DRAFT: PLEASE DO NOT CIRCULATE OR CITE

The supernatural and the commons*

Christopher Costello, B. Kelsey Jack, and Tamatoa Bambridge

November 4, 2026

Abstract

We examine whether beliefs in supernatural punishment can coordinate resource extraction in the commons. Traditional small-scale religious belief systems around the world often include restrictions on the use of natural resources, such as fish, forests and particular species, that result in supernatural punishment if violated. Motivated by specific regularities documented in the anthropological literature, and by the *Rāhui* of Polynesian cultures in particular, we introduce beliefs about supernatural punishment into a standard commons model for which the tragedy of the commons is the no-regulation equilibrium. Here, players form beliefs about whether they will be punished, for example via sickness or misfortune, for violating extraction norms, even when their extraction is private information. We show that sufficiently strong beliefs can sustain fully cooperative behavior even in the complete absence of formal rules, monitoring, or punishment. We then introduce the possibility of learning about the supernatural, where agents can violate norms and learn (albeit noisily) about true supernatural punishment. Even in this setting, certain belief structures can sustain cooperation indefinitely. Comparative statics reveal the features of a belief system under which cooperation is sustained or unravels. We develop a new database of traditional belief systems around the world and show that documented belief systems are consistent with the model's mechanisms.

Costello: University of California, Santa Barbara and NBER (costello@bren.ucsb.edu). *Jack:* University of California, Berkeley and NBER (kelseyjack@berkeley.edu). *Bambridge:* CRIOBE, and Rāhui Center (tamatoa.bambridge@criobe.pf). We thank Patrick Rochette and the staff of the Rāhui Center for collaboration and guidance. Abigail Kirk, Mauricio Collado and Takurua Parent provided outstanding research assistance. We acknowledge support from Daniel Vapnek and the Environmental Markets Lab.

1 Introduction

On the island of Tahiti, prior to European arrival, a system of beliefs called Rāhui governed fishing activity. Priests would place markers called ti'i to announce temporary restrictions on harvest. Violators of these restrictions were subject to misfortune, illness and death, caused by the spirits inhabiting the ti'i (Teuira, 1928). Similar belief systems can be found around the world, both in the historical record and – in some places – persisting to the present day. By placing restrictions on the use of natural resources, they often contain echoes of modern natural resource management. For example, use of the ti'i in Tahiti resembles a temporary closure of the fishery. A place where extractive activities are prohibited resembles a protected area; for example, in Bhutan, farming or hunting in the mountain near Tshetayna was believed to lead to misfortunes and illness (Basnet and Dendup, 2020). A requirement of rituals or gifts tied to harvest resembles a tax; for example, in Serra da Jiboia, Brazil, a failure to make offerings to the deity Caipora prior to hunting or gathering wood was believed to result in accidents in the forest (Aguiar, Neto, and Fita, 2023). Unlike modern resource management policy, which requires credible monitoring and enforcement, these belief systems instead rely on beliefs in supernatural punishment to sustain cooperation.

Even in the complete absence of formal rules, monitoring or enforcement, beliefs in supernatural punishment may give rise to incentives that can support cooperative behavior in common pool resource settings. Many small societies around the world have sustainably managed their natural resources, such as fisheries, forests and water, without over-exploitation. Prior research has focused on the importance of social norms and repeated interactions for sustaining cooperation (Ostrom, 1990; Fehr and Gächter, 2000; Ostrom, 2000; Villeval, 2020). As an alternative or complementary explanation, we focus on the role of beliefs in supernatural monitoring and enforcement.

We develop a theoretical model of community natural resource use that incorporates individual beliefs about supernatural punishment for violation of local restrictions on natural resource use. The model is motivated by the specific case of Rāhui, a well documented belief system found throughout Oceania, which involves restrictions on behaviors, including fishing, that are believed to result in punishment through divine intervention when violated (Bambridge, 2016).[1] Compliance with the Rāhui depends on beliefs. If fishers believe that violation will result in punishment, they are more likely to comply. In the absence of learning, all fishers with strong enough beliefs will always comply.

However, just as fisheries are dynamic, so are beliefs. We therefore introduce learning

[1] While the strength and pervasiveness of beliefs in Rāhui have waned since European arrival, it remains culturally important and – in some places – strictly observed.

into the model. Learning occurs when a fisher (noisily) observes her punishment (or possibly the punishment of others), following violation of Rāhui. A key feature of Rāhui, which is common across the majority of traditional belief systems that we document, is that the punishment often resembles mundane bad events, for example "falling ill" or "misfortune". This makes it difficult to learn; if a neighbor is seen fishing in violation of the Rāhui and his wife falls ill, beliefs may be reinforced. If, on the other hand, the neighbor does not violate and his wife falls ill, beliefs will not be undermined, since community members fall ill for other reasons as well. We model this formally, without taking a stand on whether beliefs are right or wrong, and show the conditions under which Rāhui beliefs will be sustained, and the resource managed sustainably.

To provide a broader empirical context for the model, we construct a new database of 180 distinct beliefs – each representing a unique community, resource, and rule combination – from 32 countries and 5 continents. While not exhaustive, the cases we compile leverage the anthropology, sociology and ethnoecology literatures to extract details relevant to our model.[2] Specifically, we categorize and code beliefs in terms of resource type, rule governing extraction or use, and punishment for violation of the rule. Based on the details available in the written record, in all but 16 of the cases, learning is obscured by one or more of the following: the punishment is vague in kind (what will happen), timing (when it will happen), target (who will be affected) or likelihood (randomness in whether it will happen).[3] Thus, while the punishment rules are, in principle, falsifiable, inference by would-be violators is made difficult by a lack of specificity. This feature of punishment stands in contrast to the "big five" religions, which concentrate most punishment and rewards in the afterlife (Moreman, 2017), a design feature that makes learning inherently difficult.

The beliefs we document share an additional feature: each survived long enough to be recorded. Cross-cultural databases of religion and practice necessarily reflect communities and beliefs present at the time of study, and underrepresent short-lived phenomena that dissolved before reaching the written record. While this is a well-recognized limitation of such databases (Lowes, 2021), we follow Henrich (2016) in treating persistence as informative.[4] Specifically, cultural beliefs that survive may do so because they are resilient to defection or falsification. Our model directly addresses one part of this logic. Here, precise punishment

[2]We use the database both to motivate modeling assumptions and corroborate implications. Because we rely on historical cases that lack consistent detail on, for example, the duration of the beliefs, we are unable to test the model directly. We discuss these limitations in greater detail in Section 2.3 and potential approaches for hypothesis testing in the conclusion.

[3]Another 44 out of the 180 cases lack detail about punishment and are left uncoded.

[4]Analogous survivorship concerns arise in, for example, paleontology, where the fossil record over-represents organisms with durable hard parts and systematically under-represents soft-bodied organisms that decayed before fossilization (Signor and Lipps, 1982).

raises the informational value of violating the Rāhui. If violating reveals weak or zero punishment, compliance unravels. While we do not take a position on the magnitude of real-world punishment, suppose it is relatively low compared to the beliefs that sustain compliance. Then our model implies that documented belief systems should tend to feature vague punishments. Consistent with this logic, we find near-universal vagueness of punishments in our database.

Our model assumes that individuals are fully rational and update as Bayesian learners.[5] These agents interact in a complex dynamic game under uncertainty. Dynamics enter in two ways. First, consistent with other models of the commons, the resource stock itself is dynamic and evolves depending on the actions of all agents. This, alone, induces incentives where one agent's behavior imposes a dynamic externality on others, so that agents compete in a non-cooperative game. In our model, this effect gives rise to overextraction relative to the social optimum. But here, dynamic incentives are also induced by learning about supernatural punishment. Players form beliefs about supernatural punishment, but also recognize the opportunity to update those beliefs by violating the Rāhui norms and observing the outcome. Agents have to weigh the tradeoff between two competing actions. They could comply, thus obtaining lower expected current payoff, but running no risk of punishment, or they could violate, which delivers a larger current return and allows them to learn about supernatural punishment, but also reduces the future resource stock and triggers a potential punishment from the supernatural. To formalize and solve this model, we combine a canonical dynamic commons model (Levhari and Mirman, 1980; Noack and Costello, 2024) with the multi-armed bandit model (Berry and Fristedt, 1985), and make use of theoretical results on the Gittins Index (Gittins and Jones, 1979; Yao, 2006) and reservation prices (Weitzman, 1979). This combination allows us to derive sharp theoretical results about the equilibrium of play between these agents and how superstition can sustain, or unravel, cooperation in the commons.

The model gives rise to several notable results. First, regardless of whether agents can learn, in each period, each agent will either violate or comply with the Rāhui. Violation involves significant over-extraction effort and compliance involves just barely complying with the Rāhui. So while each agent's choice is continuous, only two alternatives are welfare-relevant; no agent will ever violate marginally. We show that if every agent's belief in supernatural punishment exceeds a common threshold, universal compliance is a Markov

[5]Bayesian updating extracts the maximum information available from a noisy signal, and is therefore often regarded as the efficiency benchmark. Non-Bayesian alternatives, such as confirmation bias or motivated reasoning, often lead to slower learning. We show that Rāhui can persist, even under Bayesian learning, so relaxing full Bayesian rationality would, if anything, reinforce our conclusion that superstitious beliefs can survive the possibility of learning.

perfect equilibrium. Because extraction efforts are strategic substitutes, it can be shown that this is also the unique stationary equilibrium. This equilibrium can nevertheless unravel if exogenous development changes model parameters such as the discount rate or the number of players.

We then introduce learning in the spirit of Weitzman (1979) and Gittins and Jones (1979). Agents form beliefs about supernatural punishment and can update those beliefs by violating the Rāhui and observing the punishment inflicted on them. The punishment is itself noisy, both because supernatural punishment is inherently noisy, and because punishment looks similar to bad things that happen anyway. In this setting, an agent's state contains both the resource stock itself and the agent's beliefs. We show that this dynamic programming equation is additively separable in two terms: One is a classic resource allocation problem and the other is a classic optimal learning problem.

The separability insight gives rise to our next result: Even if true punishment is zero, and agents can optimally learn, full cooperation in perpetuity is a possible equilibrium. This implies that learning does not necessarily unravel the ability of beliefs to maintain cooperation. Naturally, we derive the conditions under which this result obtains, and how it can unravel, for example with changes in economic growth (which could reduce background noise) or communication technology (where outside views could reduce the precision of the prior). We also show that because extraction efforts are strategic substitutes, a violation by i lowers the incentive for j to violate, which works against complete unraveling of Rāhui.

Recognizing that peer learning may also be possible, we then allow agents to observe others' violations and the shocks that follow. Peer learning introduces incentive margins that are distinct from those present under private learning. With peer learning, information becomes a public good, so agents may free-ride on others' experimentation. However, an agent whose violation reveals a mild punishment may encourage others to experiment, which generates free information for her but also induces higher extraction from the shared resource. Public signals also synchronize beliefs, so compliance may unravel abruptly if the true punishment is small, or become entrenched if it is large. The net effect of these forces depends on the equilibrium being played.

