-\rho)}{\tilde{v}\alpha}$ and $(1 - p) - \frac{z\rho(1-\pi^{A})}{\tilde{v}\alpha[\Phi(\rho,s)]^{2}}) < 1$,

$$\left|\frac{\partial \rho^{-}}{\partial s}\right| = \frac{\rho(1 - \ \rho)(1 - z)\left[(1 - p) - \frac{z\rho(1 - \pi^{A})}{\tilde{v}\alpha[\Phi(\rho, s)]^{2}}\right]}{[\rho(1 - \pi^{A}) + (1 - \rho)(1 - p)]^{2}} \leq \frac{\rho(1 - \ \rho)(1 - z)}{\left[\frac{z(1 - \rho)}{\tilde{v}\alpha}\right]^{2}} \leq \frac{\rho(\tilde{v}\alpha)^{2}}{z^{2}(1 - \ \rho)}(1 - z).$$

Since $\frac{\rho(\tilde{v}\alpha)^{2}}{z^{2}(1-\ \rho)} \leq M^{-} \equiv \frac{4\rho(\tilde{v}\alpha)^{2}}{(1-\ \rho)}$ for $z \in (\frac{1}{2}, 1)$,

$$\left|\frac{\partial \rho^{-}}{\partial s}\right| \leq \frac{\rho(\tilde{v}\alpha)^{2}}{z^{2}(1 - \ \rho)}(1 - z) \leq M^{-}(1 - z).$$

Integrating over $[s, 1]$,

$$|\rho^{-}(s) - \rho^{-}(1)| = \int_{s}^{1} \frac{\partial \rho^{-}}{\partial \xi}\, d\xi \leq \int_{s}^{1} M^{-}(1 - z)\, d\xi = M^{-}(1 - z)(1 - s)$$

■

LEMMA (Posterior Curvature)

Fix any active $\rho_t \in \mathcal{A}$ and a finite $x_t < \infty$. The Bayesian posteriors satisfy

$$\frac{\partial^{2} \rho_{t+1}^{+}}{\partial {\rho_t}^{2}} < 0, \qquad \frac{\partial^{2} \rho_{t+1}^{-}}{\partial {\rho_t}^{2}} > 0.$$

PROOF

Since ρ_{t+1}^{+} can be simplified as follows,

$$\rho_{t+1}^{+} = \frac{\rho_t\, \Phi(\rho_t = 1, x_t)}{\rho_t\, \Phi(\rho_t = 1, x_t) + (1 - \rho_t)\, z}.$$

Since $\rho_t\,\Phi(1,x_t) + (1-\rho_t)\,z = \Phi(\rho_t, x_t)$,

$$\rho_{t+1}^{+} = \frac{\rho_t\,\Phi(1,x_t)}{\Phi(\rho_t,x_t)}.$$

Differentiate with respect to ρ_t,

$$\frac{\partial \rho_{t+1}^{+}}{\partial \rho_t} = \frac{z\Phi(1,x_t)}{\Phi(\rho_t,x_t)^2}.$$

Differentiate once more,

$$\frac{\partial^2 \rho_{t+1}^{+}}{\partial {\rho_t}^2} = -\frac{2z(1-z)\Phi(1,x_t)e^{-\gamma x_t}}{\Phi(\rho_t,x_t)^2} < 0.$$

Similarly, it can be shown that

$$\frac{\partial^2 \rho_{t+1}^{-}}{\partial {\rho_t}^2} > 0.$$

■

LEMMA 4.3 (Lipschitz Bound on V).

Let $\mathcal{A} \subset (0,1)$ denote the active region, and let $K = [\rho_a, \rho_b] \subset \mathcal{A}$ be a compact belief interval. Then there exist constants $z_K \in (1/2,1)$ and $C_K < \infty$ such that for all $z \in (z_K, 1)$ and all $\rho_1, \rho_2 \in K$,

$$|V(\rho_1; z) - V(\rho_2; z)| \le C_K(1-z)|\rho_1 - \rho_2|.$$

The Lipschitz constant of V on K is first order in the technological gap $(1-z)$.

PROOF

Recall the Bellman operator

$$(TW)(\rho) = \max_{x \ge 0}\{\tau v\, p^*(\rho,s) + \delta[p^*(\rho,s)\, W\big(\rho^+(\rho,s)\big) + \big(1-p^*(\rho,s)\big)\, W\big(\rho^-(\rho,s)\big)]\},$$

where $s = e^{-\gamma x} \in (0,1]$ and V is the unique fixed point $V = TV$.

Let $\mathrm{Lip}_K(W)$ denote the Lipschitz constant of W on K. By Lemma 4.1, for every $s \in (0,1]$, and every $\rho_1, \rho_2 \in K$, $|p^*(\rho_2,s) - p^*(\rho_1,s)| \le c_1(1-z)|\rho_2 - \rho_1|$ for some finite constant c_1.

The flow term $\tau v\, p^*(\rho,s)$ is Lipschitz on K with constant at most $\ell_0(1-z), \quad \ell_0 \equiv \tau v c_1$.

Since $z \in (1/2,1)$, $c_1 = \frac{1}{\bar{v}\alpha z}$ is uniformly bounded, so ℓ_0 can be taken independent of z.

Next, for any Lipschitz function W, define

$$G_s(\rho) = p^*(\rho, s)W\big(\rho^+(\rho, s)\big) + \big(1 - p^*(\rho, s)\big)W\big(\rho^-(\rho, s)\big).$$

Decompose the difference $G_s(\rho_2) - G_s(\rho_1)$ into a weight term and two posterior-location terms:

$$\begin{aligned} G_s(\rho_2) - G_s(\rho_1) = &\underbrace{[p^*(\rho_2, s) - p^*(\rho_1, s)][W(\rho_2^+) - W(\rho_2^-)]}_{\text{(A) weight term}} \\ &+ \underbrace{p^*(\rho_1, s)[W(\rho_2^+) - W(\rho_1^+)]}_{\text{(B) success-posterior term}} \\ &+ \underbrace{\big(1 - p^*(\rho_1, s)\big)[W(\rho_2^-) - W(\rho_1^-)]}_{\text{(C) failure-posterior term}}. \end{aligned}$$