While most of our analysis assumes the Rāhui is designed for efficiency, we also examine a version of the model where the Rāhui designer (the "selfish chief") designs the Rāhui to maximize private rent capture. We derive an equilibrium in which the selfish chief constrains all agents to low extraction (but not so low that they violate), freeing up resources to be enjoyed by the chief. This equilibrium improves the chief's welfare, but decreases the welfare of all other agents; overall welfare and the resource stock itself are also compromised.

Our analysis focuses on beliefs in supernatural punishment as an important and un-

derstudied mechanism through which traditional institutions promote cooperation. This abstraction is analytically useful, but it should not be mistaken for a complete account of institutions such as the Rāhui. In Polynesian societies, the Rāhui is embedded in a metaphysical order in which society, nature, and the sacred are not treated as separate domains. Drawing on Amer Meziane's (2021; 2023) account of metaphysics as world-making, the Rāhui is not simply a rule sustained by shared beliefs about an external threat; it is part of an ontology that constitutes persons, environments, and spiritual beings as participants in a single relational order. Under this view, supernatural sanction is not an incentive imposed on an otherwise self-interested actor from outside, but an expression of a world in which moral obligation and ecological process are already inseparable. Our model captures one analytically tractable dimension of the Rāhui (its effect on extraction incentives) while leaving aside the cosmological institutions through which this relational world is itself continuously reproduced, and which likely account for at least part of the Rāhui's – and other similar systems' – historical resilience.

We leverage rich descriptions of traditional belief systems from large literatures outside of economics as inspiration for our model, and as a resource for our database of belief systems. Descriptive work has documented the prevalence of religious or spiritual beliefs surrounding natural resources. For example, Cox, Villamayor-Tomas, and Hartberg (2014) conduct a meta-analysis of papers on the role of religion in community based resource management. They describe religion as fulfilling several governance roles in natural resource management; our model formalizes its enforcement role. Their analysis also points to a positive association between religious governance and natural resource conditions, something built into our model largely by assumption. Following North's (1990) description of social taboos, Colding and Folke, 2001 describe resource and habitat taboos as informal institutions, and inventory a large number of belief systems by their functional roles in resource management (e.g., restrictions on harvesting methods). Anthropologists have long documented spiritual beliefs pertaining to the natural environment, dating back to some of the earliest ethnographies (Tylor, 1871).

The economics literature on community management of natural resources has emphasized the role of informal institutions for cooperative outcomes (Bardhan, 1993; Baland and Platteau, 1996; Sethi and Somanathan, 1996), but has largely overlooked belief systems as informal institutions.[6] In her influential book, Ostrom (1990) describes principles associated

[6]Informal institutions include social norms. Most economic models of social norms rely either on internalized norms (Elster, 1989) or the threat of costly social sanctions (e.g., Sethi and Somanathan, 1996) to maintain cooperation, and therefore involve no dynamic updating of beliefs about uncertain punishment. In Sethi and Somanathan's (1996) sanctioning model, for instance, the relevant uncertainty concerns whether the community will actually punish a violation, and the resulting signal (punished or not) is directly ob-

with successful community management of common pool resources. A belief in supernatural punishment is not among them. An empirical literature, dominated by lab experiments, has examined the effects of communication (e.g., Hackett, Schlager, and Walker, 1994; Ostrom, Gardner, and Walker, 1994), voluntary or costly sanctions or punishment (e.g., Ostrom, Walker, and Gardner, 1992; Casari and Plott, 2003), and self-governance (e.g., Walker et al., 2000; Abatayo and Lynham, 2016) on common pool outcomes.[7] While this literature has confirmed the potential for small groups to avoid the tragedy of the commons, it has been largely silent on the role of religious or supernatural beliefs in achieving this potential. A recent exception is Deopa and Rinaldo (2025), who compare forest conservation along plausibly exogenous ethnic borders and find that historically held beliefs in spiritual protection of forests translates into modern conservation outcomes. In their static model, greater adherence to the belief system is associated with lower utility from forest extraction. Our focus is, instead, on modeling the belief system itself and its potential to sustain cooperation without requiring that participants internalize norms into their utility functions.

A separate literature asks how religions and traditional beliefs enhance or undermine prosocial behavior (see Iannaccone (1998) and Iyer (2016) for reviews on the economics of religion and Le Rossignol, Lowes, and Nunn (2022) for recent evidence that some traditional beliefs may erode prosocial behavior), but has not considered how superstitious beliefs affect the management of common pool resources.[8,9] Most similar to our paper, Nunn and Sanchez de la Sierra (2017) describe a model of "bullet-proofing" in which young men choose to go to war (a prosocial choice), after undergoing a ritual that makes them able to withstand bullets. What inhibits learning in their case is hidden action: the ritual is complicated and if the man dies, it can be assumed that he performed the ritual incorrectly. We develop the dynamics implicit in their set up and introduce an alternative obstacle to learning: a noisy signal in the punishment function. In Espín-Sánchez, Gil-Guirado, and Ryan (2026), the

served, so there is no mechanism to impede learning as there is in our framework. In the absence of other frictions, uncertainty about the norm itself would resolve quickly. More recently, models of pluralistic ignorance allow biased second-order beliefs about the prevailing norm that update in response to exogenous information (e.g., Bursztyn, Egorov, and Fiorin, 2020). These models share our focus on norm erosion, but exogenous informational shocks about others' beliefs drive the updating process, whereas in our model, it is driven by endogenous experimentation by agents who choose to violate Rāhui.

[7] In some cases, findings have also been validated outside of the lab, both through correlational analyses (e.g., Fehr and Leibbrandt, 2011; Kosfeld and Rustagi, 2015) and causal identification (e.g., Haseeb, 2024; Garcia and Belmar, 2026).

[8] Other related papers, many outside of economics, on religious beliefs and prosocial behavior have, like our set up, emphasized the importance of gods or spirits for enforcement (e.g., Johnson and Krüger, 2004; Shariff and Norenzayan, 2011; Norenzayan, 2013; Purzycki et al., 2016). These papers have not, to our knowledge, formally addressed the challenge of sustaining beliefs in a dynamic setting where violation generates information.

[9] Buccione (2023) generates experimental evidence on the importance of religious messaging in the context of water conservation in Jordan, a common pool resource in times of scarcity.

persistence of religious beliefs depends not on noise in the information feedback, but on a context where the information will tend to confirm beliefs. They consider the widespread practice of praying for rain, and show that such beliefs can only be sustained in settings where the rainfall hazard function is increasing in time since the last rain; precisely the time when religious leaders will call for prayer to induce rainfall. In this case, religion leverages nature to sustain beliefs.

In focusing on the persistence of and learning about superstitious beliefs, our paper also connects to several theoretical papers from economics. First, our initial result that beliefs can be sustained, and therefore that cooperative extraction can persist based on a belief in an off-path punishment, is a classic example of a Self-Confirming Equilibrium (Fudenberg and Levine, 1993). In an equilibrium of this type, beliefs must be correct on the equilibrium path but can be superstitious about off-path events (i.e., they need not be objectively correct). In our model, the belief that violation induces punishment is never falsified; indeed, it sustains the cooperative outcome. The persistence of a wrong or unfalsified belief that sustains an equilibrium outcome is also central to Sargent (2008). That paper shows that adaptive learning converges to a self-confirming equilibrium where beliefs about events that are never observed can remain wrong indefinitely.[10] Our mechanism is closest to that in Fudenberg and Levine (2006), who study when superstitions about off-path consequences can persist even in the presence of rational learning. In our model, an agent's expectation of being punished for violating Rāhui is what makes testing the belief costly.

Second, stable compliance with Rāhui is related to an information cascade (Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992). In a cascade, agents rationally ignore their private information because others' actions are more informative. But in our model, others' compliance provides no information (to anyone, including agent i). Here, agents comply because acquiring information is itself costly; the only way for i to learn is to violate the Rāhui and endure the punishment. Viewed that way, stable compliance is an equilibrium with rationally incomplete learning arising from optimal experimentation (Rothschild, 1974; Easley and Kiefer, 1988). In our model, a rational agent may never experiment, or may experiment for a while and optimally stop, so could hold an incorrect belief forever. The Rāhui belief system is resilient because it deters the kind of learning that could undermine it. While much of the social learning literature explores conditions for converging to the truth (e.g., Acemoglu et al., 2011), our model provides a mechanism (a costly, noisy, endogenous signal) that allows rational Bayesian agents to remain "stuck" with possibly false, though welfare-improving, beliefs.

Third, our model provides a rational-learning, micro-foundation for taboo (Bénabou and

[10]This contrasts with rational expectations, where beliefs are right in expectation.

Tirole, 2011).[11] In our model, agents comply with Rāhui not due to irrationality or social norms, but because they are fully rational and internalize a dynamic tradeoff. The core decision is an exploration-exploitation problem, and we show that a confidently held belief in severe punishment, or a noisy signal, drives rational compliance, even in the presence of learning. Agents' beliefs function as an "asset," akin to the motivated beliefs in Bénabou and Tirole (2006), which helps the community overcome a social dilemma.

Fourth, we consider peer learning, which connects directly to the literature on strategic experimentation (Bolton and Harris, 1999; Keller, Rady, and Cripps, 2005). When agents can learn from violations by their peers, information about the Rāhui's true punishment becomes a public good, but violating to generate it is privately costly via the punishment an agent expects to endure. This gives rise to the free-rider and encouragement effects identified in that literature. In addition our setting adds a different margin that is absent from those models, which arises because learning takes place within a commons. Here, if i's violation reveals a mild punishment, it may induce others to violate, and this additional extraction reduces the resource stock, thus punishing i (among all others). This effect can deter i's violations. Public signals also synchronize beliefs, which can accelerate unraveling when the true punishment is small.

2 Background and motivating examples

Our model of supernatural beliefs in a common pool resource setting is motivated by a particular case: Rāhui in French Polynesia and Oceania. We describe Rāhui, drawing from the anthropological literature. We then describe new data we compile on traditional beliefs, with global coverage, which shows that many of the features of Rāhui are common to a large number of belief systems across a wide range of cultures, geographies and natural resources. These regularities across belief systems inform our approach to modeling beliefs.

2.1 Rāhui

In French Polynesia, and throughout Oceania, communities have long depended on oceans as a food source. When Europeans first arrived in the late 1700s, anthropologists described a robust belief system centered around Rāhui. A Rāhui is a specific form of tapu (in English, *taboo*), which is a restriction around a behavior, object, or relationship, violation of which carries the implied or explicit threat of supernatural punishment. Quoting from Rigo 2016: "Rāhui is a term, with some subtle phonetic nuances, that is found in all Polynesian islands;

[11] The English language word *taboo* is derived from the Polynesian word *tapu*. Rāhui is a form of tapu.

for example, in Samoa, Hawai'i, the Marquesas, Tahiti, the Tuamotu and the Mangareva islands. It refers to the prohibition or restriction applied to the consumption of a resource... Literally, it is the imposition of a tapu brought into effect by the sacred incantation (rahu) of an ari'i or a tahu'a (an expert in the relations between the world of men and the invisible entities). It is this same person who also lifts the Rāhui."

Rāhui are applied to both the land and the ocean. In the ocean, Rāhui resemble modern fisheries management policy: restrictions on the time of year, location or type of fishing activity. Violation of the Rāhui restrictions result in divine punishment: illness of a family member, crop failure or other bad events. We conjecture that fisheries-specific Rāhui were important for avoiding over-fishing, and that the specific details of how the belief system worked were important for allowing these beliefs to be sustained.

Since European arrival, sustainable management of local fisheries has declined (Andréfouët and Adjeroud, 2019). Numerous factors, including population growth and technology contribute to over-fishing; they may also contribute to the erosion of traditional beliefs. In the model we develop, we consider both the conditions under which beliefs about Rāhui will be sustained and also the shocks and external forces that will lead beliefs to unravel.

2.2 Database of traditional beliefs

We compiled a database of traditional beliefs related to natural resources, limiting ourselves to those involving some restriction on extraction or use of the resources. Existing databases of religion, beliefs or culture (see Lowes (2021) for a review of existing data sources) lack detail on, for example, specific prohibitions or punishments for violation. Instead, we turned to the anthropological literature, and extracted details from individual case studies. We relied primarily on searches conducted through online literature search tools, such as Google Scholar, using a set of standardized keywords such as *traditional knowledge*, *ethnoecology*, and *conservation*. Snowball sampling included both cited and citing sources, and a search for additional papers on specific belief systems referenced in the source, if the initial reference provided insufficient detail.

For each identified case, we recorded the country, region, community (ethnic group or society), resource, rule or prohibition, and punishment if the rule is violated. Each specific belief – defined as a unique combination of community, resource, and rule – is recorded as a separate observation. A single source may therefore contribute multiple observations if it documents beliefs across several communities or describes distinct rules for the same resource.

An excerpt from the database is included in Appendix A,[12] along with additional detail on the search methodology, data sources, coding of entries, and a comparison with existing alternative data sources.

2.3 Traditional beliefs and common pool resources

We document widespread evidence of superstitious beliefs surrounding the use of natural resources. In Asia, Oceania, Africa and the Americas, communities and ethnic groups have held or still hold beliefs that supernatural forces monitor and enforce the use of natural resources (see Figure 1).

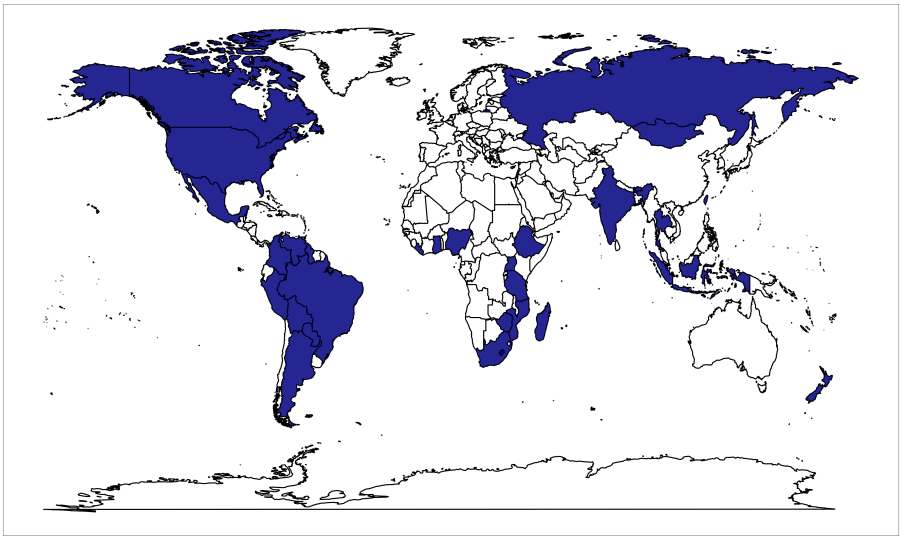


Figure 1: Geographic coverage of natural resource beliefs database

Our interest is in observing regular patterns across cases. We summarize several key features of the belief systems in our database (see Table 1). First, we categorize the natural resource that is the target of the belief into three main categories: habitat, wildlife and fish / game. Habitat is the most common (96 beliefs, 53% of observations), typically in reference to a protected area, sacred forest or otherwise restricted space. Fish and game (46 beliefs, 26%) and wildlife (38 beliefs, 21%) make up the remainder.

Second, we classify the rules according to what type of behavior is prohibited or prescribed. These rules can be assigned to three broad categories that map on to modern natural resource management: tax, restriction or prohibition. Prohibitions are the most common category (82 beliefs, 46%), followed by restrictions (65 beliefs, 36%) and tax-like prescriptions (33 beliefs, 18%). Within prohibitions, no-killing rules are most frequent (39

[12]The full database can be accessed here.

Table 1: Characteristics of natural resource beliefs

Category	Count	Description
Resource		
Wildlife	38	Animals that are sacred or taboo but not targeted for consumption.
Habitat	96	Forests, specified sites, ponds, trees, and other protected spaces.
Fish/Game	46	Plants and animals targeted for consumption.
Rule Type		
Tax		
Ritual	21	Sacred formalities required before or immediately after resource interaction.
Worship	7	Resource is treated as sacred; no specific ritual required.
Offering	5	Physical or spiritual gift required to gain use or entry approval.
Restriction		
Exploitation Restriction	31	Human activities limited by time, amount, type, or role in society.
Hunting Restriction	21	Hunting limited by time of year, amount, or role in society.
Entry Restriction	13	Entry to habitats limited by time, role in society, or other conditions.
Prohibition		
No Exploitation	27	No taking, harvesting, killing, or other human activities in a designated place.
No Killing	39	Sacred wildlife that cannot be killed.
No Entry	16	Entry to a sacred or forbidden site prohibited.
Punishment		
Illness/Death (Self or Other)	56	Physical illness or death to the defector, others, or both.
Loss of Divine Protection	18	Loss of protection from a supernatural enforcer against famine, disease, or other harm.
Loss of Crops/Livestock	12	Loss of the defector's and/or community's livestock or crops.
Loss of Hunt	12	Hunter will be unsuccessful, lose their way, or fail in future hunts.
Natural Disasters	10	Severe weather, drought, famine, or destruction of ecosystems.
Bad Luck	12	Unspecified misfortune listed as punishment in case study.
Psychological Harm	6	Family quarrels, mental illness, or other psychological effects.
Other	10	Punishments not fitting the categories above.

Note: Summary of 180 observations of traditional beliefs about natural resource use. Punishment counts exclude 44 observations with no recorded punishment information. See Appendix A for coding definitions.

beliefs); within restrictions, exploitation restrictions dominate (31 beliefs). Of the tax-like prescriptions, the most common is a requirement of ritual associated with resource use (21 beliefs).

Finally, we classify the punishments into several broad categories. Of the 136 beliefs with recorded punishment information, the most common punishment is illness or death to the violator or others (56 beliefs, 41%). Other common punishments include loss of divine protection (18), loss of crops or livestock, loss of hunt, and bad luck (12 each). Notably, none of the cases reviewed relegates punishment to the afterlife.

These punishment rules can be associated with specific impediments to learning, associated with a vagueness or lack of specificity in the rule or the punishment function. Our model makes use of these, so we describe them here, and show the resulting categorization in Figure 2.

1. Time: unclear *when* the punishment will occur after violation
2. Kind: unclear *what* the punishment will be (e.g., bad luck, illness, mental disorientation)
3. Who: unclear *to whom* the punishment will happen (e.g., friends and family, the community, unrelated people)
4. Random: the violation leads to a punishment *only probabilistically*
5. Clear: punishment rules that are specific / do not fall into any of the categories above

Our motivating example of Rāhui can be classified as follows. Though specific instances vary, it most commonly involves rules applied to either habitat or fish / game. Rules include both (temporary) limits on exploitation, when applied to specific species, and entry restrictions, when applied to fishing areas or forest areas. These fall into the broad category of restrictions. Punishments include bad luck, disease and death. These are typically not specific in *time*, *kind* and *who*.

The prevalence of vague punishment in our database – only 16 of 180 beliefs describe punishment rules specific enough not to impede learning – is consistent with a survivorship bias argument. Specifically, we observe only beliefs that have persisted long enough to be documented. Beliefs with precise, easily falsifiable punishment generate high-value opportunities to violate in order to learn. When such violations reveal weaker-than-expected punishment, the belief system can unravel quickly, and is less likely to survive long enough to be documented. In contrast, belief systems whose punishment resembles mundane events, like illness, misfortune, or crop failure, do not generate high-value learning opportunities,

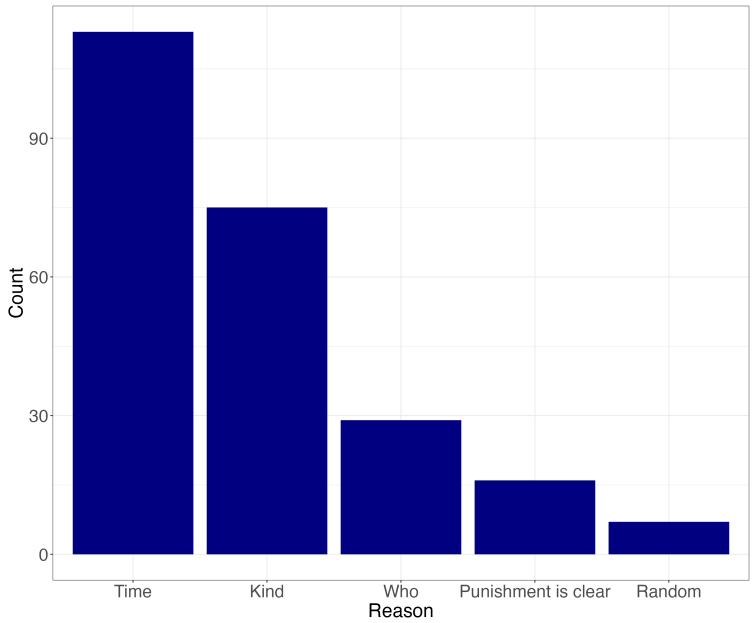


Figure 2: Sources of vagueness in punishment

Note: Each observation may exhibit more than one source of vagueness; categories are not mutually exclusive. 44 of 180 observations are excluded due to insufficient information on punishment.

because the signal from violating is noisy and weak. Our model formalizes this learning process, and we discuss in Section 6 the further step of establishing a link from learning to survival for documentation. The pattern in Figure 2 is consistent with these mechanisms, which we turn to now.