By Lemma 4.1,

$$|p^*(\rho_2, s) - p^*(\rho_1, s)| \le c_1(1 - z)|\rho_2 - \rho_1|.$$

Since $|W(\rho_2^+) - W(\rho_2^-)| \le \mathrm{Lip}_K(W)$,

$$\underbrace{[p^*(\rho_2, s) - p^*(\rho_1, s)][W(\rho_2^+) - W(\rho_2^-)]}_{\text{(A) weight term}} \le c_1(1 - z)\mathrm{Lip}_K(W)|\rho_2 - \rho_1|.$$

By the Posterior Curvature lemma, the posterior maps $\rho \mapsto \rho^{\pm}(\rho, s)$ are twice continuously differentiable in ρ on the active region $\mathcal{A}$. On the compact set K, their first derivatives are continuous and bounded. Therefore, there exists $\kappa_K < \infty$ such that, for $\rho \in K$ and all $s \in$ (0,1),

$$\left|\frac{\partial \rho^{\pm}}{\partial \rho}(\rho, s)\right| \le \kappa_K.$$

By the mean value theorem,

$$|\rho^{\pm}(\rho_2, s) - \rho^{\pm}(\rho_1, s)| \le \kappa_K|\rho_2 - \rho_1|.$$

Since W is Lipschitz on K,

$$|W\big(\rho^{\pm}(\rho_2, s)\big) - W\big(\rho^{\pm}(\rho_1, s)\big)| \le \mathrm{Lip}_K(W)|\rho^{\pm}(\rho_2, s) - \rho^{\pm}(\rho_1, s)| \le \kappa_K\mathrm{Lip}_K(W)|\rho_2 - \rho_1|.$$

Since weights $p^*(\rho_1, s)$ and $\big(1 - p^*(\rho_1, s)\big)$ are between zero and one,

$$|(B)| + |(C)| \le \kappa_K\,\mathrm{Lip}_K(W)|\rho_2 - \rho_1|.$$

Putting together,

$$|G_s(\rho_2) - G_s(\rho_1)| \le [\kappa_K + c_1(1-z)]\mathrm{Lip}_K(W)|\rho_2 - \rho_1|.$$

So for each $s \in (0,1]$,

$$\mathrm{Lip}_K(G_s) \le \eta_K(z)\,\mathrm{Lip}_K(W), \quad \eta_K(z) \equiv \kappa_K + c_1(1-z).$$

Because $\left|\frac{\partial \rho^{\pm}}{\partial \rho}(\rho, s) = \right| \le 1$, as z goes to 1,

$$\eta_K(z) = \kappa_K + c_1(1-z) \le 1 + c_1(1-z).$$

For a given W and s, the inner objective $\rho \mapsto \tau vp^*(\rho, s) + \delta G_s(\rho)$ is Lipschitz on K with constant at most $\ell_0(1-z) + \delta\eta_K(z)\,\mathrm{Lip}_K(W)$.

Taking the maximum over s, for every W,

$$\mathrm{Lip}_K(TW) \le \ell_0(1-z) + \delta\eta_K(z)\,\mathrm{Lip}_K(W).$$

Now define $W_0(\rho) \equiv 0$ and $W_{n+1} = TW_n$. Then $\mathrm{Lip}_K(W_0) = 0$, and $\mathrm{Lip}_K(W_{n+1}) \le \ell_0(1-z) + \delta\eta_K(z)\,\mathrm{Lip}_K(W_n)$.

Iterating,

$$\mathrm{Lip}_K(W_n) \le \ell_0(1-z)\sum_{j=0}^{n-1}[\delta\eta_K(z)]^j.$$

For z sufficiently close to 1,

$$\eta_K(z) \le 1 + c_1(1-z), \text{ so } \delta\eta_K(z) < 1.$$

Therefore,

$$\mathrm{Lip}_K(W_n) \le \frac{\ell_0(1-z)}{1-\delta\eta_K(z)}.$$

By Proposition 3.2.1, as $n \to \infty$,

$$W_n = T^n W_0 \to V.$$

Taking the limit, as $n \to \infty$,

$$\mathrm{Lip}_K(V) \le \frac{\ell_0(1-z)}{1-\delta\eta_K(z)}.$$

Since $\eta_K(z) \to 1$ as $z \to 1$,

choose z_K close enough to 1 that $1 - \delta\eta_K(z) \ge \frac{1-\delta}{2}$ for all $z \in (z_K, 1)$, giving

$$\mathrm{Lip}_K(V) \le \frac{2\ell_0}{1-\delta}(1-z) \equiv C_K(1-z)$$

Then, for $z \in (z_K, 1)$, and for any $\rho_1, \rho_2 \in K$

$$|V(\rho_1; z) - V(\rho_2; z)| \le C_K(1-z)|\rho_1 - \rho_2|.$$

■

PROPOSITION 4.1 (Self-Sufficiency Ceiling)

Fix any active belief $\rho \in \mathcal{A}$. Set

$$\tilde{\delta} = \frac{1}{1 + 4(2 + \frac{4\rho(\tilde{v}\alpha)^2}{(1-\ \rho)})}.$$

Then there exists $\bar{z} < 1$ such that for $\delta > \tilde{\delta}$, and for all $z \in (\bar{z}, 1)$, the unique optimal complexity choice at belief ρ is $x^* = 0$ or $s = 1$.

PROOF

Choose a compact interval $K = [\rho_a, \rho_b] \subset \mathcal{A}$. Fix $\rho \in K$. Consider the Bellman objective at belief ρ,

$$F(\rho, s) = \tau\tilde{v}p^*(\rho, s) + \delta[p^*(\rho, s)V\big(\rho^+(\rho, s)\big) + \big(1 - p^*(\rho, s)\big)V\big(\rho^-(\rho, s)\big)], \quad s \in (0,1]$$

For any $s \in (0,1)$, define the gain from simplifying the ritual as

$$\Delta(\rho, s) \equiv F(\rho, 1) - F(\rho, s).$$

I want to show that $\Delta(\rho, s) > 0$ for any $s \in (0,1)$.

[Step 1: current period flow gain is first order in $(1 - z)$]

By Corollary 4.1.1 (Cost Bound of Complexity), for any $\rho \in K$ and any $s \in (0,1)$,

$$p^*(\rho, 1) - p^*(\rho, s) \ge c_1(1-z)(1-s),$$

for some constant $c_1 = \frac{1}{\tilde{v}\alpha z}$ and $z \in (1/2,1)$.

Therefore, the flow difference satisfies

$$\tau v[p^*(\rho, 1) - p^*(\rho, s)] \ge \tau v c_1(1-z)(1-s).$$

This is strictly positive and first order in $(1-z)(1-s)$.