Our stylized model depicts a belief system governing a fishery, with restrictions on extraction (exploitation) and punishment that are conditioned on violations of these extraction restrictions. Like the large majority of cases we reviewed (64% exhibiting vagueness in timing, 42% in kind), we model punishment as difficult to distinguish from other background shocks – a design feature that, as we show, is key to sustaining beliefs.

3 A model of traditional beliefs in the commons

Motivated by these global examples, we develop a dynamic model of the commons where otherwise non-cooperative agents hold beliefs in supernatural punishment from violating locally determined norms, which we refer to as "Rāhui". We begin with a model without learning about supernatural punishment, and then allow private and eventually peer learning. Our focus is on the conditions that lead to compliance with Rāhui, and the conditions that lead to unraveling of beliefs in supernatural punishment, and therefore to unraveling of cooperation in the commons.

3.1 Model basics

N agents compete over the use of a common pool renewable resource like grazing a grassland, extracting wood from a forest, or harvesting fish from the ocean. See Table 1 for the range of natural resources covered by the cases in our database. For expositional clarity, we will refer to *extraction* from this common pool, but other forms of use could have similar interpretations. These agents are identical except in their beliefs about the supernatural world; we will return to this shortly. Absent supernatural intervention, agent i's period-t payoff is given by:

$$\pi_{it} = \log(h_{it}) \tag{1}$$

where h_{it} is i's extraction from the common pool in period t (h stands for "harvest"). We follow convention and assume that harvest is linearly homogeneous in the resource stock and extraction effort, so $h_{it} = E_{it}X_t$, where X_t is the resource stock competed for by all N agents and E_{it} can be thought of as the effort put toward extraction (fishing hours, or grazing intensity, for example). Aggregate extraction among all N extractors is

$$H_t = \sum_{i=1}^{N} h_{it}$$

The resource stock is renewable and its growth is endogenous and concave, as follows:

$$X_{t+1} = (X_t - H_t)^{\alpha} \tag{2}$$

where $\alpha \in (0, 1)$ and can be loosely interpreted as the "inverse growth rate" of the resource (because lower α implies faster growth for any given size of the resource stock). The implicit carrying capacity in this model is 1, so X_t should be interpreted as the size of the resource stock in period-t relative to its natural biological carrying capacity. Throughout, the stock satisfies $X_t > 0$, and each agent chooses $E_{it} > 0$. Because harvest is $E_{it}X_t$, aggregate effort must satisfy $\sum_j E_{jt} < 1$ (for the stock to remain positive). If aggregate effort reached or exceeded one, the resource would be exhausted and all future payoffs would be $-\infty$.[13]

So far, this model is similar to the classic fish war model from Levhari and Mirman (1980) and the more recent generalization by Noack and Costello (2024). To this basic setup we add *superstition*, that is, the *belief in supernatural punishment*. The way in which we model supernatural beliefs is inspired by Rāhui in French Polynesia, where a set of

[13]This follows from log payoffs. We will show that because an agent's continuation value contains $\log(1 - E_{it} - E_{-i})$, her best response to any $E_{-i} < 1$ is interior. All effort profiles we consider are interior. Aggregate effort is $1 - \alpha\delta$ under universal compliance, $1 - \alpha\delta(1 - E_{-i})$ under a unilateral violation, and E^v in the no-superstition benchmark (Proposition 1); all are less than one.

norms are explicitly stated to the N extractors (e.g., the rules in Table 1), along with a description of a supernatural punishment that will be inflicted on an agent should she violate the Rāhui (e.g., the punishment in Table 1). For example, on a given island there may be a Rāhui against fishing in the lagoon during spawning season, and the stated punishment is a sickness befalling a family member. Because these norms are typically stated as constraints on effort,[14] we will use effort, E_{it}, as agent i's control variable, and we assume the norm (i.e., the Rāhui) places a limit on the effort that i can exert; mathematically the Rāhui $\bar{E}$ requires that $E_{it} \leq \bar{E}$.

Agents enter a period with beliefs about the punishment from violating the Rāhui, and we examine how these beliefs affect equilibrium extraction, where our ultimate goal is to determine if, and under what conditions, these beliefs can sustain the fully cooperative extraction rate.

We assume agents behave non-cooperatively, and we seek a Markov Perfect Nash Equilibrium (MPNE). Each period, agent i chooses her own effort, E_{it}, which returns a deterministic payoff $\log(E_{it}X_t)$ and a stochastic additive shock to payoffs ϕ_{it}. This shock consists of a random component, or "luck", which is independent of i's actions and could be good or bad luck (denoted L_{it}), and a fixed supernatural punishment component that appears if and only if i violates the Rāhui (denoted $-Z_{it}$). We assume that L_{it} and Z_{it} are statistically independent.

Agents cannot distinguish between luck and punishment (that is, between L_{it} and Z_{it}), so the shock experienced by a agent violating Rāhui in period t is $\phi_{it} = L_{it} - Z_{it}$, and the shock for a compliant agent is just L_{it}.[15] The inability to precisely discern the true Rāhui punishment stems from the fact that both bad and good things can occur sporadically and independently of Rāhui compliance. For example, common forms of punishment, as summarized in Table 1, include illness or death, loss of crops or livestock and natural disasters, all of which also happen for reasons other than supernatural punishment for violating the Rāhui; that is, these could arise from bad luck (L_{it}) or punishment (Z_{it}).[16]

Superstition is represented by agents' beliefs over the punishment Z_{it}. To begin, we assume that agent i has beliefs given by the density $Z_{it} \sim f_i(Z)$, with mean μ_i in all t. In that setting, agents' beliefs do not evolve. Later in the paper we will allow agents to learn

[14] In our data (see Table 1), restrictions on exploitation and hunting are the most common type of rule, along with prohibitions on killing certain species or in specified locations.

[15] We assume punishment falls on the violating agent i. In practice, as documented in our database, punishment sometimes extends to the community or to uninvolved others. We leave this extension for future work.

[16] Note that punishment Z_{it} is a fixed bad event independent of the resource stock, rather than a consequence that scales with the state of the resource (e.g., fewer fish next season). This is true in the majority of cases in the database. There are, however, exceptions, such as beliefs held by the Sto:lo in British Columbia that a violation of rituals surrounding salmon catch would lead to poor future salmon runs.

so beliefs evolve over time. The distribution of luck is mean-zero with density $g(L)$. Thus, i's period-t payoffs are given by:

$$\pi_{it}(E_{it}|X_t) = \begin{cases} \log(E_{it}X_t) + L_{it} & \text{if} \quad E_{it} \leq \bar{E} \\ \log(E_{it}X_t) + L_{it} - Z_{it} & \text{if} \quad E_{it} > \bar{E} \end{cases} \tag{3}$$

where $L_{it} \sim g(L)$ and $Z_{it} \sim f_i(Z)$.

3.2 The no-superstition benchmark

As a benchmark, consider the case without superstition, where agents compete in a standard non-cooperative dynamic extraction game. Agent i's dynamic problem is:

$$V_i(X_t) = \max_{E_{it}} \mathbb{E}\left[\log(E_{it}X_t) + L_{it} + \delta V_i(X_{t+1}|E_{it}, E_{-it})\right] \tag{4}$$

where we note that stochastic luck contributes to payoffs, but without superstition, the punishment Z_{it} plays no role. It is straightforward to show that in the MPNE, i's value function has the form

$$V_i(X) = \frac{\log(X)}{1-\alpha\delta} + C$$

where C is a constant that is independent of X. This result substantially simplifies the equilibrium, and gives rise to the following result:

Proposition 1. *In the no-superstition game, let E_{-i} be the aggregate effort of all agents not including i. Then the optimal effort by i is given by:*

$$E_i^v = (1 - E_{-i})(1 - \alpha\delta) = \frac{1-\alpha\delta}{N - (N-1)\alpha\delta} \tag{5}$$

And aggregate effort is given by:

$$E^v = \frac{N(1-\alpha\delta)}{N(1-\alpha\delta) + \alpha\delta}$$

Proof. Substituting the form of the value function into the Bellman equation gives the FOC

$$\frac{1}{E_i^v} = \frac{\alpha\delta}{1-\alpha\delta}\left(\frac{1}{1 - E_{-i} - E_i^v}\right)$$

Rearranging delivers the first equality. Because agents are identical, we substitute $E_{-i} = (N-1)E_i^v$. Rearranging delivers the second equality. Multiplying by N and simplifying

gives the aggregate effort result. □

It is straightforward to show that aggregate effort is an increasing function of N. When the resource has a sole owner ($N = 1$), this equation returns the socially optimal level of effort $(1 - \alpha\delta)$. And when $N \to \infty$, this equation returns an aggregate effort of 1; this unfortunate outcome implies the resource is driven extinct in a single period.

3.3 Deriving the Rāhui

We showed above that in the absence of superstition and when $N > 1$, non-cooperative resource users will mine the resource more aggressively than is socially optimal. The purpose of this paper is to examine the role of beliefs in supernatural punishment in constraining that extraction. To make progress, we must specify the Rāhui itself. We define the Rāhui as a maximum level of effort that is allowed in each period by each fisher.[17] We assume that all N fishers know that they face supernatural punishment if and only if their effort exceeds the Rāhui, $\bar{E}$.

We assume that the Rāhui designer (typically an ari'i or tahu'a, the religious leader of the community) seeks to maximize extraction welfare, so sets $\bar{E}$ to exactly replicate the fully cooperative solution, assuming universal compliance.[18] Anthropological evidence suggests that the Rāhui is location and resource specific, so is endogenous to the parameters of the system. This suggests, for example, a faster growing resource, which can withstand greater extraction, might have a higher $\bar{E}$ than would a slower growing resource.

The Rāhui designer has a complicated dynamic optimization problem to solve, captured by the following infinite-horizon dynamic programming equation:

$$V(X_t) = \max_{E_t} \mathbb{E}\left[N \log(E_t X_t) + \sum_{i=1}^{N} L_{it} + \delta V\left((X_t - N E_t X_t)^{\alpha}\right)\right] \tag{6}$$

Here, the punishment Z_{it} does not appear in the dynamic problem because the Rāhui designer seeks the fully cooperative solution. The state variable is X_t and the control variable is E_t (the desired effort for each agent). The first term on the right-hand-side is the sum over all N agents of period-t utility. The second term is the stochastic luck (which is statistically independent across time and agents). The final term is the discounted continuation value, conditional on all N agents choosing the prescribed value of E_t, and accounting for the

[17] Most belief systems in our data impose a prohibition or restriction on hunting, entry, or exploitation of a resource (Table 1). These are best represented by a limit on extractive effort. A smaller set impose an implicit tax on effort by requiring, for example, a ritual that accompanies hunting. In these, the rule might be better represented by a cost tied to effort that, if not paid, results in punishment.