[Step 2: continuation value is second order in $(1 - z)$]

Write the continuation parts as

$$C(\rho, s) = p^*(\rho, s)V\big(\rho_+(\rho, s)\big) + \big(1 - p^*(\rho, s)\big)V\big(\rho_-(\rho, s)\big),$$

which can be decomposed as follows:

$$\begin{aligned}
&C(\rho, 1) - C(\rho, s) \\
&= \underbrace{[p^*(\rho, 1) - p^*(\rho, s)][V\big(\rho^+(\rho, 1)\big) - V\big(\rho^-(\rho, 1)\big)]}_{\text{(II)}} + \underbrace{p^*(\rho, s)[V\big(\rho^+(\rho, 1)\big) - V\big(\rho^-(\rho, s)\big)]}_{\text{(III)}} \\
&+ \underbrace{\big(1 - p^*(\rho, s)\big)[V\big(\rho^-(\rho, 1)\big) - V\big(\rho^-(\rho, s)\big)]}_{\text{(IV)}}.
\end{aligned}$$

Since $p^*(\rho, 1) - p^*(\rho, s) > 0$ by Lemma 4.1, V is increasing in ρ, and $\rho^+ > \rho^-$, Term (II) is non-negative, so it only helps the flow gain.

To bound (III) and (IV), first use Lemma 4.2 (Posterior Shift Bounds):

$$|\rho^+(\rho, 1) - \rho^+(\rho, s)| \le M^+(1 - z)(1 - s), \quad |\rho^-(\rho, 1) - \rho^-(\rho, s)| \le M^-(1 - z)(1 - s),$$

where $M^+ \le 2$, and $M^- \equiv \frac{4\rho(\tilde{v}\alpha)^2}{(1-\rho)}$.

Second, use Lemma 4.3 (Lipschitz Bound on V),

for any $\rho_1, \rho_2 \in K$

$$|V(\rho_1) - V(\rho_2)| \le C_K(1 - z)|\rho_1 - \rho_2|.$$

Combining,

$$|V\big(\rho^\pm(\rho, 1)\big) - V\big(\rho^\pm(\rho, s)\big)| \le C_K(1 - z) \cdot M^\pm(1 - z)(1 - s) = C_K M^\pm(1 - z)^2(1 - s).$$

Since $p^*(\rho, s) \le 1, 1 - p^*(\rho, s) \le 1$,

$$(|(\text{III})| + |(\text{IV})|) \le C_K(M^+ + M^-)(1 - z)^2(1 - s).$$

Therefore,

$$\delta|C(\rho, 1) - C(\rho, s)| \le \delta C_K(M^+ + M^-)(1 - z)^2(1 - s),$$

which is second order in $(1 - z)^2$.

[Step 3: First order beats second order for z close to 1]

Putting Steps 1 and 2 together and using (II) ≥ 0

$$\Delta(\rho, s) \ge \tau v c_1(1 - z)(1 - s) - \delta C_K(M^+ + M^-)(1 - z)^2(1 - s).$$

Factoring out $(1 - z)(1 - s) > 0$,

$$\Delta(\rho, s) \ge (1 - z)(1 - s)[\tau v c_1 - \delta C_K(M^+ + M^-)(1 - z)].$$

Define

$$\bar{z} \equiv 1 - \frac{\tau v c_1}{2\delta C_K (M^+ + M^-)}.$$

Recall all upper bounds:

$$c_1 = \frac{1}{\tilde{v}\alpha z} \leq \frac{2}{\tilde{v}\alpha},$$

$$C_K = \frac{2\ell_0}{1-\delta} = \frac{2\tau v c_1}{1-\delta} \leq \frac{4\tau \frac{v}{\tilde{v}}}{\alpha(1-\delta)},$$

$$M^+ = 2,$$

$$M^- = \frac{4\rho(\tilde{v}\alpha)^2}{(1-\ \rho)}.$$

Then,

$$\bar{z} \equiv 1 - \frac{\tau v c_1}{2\delta C_K (M^+ + M^-)} = 1 - \frac{1-\delta}{4\delta\left(2 + \frac{4\rho(\tilde{v}\alpha)^2}{(1-\ \rho)}\right)}.$$

Set

$$\tilde{\delta} = \frac{1}{1 + 4(2 + \frac{4\rho(\tilde{v}\alpha)^2}{(1-\ \rho)})}.$$

Then, $\tilde{\delta} \in (0,1)$, and for $\delta > \tilde{\delta}$, there exists $\bar{z} < 1$.

For all $z \in (\bar{z}, 1)$,

$$[\tau\tilde{v}c_1 - \delta C_K (M_+ + M_-)(1-z)] > \frac{\tau\tilde{v}c_1}{2} > 0.$$

Thus, for every $s \in (0,1)$ and $z \in (\bar{z}, 1)$,

$$\Delta(\rho, s) > 0.$$

That is, the Bellman objective is strictly higher at $s = 1$ than any $s < 1$.

Therefore, the unique optimal choice at active belief ρ is $x^* = 0$.

■

A.2 Proof for Proposition 4.2

The action space is extended to include $x = \infty$ (thereby $s = 0$), so the action set is a compact set $s \in [0,1]$. By Lemma 3.3.1, infinite complexity is never optimal in an active state, so $s = 0$ is ruled out ex post and the effective action space remains $s \in (0,1]$

LEMMA 4.2.A (Normalized Limit Operator)

Fix active belief $\rho \in \mathcal{A} \subset (0,1)$.

Define the normalized success rate as

$$a_0(\rho, s) \equiv \frac{p^*}{z}\Big|_{z\to 0} = \left(1 - \frac{1}{\tilde{v}\alpha\rho s}\right).$$

Define the limiting failure posterior as

$$R_0^-(\rho, s) \equiv \lim_{z\to 0} \rho^- = \frac{\rho(1-s) + \frac{1}{\tilde{v}\alpha}}{(1-\rho s) + \frac{1}{\tilde{v}\alpha}}.$$

Define the normalized value as

$$V^0(\rho) \equiv \lim_{z\to 0} \frac{V(\rho; z)}{z}.$$

, which satisfies the deterministic normalized Bellman equation,

$$V^0(\rho) = \max_{s\in[0,1]} \{\tau\tilde{v} a_0(\rho, s) + \delta V^0(R_0^-(\rho, s))\}.$$

PROOF

Divide the Bellman equation by z and send to the limit $z \to 0$ term by term.

Recall $\Phi(\rho, s) = z + \rho s(1 - z)$. As $z \to 0$,

$$\Phi(\rho, s) \to \rho s.$$

Since $A = \tilde{v}\alpha\Phi(\rho, s)$ and $p^* = 1 - \frac{1}{A}$, as $z \to 0$,

$$a_0(\rho, s) \equiv \frac{p^*}{z}\Big|_{z\to 0} = \left(1 - \frac{1}{\tilde{v}\alpha\rho s}\right).$$

Recall

$$\rho^- = \frac{\rho[1 - \pi^A(\rho = 1)]}{\rho[1 - \pi^A(\rho = 1)] + (1-\rho)(1-p^*)}.$$

It can be shown that

$$R_0^-(\rho, s) \equiv \lim_{z \to 0} \rho^- = \frac{\rho(1-s) + \frac{1}{\tilde{v}\alpha}}{(1-\rho s) + \frac{1}{\tilde{v}\alpha}}.$$

After dividing by z, the failure branch is, as $z \to 0$,

$$\delta(1-p^*)\frac{V(\rho^-; z)}{z} \to \delta V_0\big(R_0^-(\rho, s)\big),$$

since $1 - p^* \to 1$.

On the other hand, the success branch is, as $z \to 0$,

$$\delta \frac{p^*}{z} V(\rho^+; z) \to 0,$$

since $\frac{p^*}{z}$ goes to a finite $a_0(\rho, s)$ and $\lim_{z \to 0} V(\rho; z) \leq \lim_{z \to 0} \frac{\tau \tilde{v} z}{1-\delta} = 0$ (by Corollary 3.2.1)

Collecting terms,

$$V^0(\rho) = \max_{s \in [0,1]} \{\tau \tilde{v} a_0(\rho, s) + \delta V^0(R_0^-(\rho, s))\}.$$

■

LEMMA 4.2.B (Uniform Convergence of Normalized Values)

Let $K \subset \mathcal{A}$ be a compact set of active beliefs. For each $z \in (0, \underline{z}] \subset (0,1)$, define the normalized Bellman operator $T_z : C(K) \to C(K)$ by

$$(T_z W)(\rho) = \max_{s \in [0,1]} \left\{ \tau \tilde{v} \cdot \frac{p^*}{z} + \delta [p^* \, W(\rho^+) + (1 - p^*) \, W(\rho^-)] \right\},$$

where $\frac{p^*}{z} = 1 - \frac{1}{A}$, $A = \tilde{v}\alpha\Phi$, $\Phi = z + \rho s(1-z)$.

Define the limiting operator $T_0 : C(K) \to C(K)$ by

$$(T_0 W)(\rho) = \max_{s \in [0,1]} \{\tau \tilde{v} \cdot a_0(\rho, s) + \delta W(R_0^-(\rho, s))\},$$

where $a_0(\rho, s) \equiv \frac{p^*}{z}\big|_{z \to 0} = \left(1 - \frac{1}{\tilde{v}\alpha\rho s}\right)$, $R_0^-(\rho, s) \equiv \lim_{z \to 0} \rho^- = \frac{\rho(1-s) + \frac{1}{\tilde{v}\alpha}}{(1-\rho s) + \frac{1}{\tilde{v}\alpha}}$.

Then it follows that

(1) T_z and T_0 are a contraction mapping on $(C(K), \|\cdot\|_\infty)$ with the same modulus $\delta \in (0,1)$: for $W_1, W_2 \in C(K)$,

$$\|T_z W_1 - T_z W_2\|_\infty \leq \delta \|W_1 - W_2\|_\infty,$$

$$\|T_0 W_1 - T_0 W_2\|_\infty \le \delta \|W_1 - W_2\|_\infty.$$

(2) The operators converge uniformly as $z \to 0$, for $W \in C(K)$,

$$\|T_z W - T_0 W\|_\infty \to 0.$$

Consequently,

$$\left\| \frac{V(\cdot; z)}{z} - V_0 \right\|_\infty \le \frac{1}{1-\delta} \|T_z V_0 - T_0 V_0\|_\infty \underset{z \to 0}{\to} 0.$$

PROOF

[Step 1: T_z and T_0 are a contraction mapping with the same modulus $\delta \in (0, 1)$]

Take $W_1, W_2 \in C(K)$ and fix $\rho \in K$. For T_z, the only dependence on W is via $\delta[p^* W(\rho^+) + (1 - p^*) W(\rho^-)]$.

So for each s,

$|(T_z W_1)(\rho) - (T_z W_2)(\rho)| = |F_z(\rho, s; W_1) - F_z(\rho, s; W_2)| = \delta |p^* \big(W_1(\rho^+) - W_2(\rho^+)\big) + (1 - p^*) \big(W_1(\rho^-) - W_2(\rho^-)\big)| \le \delta [p^* |\big(W_1(\rho^+) - W_2(\rho^+)\big)| + (1 - p^*) |\big(W_1(\rho^-) - W_2(\rho^-)\big)|] \le \delta \|W_1 - W_2\|_\infty$.

Taking max over s cannot increase the gap:

$$|(T_z W_1)(\rho) - (T_z W_2)(\rho)| \le \delta \|W_1 - W_2\|_\infty.$$

Taking sup over $\rho \in K$:

$$|(T_z W_1) - (T_z W_2)| \le \delta \|W_1 - W_2\|_\infty.$$

Exactly the same argument applies to T_0:

$$|(T_0 W_1) - (T_0 W_2)| \le \delta \|W_1 - W_2\|_\infty.$$

[Step 2: Uniform Convergence $T_z W \to T_0 W$]

Fix $W \in C(K)$. For each (ρ, s), define

$$F_z(\rho, s) = \tau v \cdot \frac{p^*(\rho, s; z)}{z} + \delta [p^*(\rho, s; z)\, W\big(\rho^+(\rho, s; z)\big) + \big(1 - p^*(\rho, s; z)\big)\, W\big(\rho^-(\rho, s; z)\big)],$$

$$F_0(\rho, s) = \tau v\, a_0(\rho, s) + \delta\, W\big(R_0^-(\rho, s)\big).$$

Recall that as $z \to 0$,

$$\frac{p^*(\rho,s;z)}{z} \to a_0(\rho,s),$$

$$p^*(\rho,s;z) \to 0,$$

$$p^*W(\rho^+) \to 0,$$

$$\rho^-(\rho,s;z) \to R_0^-(\rho,s),$$

all of which are continuous in ρ, s, z on the compact set $K \times [0,1] \times [0,\bar{z}]$.