[18] In Section 5.2 we consider a leader who instead seeks to maximize his or her own private payoff.

growth of the stock. From this dynamic programming equation we can derive a closed-form expression for the optimal effort, summarized as follows:

Proposition 2. *The cooperative Rāhui model has solution*

$$\bar{E} = \frac{1 - \alpha\delta}{N} \tag{7}$$

which is independent of calendar time and independent of the stock X_t.

Proof. The derivation is similar to the previous result, so we omit it here. □

Notice that if all N agents comply with the Rāhui, aggregate effort is $N\bar{E} = 1 - \alpha\delta$, which is precisely the sole owner solution derived above. Proposition 2 implies that $\bar{E}$, is decreasing in N, δ, and α, so the designer sets tighter per-agent restrictions when the community is larger, the designer is more patient, and the resource is slower growing. An important implication of Proposition 2 is that the terms of the Rāhui, given by $\bar{E}$, are fixed over time, but will differ from location-to-location. The conclusion that $\bar{E}$ is fixed over time is, importantly, a result of our model, not an assumption.[19]

3.4 Superstition without learning

Suppose, then, that Rāhui $\bar{E}$ is announced. Any agent who supplies $E_{jt} \leq \bar{E}$ knows that she is free from supernatural punishment in period t, and any agent who supplies $E_{it} > \bar{E}$ knows she will receive punishment Z_{it} in period t. Such a violating agent also knows that Z_{it} is only partially observable because it cannot be distinguished from L_{it}. To begin, we assume no learning over time, so agent i's beliefs about Z_{it} are completely characterized by the stationary distribution $f_i(Z)$, with $\mathbb{E}(Z_{it}) = \mu_i$. Each agent brings her own beliefs about this punishment (which of course, could be 0) and we seek the MPNE across agents.

In particular, we are interested in whether these beliefs are sufficient to sustain universal compliance with the Rāhui by all N agents as an equilibrium outcome. We therefore focus on the following equilibrium:

Definition 1. *A* universal compliance equilibrium *is an MPNE in which every agent complies with the Rāhui in every period and at every resource stock; that is, $E_{it} \leq \bar{E}$ for all i and t.*

To determine when such an equilibrium exists, it suffices to check whether any single agent would deviate when all others comply. Suppose, then, that $N - 1$ agents comply with the

[19] This is difficult to verify in our database since cases are often recorded only at a single point in time – however, we note that anecdotally many Rāhui in French Polynesia were unchanged for centuries.

Rāhui. While those cooperating agents could choose any $E_j \leq \bar{E}$, it is straightforward to see that they will all choose $E_j = \bar{E}$. This is because any choice $E_j \leq \bar{E}$ is free from punishment, but present value profit is increasing in E_j.[20] So from the perspective of the potentially violating agent i, aggregate effort by all others is:

$$E_{-i} = \frac{N-1}{N}(1-\alpha\delta)$$

Given this, what is i's optimal choice of effort? If i decides to comply, she will choose $E_i = \bar{E}$ (by the argument above). Instead, invoking Proposition 1, if i decides to violate, she will choose:

$$\hat{E} = (1-E_{-i})(1-\alpha\delta) \tag{8}$$

$$= \frac{1+\alpha\delta(N-1)}{N}(1-\alpha\delta) \tag{9}$$

Importantly, the punishment distribution, $f_i(Z)$, will play a pivotal role in determining whether i complies or violates, but plays no role in determining the optimal violation effort, $\hat{E}$. This result arises because, by assumption, punishment induces no marginal incentive for effort. Throughout the paper we assume that the punishment Z_{it} is triggered by any violation $E_{it} > \bar{E}$, independent of the magnitude of violation. This is consistent with the database of beliefs: of the 136 beliefs with relatively well documented punishment, none clearly describes punishment that increases in the degree of violation. It is easy to show that i's optimal violation effort when all other agents comply is larger than i's violation effort when all other agents violate. This result arises because i's optimal violation effort is decreasing in the aggregate effort of all others (Equation 5), and the aggregate effort of the others is larger when they violate than when they comply.

Agent i's dynamic programming equation is:

$$V_i(X_t) = \max_{E_{it}} \mathbb{E}\left[\log(E_{it}X_t) + L_{it} - Z_{it}\,\mathbb{I}_{E_{it}>\bar{E}} + \delta V_i(X_{t+1}|E_{it}, E_{-it})\right] \tag{10}$$

where $\mathbb{I}_{E_{it}>\bar{E}}$ (henceforth, $\mathbb{I}$ for short) is an indicator that equals 1 if i violates the Rāhui and 0 if i complies. The key thing to notice is that the punishment for violating depends on E_{it} but is independent of X_t. The implication is that the value function has the same form as above:

$$V_i(X_t) = \frac{\log(X_t)}{1-\alpha\delta} + C \tag{11}$$

where the constant C is independent of X_t but depends on i's effort choice and on the

[20] Agent j's payoff is increasing in E_j up to (and in fact beyond) $\bar{E}$.

expected punishment from violating, μ_i.

Because the constant C is independent of X_t, the agent's choice of E_{it} is independent of X_t. Together with the violation and compliance levels of effort derived above, this gives rise to an explicit expression for i's value function, as follows:

$$V_i(X_t) = \begin{cases} \frac{\log(X_t)}{1-\alpha\delta} + C_0 & \text{if} \quad E_{it} = \bar{E} \quad \forall t \\ \frac{\log(X_t)}{1-\alpha\delta} + C_1 & \text{if} \quad E_{it} = \hat{E} \quad \forall t \end{cases} \tag{12}$$

Where we derive the infinite horizon constants in the MPNE as follows:

$$C_0 = \frac{\log(\bar{E})}{(1-\delta)} + \frac{\alpha\delta\log(1-\bar{E}-E_{-i})}{(1-\alpha\delta)(1-\delta)} \tag{13}$$

$$C_1 = \frac{\log(\hat{E}) - \mu_i}{(1-\delta)} + \frac{\alpha\delta\log(1-\hat{E}-E_{-i})}{(1-\alpha\delta)(1-\delta)} \tag{14}$$

Effectively, agent i's choice of whether to comply with the Rāhui (and earn present value $\frac{\log(X_t)}{1-\alpha\delta} + C_0$) or to violate from the Rāhui (and earn present value $\frac{\log(X_t)}{1-\alpha\delta} + C_1$) boils down to $C_0 \lessgtr C_1$. This inequality can be manipulated to derive the following result:

Proposition 3. *In the infinite horizon common pool extraction game with superstition (but no learning), universal compliance (Definition 1) is a MPNE if and only if $\mu_i \geq \mu^c$ for all i, where*

$$\mu^c = \log(1+\alpha\delta(N-1)) + \frac{\alpha\delta}{(1-\alpha\delta)}\log\left(1 - \frac{N-1}{N}(1-\alpha\delta)\right) \tag{15}$$

In this equilibrium, every agent chooses $E_{it} = \bar{E}$. If $\mu_i > \mu^c$ for all i, compliance is each agent's strict best response.

Proof. Consider arbitrary agent i and suppose every $j \neq i$ chooses $\bar{E}$ at every state, so $E_{-i} = (N-1)\bar{E}$ is constant. Then i's value function takes the form $\frac{\log(X_t)}{1-\alpha\delta} + C$, and i's period-t effort affects her value only through the period-t term $\log(E_{it}) + \frac{\alpha\delta}{1-\alpha\delta}\log(1-E_{it}-E_{-i}) - \mu_i \mathbb{I}_{E_{it}>\bar{E}}$. Agent i's problem is therefore the same static comparison in every period and at every stock. Conditional on compliance, this term is maximized at $\bar{E}$, and conditional on violation, it is maximized at $\hat{E}$, so i's best response is one of the two stationary policies with values C_0 and C_1. Noting that C_0 is independent of μ_i and C_1 is decreasing in μ_i, compliance is a best response if and only if $C_0 \geq C_1$. Setting $C_0 = C_1$ delivers the critical value:

$$\mu^c = \log\left(\frac{\hat{E}}{\bar{E}}\right) + \frac{\alpha\delta}{(1-\alpha\delta)}\log\left(\frac{1-\hat{E}-E_{-i}}{1-\bar{E}-E_{-i}}\right) \tag{16}$$

$$= \log(1+\alpha\delta(N-1)) + \frac{\alpha\delta}{(1-\alpha\delta)}\log\left(1 - \frac{N-1}{N}(1-\alpha\delta)\right) \tag{17}$$

so compliance is a best response for i if and only if $\mu_i \geq \mu^c$, and the strict best response if $\mu_i > \mu^c$. Universal compliance is a MPNE if and only if this holds for every i. □

The critical value of mean beliefs, μ^c, is positive for $N > 1$. When $N = 1$ we obtain $\mu^c = 0$ because a sole owner requires no supernatural beliefs to extract at the socially optimal level. By Proposition 3, μ^c is the belief at which a single agent is just barely willing to comply when all others comply. If every agent's belief exceeds μ^c, universal compliance with the Rāhui is an equilibrium. If some agent holds belief $\mu_i < \mu^c$, universal compliance is not an equilibrium, because that agent would optimally deviate to $\hat{E}$. In that case, the resulting equilibrium will depend on how the remaining agents respond to her violation, which we address next.

Proposition 3 assumes all other agents comply. Suppose instead that the other agents' aggregate effort is constant at arbitrary effort E_{-i}, leaving a maximum possible effort by i given by $R \equiv 1 - E_{-i}$, where R stands for "residual". By Proposition 1, i's optimal violation is $\hat{E}(R) = (1 - \alpha\delta)R$, i complies if and only if $\mu_i \geq \mu^c(R)$, where

$$\mu^c(R) = \log\left(\frac{(1-\alpha\delta)R}{\bar{E}}\right) + \frac{\alpha\delta}{1-\alpha\delta}\log\left(\frac{\alpha\delta R}{R-\bar{E}}\right)$$

Under universal compliance, $R = \bar{R} \equiv \frac{1+\alpha\delta(N-1)}{N}$ and $\mu^c(\bar{R}) = \mu^c$. Differentiating,

$$\frac{d\mu^c(R)}{dR} = \frac{\hat{E}(R) - \bar{E}}{(1-\alpha\delta)R(R-\bar{E})} > 0 \quad \text{whenever} \quad \hat{E}(R) > \bar{E}.$$

This implies that compliance is easier to achieve when others extract more. The intuition makes use of the fact that efforts are strategic substitutes, so violations by others make compliance easier to sustain for the remaining agents.

We have shown that if every agent's belief in supernatural punishment exceeds μ^c, universal compliance with the welfare-maximizing effort level is a MPNE, even in the absence of any kind of formal institutions governing extraction.

4 Learning about supernatural punishment

So far we have assumed that while each agent may hold different beliefs about the supernatural, each agent's beliefs are fixed over time. While this helps with model tractability, it obviates learning and so suppresses a potentially important feature of superstitious beliefs. In this section, we incorporate learning into the model presented above and ask whether allowing learning could cause superstitious beliefs to unravel. We begin with learning only by

doing (which we refer to as *private learning*) and later add *peer learning*. As in Section 3.4, our focus is on universal compliance. We ask when universal compliance can be sustained as an equilibrium outcome once agents can learn by violating.[21]

Under private learning, an agent's beliefs change only when she herself violates. Thus if all agents other than i comply, any deviation by i, including a sequence of experimental violations, leaves the other agents' information states unchanged. So, to establish a universal compliance equilibrium, we need only examine agent i's learning problem holding the others' aggregate effort fixed at $E_{-i} = (N-1)\bar{E}$. We maintain this structure throughout Sections 4.1–4.3, and return to settings in which others' behavior responds to violations when we consider peer learning in Section 4.4.

Agent i can learn about punishment (Z_i) by violating the Rāhui and (noisily) observing Z_i. When agent i complies with the Rāhui, she knows for sure that she is free from supernatural punishment, so cannot learn anything new about the distribution of Z_i. Instead, when i violates the Rāhui, she knows that the shock she incurs is part random and part punishment; this allows her to update her beliefs about $f_i(Z)$.

To operationalize learning about the supernatural, we adopt a Bayesian approach. Following from above, let the true (but unknown) mean supernatural punishment for i be μ_i, which is fixed for all time. We assume that in any period t, a violating agent's punishment is drawn from a normal distribution with mean μ_i:

$$Z_{it} \sim N(\mu_i, \tau_i)$$

where the precision[22] (τ_i) is known to agent i, but the mean (μ_i) is not.[23]

Agent i can learn about μ_i by violating the Rāhui. We assume that i enters period t with a normal prior distribution over μ_i, as follows:

$$\mu_i \sim N(m_{it}, \eta_{it}) \quad \text{in period } t$$

where m_{it} captures i's period-t belief about the mean μ_i, and η_{it} reflects the precision of that belief.

[21] Learning by violating is particularly relevant in the cases recorded in our database, none of which emphasize punishment in the afterlife. Thus, unlike other major belief systems, these traditional beliefs are – in principle – falsifiable.

[22] Inverse of the variance.

[23] We assume τ_i (and, in what follows, the precision of luck ω_i) is known to agent i while the mean, μ_i, is not. This is a common assumption for conjugate prior models and helps with tractability. Here, agents learn about a single unknown parameter, so that beliefs are summarized by the two-dimensional state (m_{it}, η_{it}) and the learning problem is a normal bandit with known variance. If the precision were also unknown, beliefs would follow a richer posterior with additional state variables.

Learning requires a bit more work because Z_{it} cannot be directly observed, as in the standard Bayesian updating model. When i violates the Rāhui in period t, the shock she experiences is $\phi_{it} = L_{it} - Z_{it}$. Importantly, this shock ϕ_{it} is observed, but i cannot distinguish the "luck" component (L_{it}) from the "punishment" component ($-Z_{it}$). Punishment is non-specific in a variety of ways and Figure 2 summarizes features of punishment observed in our data that make it hard to distinguish punishment from luck.[24]

We assume that luck has distribution $L_{it} \sim N(0, \omega_i)$, where ω_i is the precision of the luck distribution for i, which allows for the possibility that one agent's luck distribution is more variable than another's. Thus, a Rāhui-violating agent i's overall shock is the sum of two normally distributed random variables and has the following distribution:

$$\phi_{it} \sim N\left(-\mu_i, \nu_i\right)$$

where ν_i is the precision of ϕ_{it}.[25] In other words, a violating agent receives a shock with unknown mean $-\mu_i$ and known precision ν_i.[26] That agent has a prior distribution over μ_i, and knows that if she violates the Rāhui, she will observe that period's shock ϕ_{it} which allows her to update her prior. It turns out that because both the prior over μ_i and signals ϕ_{it} are normally distributed with known precision, we can invoke the conjugacy of the prior to show that the posterior distribution over μ_i is also normally distributed. This posterior then becomes i's prior distribution over μ_i in $t+1$. Putting all this together, agent i's updated beliefs are summarized as follows:

$$\mu_i \sim N(m_{it+1}, \eta_{it+1}) \quad \text{in period } t+1, \text{ where} \tag{18}$$

$$m_{it+1} = m_{it}\left(\frac{\eta_{it}}{\eta_{it}+\nu_i}\right) - \phi_{it}\left(\frac{\nu_i}{\eta_{it}+\nu_i}\right) \quad \text{and} \tag{19}$$

$$\eta_{it+1} = \eta_{it} + \nu_i \tag{20}$$

A violating agent receives a shock ϕ_{it} and updates her beliefs about μ_i. The mean of this posterior is a precision-weighted average of the prior mean (m_{it}) and the negative of the observed shock ($-\phi_{it}$). Higher precision on the prior (η_{it}) lends more weight to the prior

[24] If, for example, the punishment stated that a violator's first born child would develop a rash on their left leg within 2 hours of the violation, then the punishment could be easily isolated even if kids get rashes frequently.

[25] Because the sum of normal distributions is also normal, it is straightforward to show that

$$\nu_i = \frac{1}{\frac{1}{\tau_i} + \frac{1}{\omega_i}}$$

[26] For a Rāhui-violator, it is helpful to think of μ_i as a positive number, so that ϕ_{it} is negative in expectation. Because luck is also present, any particular realization of ϕ_{it} may nonetheless be positive.

mean, while higher precision on the shocks themselves (ν_i) lends more weight to the observed shock. Thus, a worse outcome after violating (a more negative ϕ_{it}) raises the agent's belief about the mean punishment, while a mild or favorable outcome lowers it. The negative sign in Equation 19 arises because ϕ_{it} has mean $-\mu_i$. The updated precision is the sum of the prior's precision (η_{it}) and the precision of the overall shock (ν_i). The payoffs, dynamics, and learning are summarized in Figure 3.

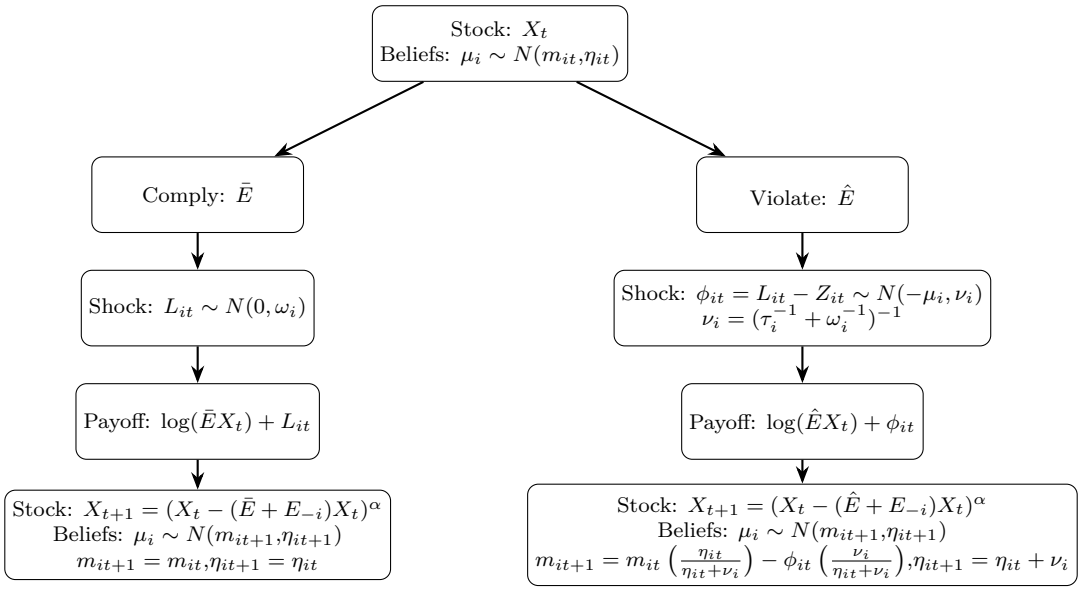


Figure 3: Decision tree to comply vs. violate the Rāhui

4.1 Learning as a dynamic program

We now fold Bayesian learning about supernatural punishment into the dynamic game presented above. The period-t state for agent i is now a vector composed of the resource stock (X_t), the mean of the prior (m_{it}), and the precision of the prior (η_{it}). Agent i's dynamic programming equation is:

$$V_i(X_t, m_{it}, \eta_{it}) = \max_{E_{it}} \mathbb{E}\left[\log(E_{it}X_t) + L_{it} - Z_{it}\,\mathbb{I}_{E_{it}>\bar{E}} + \delta V_i(X_{t+1}, m_{it+1}, \eta_{it+1} | E_{it}, E_{-i})\right] \tag{21}$$

where $E_{-i} = (N-1)\bar{E}$ is the (constant) aggregate effort of the other, Rāhui-compliant agents, and the equations of motion are given by Equations 2, 19, and 20. Our first observation is that even in this learning setting, agent i will choose one of two effort levels in period t: $\bar{E}$ (compliance with the Rāhui) or $\hat{E}$ (a given level of effort that is optimal if she decides to violate the Rāhui). This is straightforward to confirm because, as we established without learning, if i complies, there is no marginal incentive to *over-comply*. If i violates, she knows

she will receive a punishment (and therefore learn) regardless of the extent of her violation, so learning does not affect the optimal extent of violation. Thus, given that the other agents comply, both $\bar{E}$ and $\hat{E}$ are independent of the state (X_t, m_{it}, η_{it}), and are given in Equations 7 and 9.

Intuitively, though, the period-t state does affect the incentives for i to select $\bar{E}$ versus $\hat{E}$. To explore these incentives, we begin with the following observation: while the size of the resource stock (X_t) surely affects payoffs, it turns out to affect payoffs equally whether the agent complies or violates. For that reason, the dynamic programming problem can be separated into a deterministic component that depends only on X_t, and a stochastic component that depends on E_{it} and the information state (m_{it} and η_{it}), summarized by the following proposition:

Proposition 4. *Suppose all agents $j \neq i$ comply in every period. Then agent i's dynamic programming equation can be written in an additively separable manner, as follows:*

$$V_i(X_t, m_{it}, \eta_{it}) = \frac{\log(X_t)}{1 - \alpha\delta} + J_i(m_{it}, \eta_{it}) \tag{22}$$

where the new value function $J_i(m_{it}, \eta_{it})$ is a time-invariant function of the information state, is independent of X_t, and depends on model parameters and on the optimal choice of E_{it}.