Since it is continuous on the compactness set, the convergence $F_z(\rho,s) \to F_0(\rho,s)$ is uniform on $K \times [0,1]$: as $z \to 0$,

$$\sup_{(\rho,s)\in K\times[0,1]} |F_z(\rho,s) - F_0(\rho,s)| \to 0.$$

Note that

$$(T_zW)(\rho) = \max_{s\in[0,1]} F_z(\rho,s)\,, \qquad (T_0W)(\rho) = \max_{s\in[0,1]} F_0(\rho,s).$$

For each ρ,

$$|(T_zW)(\rho) - (T_0W)(\rho)| = \left|\max_{s\in[0,1]} F_z(\rho,s) - \max_{s\in[0,1]} F_0(\rho,s)\right| \le \max_{s\in[0,1]} |F_z(\rho,s) - F_0(\rho,s)|$$
$$\le \sup_{s\in[0,1]} |F_z(\rho,s) - F_0(\rho,s)|.$$

Taking sup over $\rho \in K$,

$$\|T_zW - T_0W\|_\infty \le \sup_{(\rho,s)\in K\times[0,1]} |F_z(\rho,s) - F_0(\rho,s)| \to 0.$$

Therefore,

$$T_zW \to T_0W \text{ uniformly on } K.$$

[Step 3: Uniform Convergence of Fixed Points]

Let V_z and V_0 denote the fixed points where $V_z = T_zV_z, V_0 = T_0V_0$.

Then,

$$\|V_z - V_0\|_\infty = \|T_zV_z - T_0V_0\|_\infty = \|T_zV_z - T_zV_0 + T_zV_0 - T_0V_0\|_\infty$$
$$\le \|T_zV_z - T_zV_0\|_\infty + \|T_zV_0 - T_0V_0\|_\infty.$$

By the contraction property of T_z,

$$\|T_zV_z - T_zV_0\|_\infty \le \delta\|V_z - V_0\|_\infty.$$

It follows that

$$\|V_z - V_0\|_\infty \leq \|T_z V_z - T_z V_0\|_\infty + \|T_z V_0 - T_0 V_0\|_\infty \leq \delta \|V_z - V_0\|_\infty + \|T_z V_0 - T_0 V_0\|_\infty$$

Rearranging,

$$\|V_z - V_0\|_\infty \leq \frac{\|T_z V_0 - T_0 V_0\|_\infty}{(1-\delta)}.$$

Since $\|T_z V_0 - T_0 V_0\|_\infty \to 0$,

$$\|V_z - V_0\|_\infty \to 0.$$

■

LEMMA 4.2.C (Marginal Current Loss)

Fix active belief $\rho \in \mathcal{A} \subset (0,1)$. At $s = 1$ (or equivalently, $x = 0$), starting a ritual is locally costly, with marginal loss

$$MC(\rho) = -\frac{\partial}{\partial x}\left[\tau v \cdot a_0(\rho, e^{-\gamma x})\right]\Big|_{x=0} = \frac{\gamma\tau}{\alpha\rho(1-\tau)} > 0.$$

PROOF

Recall that

$$a_0(\rho, s) \equiv \frac{p^*}{z}\Big|_{z\to 0} = \left(1 - \frac{1}{\tilde{v}\alpha\rho s}\right).$$

Substitute $s = e^{-\gamma x}$,

$$a_0(\rho, s = e^{-\gamma x}) = \left(1 - \frac{e^{\gamma x}}{\tilde{v}\alpha\rho}\right).$$

The current payoff is

$$\tau v a_0 = \tau v \left(1 - \frac{e^{\gamma x}}{\tilde{v}\alpha\rho}\right).$$

Take derivative with respect to x

$$\frac{\partial}{\partial x_t}\left[\tau v \left(1 - \frac{e^{\gamma x}}{\tilde{v}\alpha\rho}\right)\right]\Big|_{x=0} = -\frac{\gamma\tau v}{\tilde{v}\alpha\rho} = -\frac{\gamma\tau}{\alpha\rho(1-\tau)} < 0.$$

Therefore,

$$MC(\rho) = -\frac{\partial}{\partial x_t}\left[\tau v \left(1 - \frac{e^{\gamma x}}{\tilde{v}\alpha\rho}\right)\right]\Big|_{x=0} = \frac{\gamma\tau}{\alpha\rho(1-\tau)}.$$

■

LEMMA 4.2.D.1 (Lower Bound Slope of Normalized Value)

Let $[a', b'] \in \mathcal{A}$ be the compact active belief region. Then V_0 is increasing on $[a', b']$ with uniform lower-bound slope:

$$V_0(r_2) - V_0(r_1) \geq \frac{\tau v}{\tilde{v}\alpha(b')^2}(r_2 - r_1),$$

where $a' \leq r_1 \leq r_2 \leq b'$.

PROOF

Fix a compact belief interval $[a', b'] \in \mathcal{A}$.

Let $r_1, r_2 \in [a', b']$ with $r_2 \geq r_1$.

Let $s_1 = s_1(r_1) \in \max_{s\in[0,1]}\{\tau v \cdot a_0(r_1, s) + \delta V_0\big(R_0^-(r_1, s)\big)\}$.

Recall that the normalized value function is defined as

$$V_0(r) = \max_{s\in[0,1]}\{\tau v \cdot a_0(r, s) + \delta V_0\big(R_0^-(r, s)\big)\}.$$

Because r_1 is active, $s = 0$ couldn't be optimal; otherwise current tribute would be zero and the state unchanged, contradicting optimality.

The Bellman equation gives

$$V_0(r_1) = \tau v\big(a_0(r_1, s_1)\big) + \delta V_0\big(R_0^-(r_1, s_1)\big).$$

The same action s_1 is also feasible at r_2.

So,

$$V_0(r_2) \geq \tau v \cdot a_0(r_2, s_1) + \delta V_0\big(R_0^-(r_2, s_1)\big).$$

Subtracting,

$$V_0(r_2) - V_0(r_1) \geq \tau v[a_0(r_2, s_1) - a_0(r_1, s_1)] + \delta[V_0\big(R_0^-(r_2, s_1)\big) - V_0\big(R_0^-(r_1, s_1)\big)].$$

Since $R_0^-(, s_1)$ is increasing in r, and V_0 is non-decreasing (established by Proposition 3.2.1 and Lemma 4.2.B)

$$V_0\big(R_0^-(r_2, s_1)\big) - V_0\big(R_0^-(r_1, s_1)\big) \geq 0.$$

It follows that

$$V_0(r_2) - V_0(r_1) \geq \tau v\big(a_0(r_2, s_1) - a_0(r_1, s_1)\big).$$

By Lemma 4.2.A,

$$a_0(r, s_1) = \left(1 - \frac{1}{\tilde{v}\alpha r s_1}\right).$$

Holding s_1 fixed, the mean-value theorem yields some $\check{r} \in (r_1, r_2)$ such that

$$\tau v\big(a_0(r_2, s_1) - a_0(r_1, s_1)\big) = \tau v \frac{1}{\tilde{v}\alpha s_1 (\check{r})^2}(r_2 - r_1).$$

Since $s_1 \leq 1$ and $\check{r} \leq b'$,

$$\tau v\big(a_0(r_2, s_1) - a_0(r_1, s_1)\big) \geq \frac{\tau v}{\tilde{v}\alpha (b')^2}(r_2 - r_1).$$

Therefore,

$$V_0(r_2) - V_0(r_1) \geq \frac{\tau v}{\tilde{v}\alpha (b')^2}(r_2 - r_1).$$

■

LEMMA 4.2.D.2 (Marginal Continuation Gain)

Fix an active belief $\rho \in (0,1)$ with $\tilde{v}\alpha > 1$. Define

$$b^\star \equiv R_0^-(\rho, 1) = \frac{1}{\tilde{v}\alpha(1-\rho) + 1}.$$

At $x = 0$,

$$MB(\rho) \equiv \delta \cdot \lim_{x \to 0^+} \frac{V_0\big(R_0^-(\rho, e^{-\gamma x})\big) - V_0(b^\star)}{x} \geq MC(\rho) \cdot L_0(\rho),$$

where $L_0(\rho) = \delta(\tilde{v}\alpha)^2 \rho^2 (1-\rho)$.

PROOF

For small $x > 0$, let

$$s \equiv e^{-\gamma x}.$$

From Lemma 4.2.A, define

$$b^\star \equiv R_0^-(\rho, 1) = \frac{1}{\tilde{v}\alpha(1-\rho) + 1}.$$

Since $R_0^-(\rho, s)$ is decreasing in s, and $s \equiv e^{-\gamma x}$ is decreasing in x,

For small $x > 0$,

$$R_0^-(\rho, s = e^{-\gamma x}) \geq b^\star \equiv R_0^-(\rho, 1).$$

Choose $\bar{b} > b^\star$ such that $[b^\star, \bar{b}] \subset \mathcal{A}$.

Set a small enough $x > 0$ such that

$$R_0^-(\rho, s = e^{-\gamma x}) \in [b^\star, \bar{b}].$$

Applying Lemma 4.2.D.1 on $[b^\star, \bar{b}]$:

$$V_0\big(R_0^-(\rho, s = e^{-\gamma x})\big) - V_0\big(R_0^-(\rho, 1)\big) \geq \frac{\tau v}{\tilde{v}\alpha(\bar{b})^2}\big(R_0^-(\rho, s = e^{-\gamma x}) - R_0^-(\rho, 1)\big)$$

Now divide by $x > 0$:

$$\frac{V_0\big(R_0^-(\rho, s = e^{-\gamma x})\big) - V_0\big(R_0^-(\rho, 1)\big)}{x} \geq \frac{\tau v}{\tilde{v}\alpha(\bar{b})^2}\frac{\big(R_0^-(\rho, s = e^{-\gamma x}) - R_0^-(\rho, 1)\big)}{x}.$$

With some algebraic manipulation,

it can be shown that

$$\begin{aligned}
&R_0^-(\rho, s = e^{-\gamma x}) - R_0^-(\rho, 1) \\
&= \frac{\rho(1-s) + \frac{1}{\tilde{v}\alpha}}{(1-\rho s) + \frac{1}{\tilde{v}\alpha}} - \frac{\frac{1}{\tilde{v}\alpha}}{(1-\rho) + \frac{1}{\tilde{v}\alpha}} \\
&= \frac{\rho(1-\rho)(1-s)}{\big(1 - \rho s + \frac{1}{\tilde{v}\alpha}\big)\big(1 - \rho + \frac{1}{\tilde{v}\alpha}\big)}.
\end{aligned}$$

Then,

$$\lim_{x \to 0^+} \frac{R_0^-(\rho, s = e^{-\gamma x}) - R_0^-(\rho, 1)}{x} = \frac{\rho(1-\rho)\frac{(1 - e^{-\gamma x})}{x}}{\big(1 - \rho e^{-\gamma x} + \frac{1}{\tilde{v}\alpha}\big)\big(1 - \rho + \frac{1}{\tilde{v}\alpha}\big)}.$$

Since $\lim_{x \to 0^+} \frac{(1 - e^{-\gamma x})}{x} = \gamma$,

$$\lim_{x \to 0^+} \frac{R_0^-(\rho, s = e^{-\gamma x}) - R_0^-(\rho, 1)}{x} = \frac{\gamma\rho(1-\rho)}{\big(1 - \rho + \frac{1}{\tilde{v}\alpha}\big)^2}.$$

Take the limit,

$$\lim_{x\to 0^+}\frac{V_0\big(R_0^-(\rho,s=e^{-\gamma x})\big)-V_0\big(R_0^-(\rho,1)\big)}{x}\geq\frac{\tau v}{\tilde{v}\alpha(\bar{b})^2}\frac{\gamma\rho(1-\rho)}{\Big(1-\rho+\frac{1}{\tilde{v}\alpha}\Big)^2}.$$

Since $b^\star=\frac{\frac{1}{\tilde{v}\alpha}}{(1-\rho)+\frac{1}{\tilde{v}\alpha}}$,

$$\frac{\gamma\rho(1-\rho)}{\Big(1-\rho+\frac{1}{\tilde{v}\alpha}\Big)^2}=\frac{\gamma\rho(1-\rho)}{\left(\frac{\frac{1}{\tilde{v}\alpha}}{b^\star}\right)^2}=\gamma\rho(1-\rho)(\tilde{v}\alpha)^2b^{\star 2}.$$

Therefore,

$$\lim_{x\to 0^+}\frac{V_0\big(R_0^-(\rho,s=e^{-\gamma x})\big)-V_0\big(R_0^-(\rho,1)\big)}{x}\geq\frac{\tau v}{\tilde{v}\alpha(\bar{b})^2}\gamma\rho(1-\rho)(\tilde{v}\alpha)^2b^{\star 2}.$$