Proof. We proceed by induction. First, guess that the value function for all future periods takes the separable form:

$$V_i(X_{t+1}, m_{it+1}, \eta_{it+1}) = \frac{\log(X_{t+1})}{1 - \alpha\delta} + J_i(m_{it+1}, \eta_{it+1})$$

Now substitute this assumed form into the dynamic programming equation, giving:

$$V_i(X_t, m_{it}, \eta_{it}) = \max_{E_{it}} \mathbb{E}\left[\log(E_{it}X_t) + L_{it} - Z_{it}\,\mathbb{I}_{E_{it}>\bar{E}} + \delta\left(\frac{\log(X_{t+1})}{1 - \alpha\delta} + J_i(m_{it+1}, \eta_{it+1})\right)\right]$$

The stock-dependent terms can be expanded and substituted into the dynamic programming equation, giving:

$$\max_{E_{it}} \mathbb{E}\Bigg[\log(E_{it}) + \log(X_t) + L_{it} - Z_{it}\,\mathbb{I}_{E_{it}>\bar{E}} + \frac{\delta\left(\alpha\log(X_t) + \alpha\log(1 - E_{it} - E_{-i})\right)}{1 - \alpha\delta} + \delta J_i(m_{it+1}, \eta_{it+1})\Bigg] \tag{23}$$

Because X_t is a known state variable, we group the $\log(X_t)$ terms, and note (as we showed

above) that the $\log(X_t)$ component simplifies to $\frac{\log(X_t)}{1-\alpha\delta}$. This gives the desired dynamic programming equation:

$$V_i(X_t, m_{it}, \eta_{it}) = \frac{\log(X_t)}{1-\alpha\delta} + \underbrace{\max_{E_{it}} \mathbb{E}\left[\log(E_{it}) + L_{it} - Z_{it}\,\mathbb{I}_{E_{it}>\bar{E}} + \frac{\delta\alpha\log(1-E_{it}-E_{-i})}{1-\alpha\delta} + \delta J_i(m_{it+1}, \eta_{it+1})\right]}_{J_i(m_{it},\eta_{it})} \quad (24)$$

Because $E_{-i} = (N-1)\bar{E}$ is constant, the bracketed expression depends on time only through (m_{it}, η_{it}), so J_i is a time-invariant function of the information state, which verifies the desired separability. □

Proposition 4 shows that we can additively separate the resource-stock problem from the information-state problem. The crux of the dynamic decision is captured by $J_i(m_{it}, \eta_{it})$, which is the expected present value of payoffs from the optimal learning and extraction strategy, net of the resource stock.

What does all of this imply about whether to comply or violate the Rāhui? The agent's problem is to choose between two "arms" in each period: comply ($\bar{E}$) or violate ($\hat{E}$). The value function J_i is just the larger of the comply vs. violate value functions, as follows:

$$J_i(m_{it}, \eta_{it}) = \max\{J_i^C(m_{it}, \eta_{it}), J_i^V(m_{it}, \eta_{it})\} \quad (25)$$

If agent i complies, her payoff is stochastic only due to the luck term L_{it}, which has an expected value of zero. She receives no punishment and does not learn, so her information state next period is unchanged. In that case, her learning dynamic programming equation is:

$$J_i^C(m_{it}, \eta_{it}) = \left(\log(\bar{E}) + \frac{\alpha\delta}{1-\alpha\delta}\log(1-\bar{E}-E_{-i})\right) + \delta J_i(m_{it}, \eta_{it}) \quad (26)$$

Instead, if i violates, she receives a different payoff, an expected punishment of m_{it}, and learns. Her information state evolves to (m_{it+1}, η_{it+1}) based on the shock ϕ_{it} received. In that case, her learning dynamic programming equation is:

$$J_i^V(m_{it}, \eta_{it}) = \left(\log(\hat{E}) + \frac{\alpha\delta}{1-\alpha\delta}\log(1-\hat{E}-E_{-i}) - m_{it}\right) + \delta\,\mathbb{E}\left[J_i(m_{it+1}, \eta_{it+1})\right] \quad (27)$$

4.2 Learning about superstition as a Bandit Problem

We have separated the value function into a resource stock component and a dynamic learning component. The former is a classic dynamic fisheries problem and the latter is a classic "two-armed normal bandit" problem, where one arm (comply) is a safe arm with a known payoff, and the other (violate) is an experimental arm with an unknown payoff that the agent can learn about by pulling it. Invoking the bandit relies on the fact that because the other agents are assumed to comply, E_{-i} is constant. Payoffs from both arms are therefore stationary, and the state of the violate arm changes only when i pulls it. These are the conditions required to apply the classic index results from the multi-armed bandit literature, which we describe below.

Beginning with the comply arm, if it was optimal to comply this period, it will be optimal next period (since the state is unchanged). Therefore, J_i^C represents the value of perpetual compliance and the value function can be written

$$J_i^C(m_{it}, \eta_{it}) = \underbrace{\left(\log(\bar{E}) + \frac{\alpha\delta}{1-\alpha\delta}\log(1-\bar{E}-E_{-i})\right)}_{K^C} + \delta J_i^C(m_{it}, \eta_{it}) \tag{28}$$

where K^C is the constant single-period payoff from complying. Then,

$$J_i^C = \frac{K^C}{1-\delta}$$

The key point is that the present value of the comply arm is a constant value; it is just the discounted present value of the per-period compliance payoffs and does not depend on the agent's beliefs (m_{it}, η_{it}).

The violate arm is more complicated because it returns both a value from the violation itself and a value of information. Furthermore, if it is optimal to pull the violate arm, it may subsequently be optimal to pull the violate arm again, or to switch to the comply arm; this fact induces interesting learning dynamics. The violation dynamic programming equation is:

$$J_i^V(m_{it}, \eta_{it}) = \underbrace{\left(\log(\hat{E}) + \frac{\alpha\delta}{1-\alpha\delta}\log(1-\hat{E}-E_{-i})\right)}_{K^V} - m_{it} + \delta\,\mathbb{E}\left[J_i(m_{it+1}, \eta_{it+1})\right]$$

where K^V is the constant single-period payoff from violating, excluding the expected punishment. Comparing with Equation 16, note that $\mu^c = K^V - K^C$. This value depends on the agent's information state (beliefs about supernatural punishment) in two distinct ways.

First, the current payoff depends on expected punishment ($-m_{it}$). Second, by violating, the agent knows she will have improved information in the next period; the information gained by violating the Rāhui will, in general, increase the continuation value, so returns a value of information. The $\delta\,\mathbb{E}\left[J_i(m_{it+1}, \eta_{it+1})\right]$ term is the expected continuation value, given that i violates today and proceeds optimally tomorrow with her new beliefs. The decision about whether to comply vs. violate can be summarized as follows:

Proposition 5. *Suppose all agents $j \neq i$ comply in every period. Then in the superstition model with learning, complying in period t is a best response for agent i if and only if her mean punishment belief satisfies:*

$$m_{it} \geq \mu^c + \Delta_i(m_{it}, \eta_{it}), \quad \textit{where} \quad \Delta_i(m_{it}, \eta_{it}) \equiv \delta\left(\mathbb{E}\left[J_i(m_{it+1}, \eta_{it+1})\right] - J_i^C\right)$$

is the value of information (VOI) from violating, μ^c is defined in Equation 16, and the expectation is taken over the shock ϕ_{it} that i observes upon violating. If complying is optimal in period t, it is optimal in every subsequent period.

Proof. Agent i complies if $J_i^C \geq J_i^V(m_{it}, \eta_{it})$, i.e.:

$$J_i^C \geq \left(\log(\hat{E}) + \frac{\alpha\delta}{1-\alpha\delta}\log(1 - \hat{E} - E_{-i}) - m_{it}\right) + \delta\,\mathbb{E}\left[J_i(m_{it+1}, \eta_{it+1})\right]$$

Using the fact that $J_i^C = K^C + \delta J_i^C$, grouping terms, and invoking the definition of μ^c delivers the condition. For the second claim, complying leaves (m_{it}, η_{it}) unchanged and, with the others complying, leaves E_{-i} unchanged, so the same condition holds in $t+1$ and, by induction, in every subsequent period. □

Corollary 1. *With private learning, universal compliance is an equilibrium if and only if every agent's initial beliefs satisfy*

$$m_{i0} \geq \mu^c + \Delta_i(m_{i0}, \eta_{i0}) \quad \textit{for all } i.$$

Along the equilibrium path no agent violates, beliefs never change, and every agent chooses $\bar{E}$ in every period.[27]

Proof. If the condition fails for some i, then by Proposition 5 agent i's best response to the others' compliance is to violate in period 0, thus proving necessity. Sufficiency is guaranteed

[27] Characterizing a full MPNE would require specifying best responses at belief profiles in which some agent prefers to violate; there, others' effort varies with their learning, so the bandit results we invoke may not apply. We do not characterize these partial-compliance dynamics.

because a deviation by i does not alter any other agent's information state, so all $j \neq i$ continue to comply and the aggregate effort of others is $E_{-i} = (N-1)\bar{E}$. Agent i's best response is then the optimal policy of her bandit problem, which by Proposition 5 implies compliance in period 0 and in every subsequent period. □

The μ^c term is the static threshold on supernatural beliefs that is needed to make the agent indifferent based only on the single-period payoffs. The VOI term, which appears only in the learning setting, is the extra benefit the agent expects to earn from violating. This value arises because by violating, the agent expects to have a new, more precise supernatural belief next period. Since the future value J_i is the maximum of the two arms, $\mathbb{E}[J_i]$ must be greater than or equal to J_i^C (the value of just one of the arms), which implies that the value of information is positive. This means the critical belief threshold required to induce compliance in the learning model is higher than in the static model. In other words, it is harder to sustain Rāhui compliance with learning than without it.

The additive separability of our problem implies that the "information" part of our problem, captured by the value function $J_i(m_{it}, \eta_{it})$, is precisely a two armed Normal Bandit. In that classic problem, the central tension is between exploration (violating Rāhui in our setting) vs. exploitation (complying with Rāhui in our setting). In Bandit theory, this canonical problem remained unsolved for decades until the insight by Gittins and Jones (1979), and, in parallel, by Weitzman (1979). The insight was that in each time step t, it is possible to independently calculate an "index" for each arm, based only on the state of that arm, and completely independently of the state of the other arms. Gittins and Jones referred to this as the "Gittins Index" and Weitzman referred to this as the "Reservation Price". The discovery implied that the optimal search among all bandits over an infinite horizon simply boils down to the stepwise comparison of the index of each arm. The Gittins Index notation for a Normal Multi-Armed Bandit is denoted $\lambda(u, v, \sigma^2, \beta)$,[28] where u is the mean of the prior, v is the variance of the prior, σ^2 is the variance of the payoff, and β is the discount factor. Gittins and Jones (1979) showed that in any period t, $\lambda(u, v, \sigma^2, \beta)$ can be calculated for all k arms, and it is optimal to sample (i.e. "pull") the arm with the largest Gittins Index. In our setting, the comply arm yields a known payoff K^C, so its index is simply K^C. The violate arm yields payoff $K^V + \phi_{it}$, whose unknown mean $K^V - \mu_i$ has prior $N(K^V - m_{it}, 1/\eta_{it})$ and whose observation variance $1/\nu_i$ is known, so its index is $\lambda(K^V - m_{it}, 1/\eta_{it}, 1/\nu_i, \delta)$. Weitzman (1979) independently derived a similar result in his "Pandora's box" paper, which seeks to open boxes with uncertain rewards in an optimal sequence to maximize the expected return. Weitzman referred to the result as calculating

[28] See Yao (2006).

a reservation price for each box in each timestep, and to open the box with the largest reservation price.