Taking $\varepsilon\to 0$, then,

$$\bar{b}\to b^\star.$$

$$\lim_{x\to 0^+}\frac{V_0\big(R_0^-(\rho,s=e^{-\gamma x})\big)-V_0\big(R_0^-(\rho,1)\big)}{x}\geq\tau v\gamma\rho(1-\rho)\tilde{v}\alpha$$

Multiplying by $\delta\in(0,1)$,

$$\delta\lim_{x\to 0^+}\frac{V_0\big(R_0^-(\rho,s=e^{-\gamma x})\big)-V_0\big(R_0^-(\rho,1)\big)}{x}\geq\delta\tau v\gamma\rho(1-\rho)\tilde{v}\alpha.$$

Then,

$$MB(\rho)\geq\delta\tau v\gamma\rho(1-\rho)\tilde{v}\alpha.$$

Since $MC(\rho)=\frac{\gamma\tau}{\alpha\rho(1-\tau)}=\frac{\tau v\gamma}{\alpha\rho\tilde{v}}$, it can be rewritten as:

$$MB(\rho)\geq\delta\tau v\gamma\rho(1-\rho)\tilde{v}\alpha=\delta MC(\rho)\rho^2(1-\rho)(\tilde{v}\alpha)^2,$$

$$\frac{MB(\rho)}{MC(\rho)}\geq\delta\rho^2(1-\rho)(\tilde{v}\alpha)^2.$$

Setting $L_0(\rho)\equiv\delta\rho^2(1-\rho)(\tilde{v}\alpha)^2$,

$$\frac{MB(\rho)}{MC(\rho)}\geq L_0(\rho).$$

■

LEMMA 4.2.E.1 (Floor Condition at $x=0$)

At the limit $z = 0$, starting from $x = 0$, ritual is strictly beneficial at the margin, $MB(\rho) > MC(\rho)$, if and only if $L_0(\rho) = \delta(\tilde{v}\alpha)^2\rho^2(1-\rho) > 1$.

LEMMA 4.2.E.2 (Sufficient Patience Threshold)

Suppose the floor condition $L_0(\rho) = \delta(\tilde{v}\alpha)^2\rho^2(1-\rho) > 1$ is satisfied. For a fixed active belief $\rho \in \mathcal{A}$, the threshold is defined as

$$\overline{\delta}(\rho) = \frac{1}{(\tilde{v}\alpha)^2\rho^2(1-\rho)}.$$

For $\delta > \overline{\delta}(\rho)$, the optimal ritual complexity is strictly positive $x^* > 0$.

PROPOSITIN 4.2 (Technological Floor)

Fix active belief $\rho \in \mathcal{A}$. If the shaman is sufficiently patient, $\delta > \overline{\delta}(\rho) = \frac{1}{(\tilde{v}\alpha)^2\rho^2(1-\rho)}$, as in Lemma 4.2.E.2, then there exists a threshold $\underline{z} > 0$ such that for all $z \in (0, \underline{z})$, the optimal ritual complexity is strictly positive $x^* > 0$.

PROOF.

By Lemma 4.2.A,

the normalized value $V_0(\rho_t)$ satisfies:

$$V_0(\rho) = \max_{s\in[0,1]}\{\tau v a_0(\rho, s) + \delta V^0(R_0^-(\rho, s))\}.$$

By Lemma 4.2.C, at $z = 0, x_t = 0$,

$$MC(\rho) = \frac{\gamma\tau}{\alpha\rho(1-\tau)} > 0.$$

By Lemma 4.2.D.2,

$$MB(\rho) \geq MC(\rho) \cdot L_0(\rho),$$

where $L_0(\rho) = \delta(\tilde{v}\alpha)^2\rho^2(1-\rho)$.

For $\delta > \overline{\delta_1}$,

by Lemma 4.2.E.2, there exists ρ such that at $x = 0$,

$$L_0(\rho) > 1, \text{ and } MB(\rho) > MC(\rho).$$

By Lemma 4.2.B, as $z \to 0$,

$$\left\| \frac{V(\cdot; z)}{z} - V_0 \right\|_\infty \to 0,$$

, so there exists a small enough $\underline{z}$ such that for $z \in (0, \underline{z})$,

$$MB(\rho) > MC(\rho).$$

Therefore, for $z \in (0, \underline{z})$,

$$x_t^* > 0.$$

■

LEMMA (Per-period Payoff Is Concave in ρ_t)

Fix any active $\rho_t \in \mathcal{A}$ and a finite $x_t < \infty$. The current tribute revenue $\tau \tilde{v} p_t^*$ is strictly concave in ρ_t.

PROOF

Recall that

$$p_t^* = z - \frac{z}{\tilde{v}\alpha\Phi}.$$

Differentiating with ρ_t,

$$\frac{\partial \rho_t}{\partial p_t^*} = \frac{z}{\tilde{v}\alpha} \cdot \frac{1}{\Phi^2} \cdot \frac{\partial \Phi}{\partial \rho_t}.$$

Since $\frac{\partial \Phi}{\partial \rho_t} = e^{-\gamma x_t}(1-z)$,

$$\frac{\partial p_t^*}{\partial \rho_t} = \frac{z(1-z)e^{-\gamma x_t}}{\tilde{v}\alpha\Phi^2} > 0.$$

The second derivative with respect to ρ_t yields to

$$\frac{\partial^2 p_t^*}{\partial \rho_t^2} = -\frac{2z(1-z)^2 e^{-2\gamma x_t}}{\tilde{v}\alpha\Phi^3} < 0.$$

Therefore, the current tribute revenue $\tau \tilde{v} p_t^*$ is strictly concave in ρ_t.

■

A.3 Deriving Euler Equation

Euler equation is obtained under the maintained assumption that V is differentiable for the purpose of interpreting the economic forces governing x_t.

[Step 1] Differentiate the Bellman Equation with respect to x_t

Formally, the first-order condition with respect to x_t equates the marginal cost of complexity against its marginal benefit:

$$\underbrace{\tau v\left[\frac{\partial p_t^*}{\partial x_t}\right]}_{\text{lost current tribute }(-)} + \underbrace{\delta\frac{\partial p_t^*}{\partial x_t}[V(\rho_{t+1}^+)-V(\rho_{t+1}^-)]}_{\text{lower chance of good future state }(-)} + \underbrace{[\delta p_t^* V'(\rho_{t+1}^+)\frac{\partial \rho_{t+1}^+}{\partial x_t}}_{\text{diluted credit after success }(-)}$$

$$+ \underbrace{\delta(1-p_t^*)V'(\rho_{t+1}^-)\frac{\partial \rho_{t+1}^-}{\partial x_t}}_{\text{slower belief erosion after failure }(+)}] = 0.$$

The first term $\tau v\left[\frac{\partial p_t^*}{\partial x_t}\right]$ is the current revenue loss. A more demanding ritual suppresses current effort $\frac{\partial p_t^*}{\partial x_t} < 0$.

The second term $\delta\frac{\partial p_t^*}{\partial x_t}[V(\rho_{t+1}^+)-V(\rho_{t+1}^-)]$ is the success probability cost. Lower effort today reduces the chance of a good harvest $\frac{\partial p_t^*}{\partial x_t} < 0$, making the high-belief state ρ_{t+1}^+ less likely to be reached. Since $\rho_{t+1}^+ > \rho_{t+1}^-$, the shaman sacrifices expected continuation value by $V(\rho_{t+1}^+)-V(\rho_{t+1}^-) > 0$.

The third term $\delta p_t^* V'(\rho_{t+1}^+)\frac{\partial \rho_{t+1}^+}{\partial x_t}$ is the diluted affirmation cost. By Proposition 2.7.3, workers attribute less of the credit to the shaman $\frac{\partial \rho_{t+1}^+}{\partial x_t} < 0$. Success is less confirming of the shaman's divine connection.

The fourth term $\delta(1-p_t^*)V'(\rho_{t+1}^-)\frac{\partial \rho_{t+1}^-}{\partial x_t}$ is the only positive term. By Proposition 2.7.2, reputation declines more slowly against the shaman. This is the plausible deniability benefit that makes complexity worthwhile.

The shaman chooses x_t until the marginal value of slowing belief erosion after failure exactly covers the three costs: lost current tribute, a lower probability of reaching the high-belief state, and weaker confirmation of authority when success does occur.

[Step 2] Apply the Envelope theorem.

Differentiating the Bellman equation with respect to the state ρ_t, holding $x_t = x_t^*$, gives the Envelope condition:

$$
\begin{aligned}
&V'(\rho_t) \\
&= \underbrace{\tau v \frac{\partial p_t^*}{\partial \rho_t}}_{(i)>0} + \delta \underbrace{\frac{\partial p_t^*}{\partial \rho_t}[V(\rho_{t+1}^+) - V(\rho_{t+1}^-)]}_{(ii)>0} + \delta \underbrace{\left\{ p_t^* V'(\rho_{t+1}^+) \frac{\partial \rho_{t+1}^+}{\partial \rho_t} + (1 - p_t^*) V'(\rho_{t+1}^-) \frac{\partial \rho_{t+1}^-}{\partial \rho_t} \right\}}_{(iii)>0}
\end{aligned}
$$

Evaluate at $t+1$,

For successful state ρ_{t+1}^+,

$$
\begin{aligned}
&V'(\rho_{t+1}^+) \\
&= \tau v \frac{\partial (p_{t+1}^*)}{\partial \rho_{t+1}^+} + \delta \frac{\partial p_{t+1}^*}{\partial \rho_{t+1}^+}[V(\rho_{t+2}^+) - V(\rho_{t+2}^-)] + \delta p_{t+1}^* V'(\rho_{t+2}^+) \frac{\partial \rho_{t+2}^+}{\partial \rho_{t+1}^+} \\
&+ \delta (1 - p_{t+1}^*) V'(\rho_{t+2}^-) \frac{\partial \rho_{t+2}^-}{\partial \rho_{t+1}^+}.
\end{aligned}
$$

For unsuccessful state ρ_{t+1}^-,

$$
\begin{aligned}
&V'(\rho_{t+1}^-) \\
&= \tau v \frac{\partial (p_{t+1}^*)}{\partial \rho_{t+1}^-} + \delta \frac{\partial p_{t+1}^*}{\partial \rho_{t+1}^-}[V(\rho_{t+2}^+) - V(\rho_{t+2}^-)] + \delta p_{t+1}^* V'(\rho_{t+2}^+) \frac{\partial \rho_{t+2}^+}{\partial \rho_{t+1}^-} \\
&+ \delta (1 - p_{t+1}^*) V'(\rho_{t+2}^-) \frac{\partial \rho_{t+2}^-}{\partial \rho_{t+1}^-}.
\end{aligned}
$$

[Step 3] Iterate forward and substitute into the third and fourth terms of the FOC to get the Euler condition under the transversality condition $\lim_{t \to \infty} \delta^t V'(\rho_t) = 0$,

$$
\begin{aligned}
&- \left\{ \underbrace{\tau v \frac{\partial p_t^*}{\partial x_t}}_{\text{lost revenue } (-)} + \delta \underbrace{\frac{\partial p_t^*}{\partial x_t}[V(\rho_{t+1}^+) - V(\rho_{t+1}^-)]}_{\text{lower chance of the good state } (-)} \right\} \\
&= \underbrace{\sum_{k=1}^{\infty} \delta^k E_t \left[\left\{ \tau v \frac{\partial p_{t+k}^*}{\partial \rho_{t+k}} + \delta \frac{\partial p_{t+k}^*}{\partial \rho_{t+k}}[V(\rho_{t+k+1}^+) - V(\rho_{t+k+1}^-)] \right\} \cdot \frac{\partial \rho_{t+k}}{\partial x_t} \right]}_{\text{Present vaule of future net revenue gains from shifting beliefs (?)}},
\end{aligned}
$$

where $\frac{\partial p_{t+k}^*}{\partial \rho_{t+k}}$ is the marginal revenue of belief and $\frac{\partial \rho_{t+k}}{\partial x_t}$ captures how today's ritual complexity propagates through the belief process forward.

Remark 1 The optimal complexity policy x_t is an information design problem. The shaman controls how diagnostic each harvest outcome is about his own legitimacy. The Euler equation characterizes the optimal interior complexity policy, trading off the informativeness of publicly observable failure $\frac{\partial \rho_{t+s}}{\partial x_t}$ (and thus the preservation of future extraction gains $\frac{\partial p^*_{t+s}}{\partial \rho_{t+s}}$ through preserving future attachment and effort gains) against current extraction losses.