The Gittins Index (or Reservation Price) for an arm can be interpreted as the sure payoff that would be required such that the agent is indifferent between the sure payoff and the uncertain payoff with learning. If an arm has a larger Gittins Index, then sampling that arm is more lucrative because it would require a larger sure return to dissuade the agent from pulling that arm.

This structure allows us to fully characterize the Rāhui compliance decision. To do so, we make use of a key property of the Normal Gittins index called location invariance (additive shift by the mean of the prior): $\lambda(u, v, \sigma^2, \beta) = u + \lambda(0, v, \sigma^2, \beta)$ (Gittins, Glazebrook, and Weber, 2011). The index of the violate arm can therefore be decomposed into its expected payoff, $K^V - m_{it}$, plus an *exploration premium* that reflects only the value of learning.

Proposition 6. *Suppose all agents $j \neq i$ comply in every period. Then complying in period t is a best response for agent i if and only if i's belief satisfies:*

$$m_{it} \geq \mu^c + B_i(\eta_{it}, \nu_i, \delta), \quad \textit{where} \quad B_i(\eta_{it}, \nu_i, \delta) \equiv \lambda\left(0, \tfrac{1}{\eta_{it}}, \tfrac{1}{\nu_i}, \delta\right)$$

is the Gittins index of a normal arm with zero prior mean, prior variance $1/\eta_{it}$, and observation variance $1/\nu_i$ (the exploration premium). The exploration premium is nonnegative, and strictly positive whenever $\eta_{it} < \infty$, $\nu_i > 0$, and $\delta > 0$.

Proof. By Proposition 4, agent i's problem is a two-armed bandit with stationary, independent arms. The comply arm yields a known payoff K^C and leaves the information state unchanged, so its index is K^C. By location invariance, the violate arm has index $K^V - m_{it} + B_i(\eta_{it}, \nu_i, \delta)$. By the Gittins index theorem (Gittins and Jones, 1979), complying is optimal if and only if $K^C \geq K^V - m_{it} + B_i$. Rearranging and using $\mu^c = K^V - K^C$ gives the result. Nonnegativity follows because the agent can always stop pulling the violate arm, so the index of an arm is at least its expected payoff. Strict positivity follows because, with a non-degenerate prior and an informative signal, a single pull has positive learning value. □

Propositions 5 and 6 characterize the same compliance region, but in different ways. The value of information $\Delta_i(m_{it}, \eta_{it})$ is evaluated at the agent's current belief, and so depends on m_{it}. The exploration premium B_i depends only on the precision of the prior, the precision of the signal, and the discount factor. At the critical belief $m^*_{it} = \mu^c + B_i$ the two coincide, $\Delta_i(m^*_{it}, \eta_{it}) = B_i$. In particular, the condition in Corollary 1 can equivalently be written $m_{i0} \geq \mu^c + B_i(\eta_{i0}, \nu_i, \delta)$ for all i. Because B_i is independent of the agent's mean belief and of the effort levels $\bar{E}$, $\hat{E}$ and E_{-i}, we use it for the comparative statics that follow.

Proposition 6 also shows that compliance requires a substantially higher current payoff than violation. The payoff from complying must exceed the expected payoff from violating by at least the exploration premium, $K^C - (K^V - m_{it}) \geq B_i$. If the two payoffs were equal and $B_i > 0$, the agent would strictly prefer to violate, because violating also yields information. The exploration premium has several useful properties:

Proposition 7. *The exploration premium $B_i(\eta_{it}, \nu_i, \delta)$ has the following properties:*

(a) B_i is independent of the mean belief m_{it} and of the effort levels $\bar{E}$, $\hat{E}$ and E_{-i}

(b) B_i is decreasing in the precision of the prior (η_{it})

(c) B_i is nondecreasing in the precision of the signal (ν_i)

(d) B_i is increasing in the discount factor (δ)

Proof. (a) Self-evident from the definition of B_i and location invariance of the Normal Gittins index.

(b) The Normal Gittins index satisfies the scaling property $\lambda(0, v, \sigma^2, \beta) = \sigma\, \lambda(0, v/\sigma^2, 1, \beta)$, and $\lambda(0, v, 1, \beta)$ is increasing in the prior variance v (Gittins, Glazebrook, and Weber, 2011). Holding ν_i fixed, B_i is therefore increasing in $1/\eta_{it}$, i.e., decreasing in η_{it}.

(c) The proof makes use of the observation that an agent with better information can always pretend to have worse information, so can never do worse. Let $\nu' > \nu$. A signal with precision ν can be generated from a signal with precision ν' by adding independent mean-zero normal noise (with variance $1/\nu - 1/\nu'$). An agent observing the more precise signals can therefore replicate any policy available to an agent observing the less precise signals. Because the expected payoff of the violate arm, conditional on μ_i, does not depend on signal precision, the replicating policy earns the same expected payoff. The value of continuing to pull the violate arm, and thus its index, is therefore weakly higher under ν'.

(d) See Gittins, Glazebrook, and Weber (2011): the exploration premium of a normal arm is increasing in the discount factor.

□

Together with Proposition 6, these properties deliver the main comparative statics of the learning model. Part (a) implies that a higher mean belief m_{it} raises the left-hand side of the compliance condition one-for-one without changing the threshold, so stronger beliefs always favor compliance. For (b), when the prior is very precise, the agent is extremely sure of her

beliefs, so there is little to learn, which tips the balance toward compliance. For (c), consider the extreme where the data have perfect precision, so a single violation reveals the true punishment; then violating is extremely informative and the exploration premium is large. This relates to our idea of "background noise": higher ν_i translates to lower background noise. Since ν_i depends positively on τ_i (precision of the punishment Z_{it}) and ω_i (precision of luck L_{it}), more precise punishments and/or more predictable luck mean less background noise and a higher likelihood of violating. For (d), a more patient agent places more weight on the future benefits of learning, so she is more willing to experiment. In our model, η_{it} captures how certain agents are about supernatural punishment, and can be thought of as a measure of dogmatism. Dogmatism alone, however, does not guarantee compliance. As $\eta_{it} \to \infty$, the prior becomes degenerate and there is nothing left to learn, so the index of the violate arm converges to its expected payoff and $B_i \to 0$. The compliance condition in Proposition 6 then converges to the static condition $m_{it} \geq \mu^c$ of Proposition 3. A perfectly dogmatic agent thus has no motive to experiment, but she complies only if her (now perfectly precise) belief is that punishment is sufficiently severe. Conversely, an agent with a diffuse prior (small η_{it}) is tempted to experiment (because there is much to learn), so may violate even when her mean punishment belief exceeds μ^c. Overall, then, compliance depends both on the strength of beliefs (m_{it}) and the confidence with which they are held (η_{it}).

4.3 Interaction effects and the resilience of beliefs

How do the resource stock (X_t) and the information state (m_{it}, η_{it}) interact to affect agents' decisions? We showed above that when the other agents comply, agent i's value function is additively separable, which implies that her decision to comply or violate is independent of the level of the resource stock.

However, the resource stock and beliefs do interact in the determination of the physical harvest. Recall that harvest is given by $h_{it} = E_{it}X_t$. Invoking the results above, we conclude that: (1) Agent i's choice of effort share E_{it} depends on the information state (m_{it}, η_{it}) but is independent of the stock X_t; and (2) Agent i's harvest scales linearly with X_t.

Consequently, the cross-derivatives of the harvest function reveal how beliefs scale with the physical resource stock:

$$\frac{\partial^2 h_{it}}{\partial X_t \partial m_{it}} = \frac{\partial E_{it}}{\partial m_{it}} \leq 0 \tag{29}$$

$$\frac{\partial^2 h_{it}}{\partial X_t \partial \eta_{it}} = \frac{\partial E_{it}}{\partial \eta_{it}} \leq 0 \tag{30}$$

The signs of these expressions follow from Propositions 6 and 7. Because E_{it} is in practice

a discrete choice ($\bar{E}$ or $\hat{E}$), these expressions should be interpreted as the direction of the change in equilibrium effort accompanying a sufficiently large change in m_{it} or η_{it}, rather than as conventional point derivatives.

Regarding Equation 29, as the expected punishment m_{it} increases, the agent is more likely to comply (choose $\bar{E}$ rather than $\hat{E}$).[29] Because $\bar{E} < \hat{E}$, stricter beliefs lead to lower effort. This effect scales with the stock X_t: when the resource stock is abundant, agent i's harvest is more responsive to a large m_{it} than when the resource stock is low.

Regarding Equation 30, as the precision of beliefs η_{it} increases, the exploration premium B_i decreases (Proposition 7(b)). The incentive to violate for the sake of learning diminishes, increasing the likelihood of compliance and reducing extraction effort.

This separability, where the stock X_t affects the level of harvest but not the strategy E_{it}, harkens back to the early literature on the dynamics of economic growth. Our finding that the optimal effort share is independent of the resource stock is the resource economics analogue to canonical growth model by Brock and Mirman, 1972, also summarized and extended by Acemoglu, 2009.

This result has implications for the interaction between the state of the environment and the resilience of a belief system. In our model, low resource abundance, by itself, does not induce violation of a system of norms. Conversely, high resource abundance does not necessarily erode a system of norms. Rather, because the policy function E_{it} is independent of X_t, a community facing a resource collapse is no more likely to violate the Rāhui than a community enjoying a resource boom. The Rāhui is robust to shocks in the natural capital stock, unraveling only when the beliefs themselves (m_{it}, η_{it}) are conducive to violation.

4.4 Peer learning

In the model above, learning is a purely private endeavor; the only way for an agent to update her beliefs about supernatural punishment is to violate the Rāhui herself. That setup is attractive, in part, because it does not require that violations or punishments are public to sustain cooperation with Rāhui. However, in many real-world settings, violations, and possibly punishments, may be visible to others. If a neighbor fishes in violation of the Rāhui and his family falls ill, other members of the community may notice. Here we describe how allowing agents to learn from one another's violations changes the incentives identified above. Characterizing the full equilibrium of the resulting game is beyond the scope of this paper; it is related to problems of strategic experimentation studied by Bolton and Harris

[29] Strictly speaking, because the choice set is discrete ($\bar{E}$ vs $\hat{E}$), this derivative will typically be zero, but will jump from $\hat{E}$ down to $\bar{E}$ when m_{it} (or η_{it}) is sufficiently large.

(1999) and Keller, Rady, and Cripps (2005), though these papers do not study a dynamic natural resource. Allowing for peer learning introduces several interesting incentive margins that are not present under private learning, which we characterize here.

To make progress, we must make some new assumptions about the peer learning environment. We assume that all agents observe whether each other agent complies with the Rāhui, and observe each other's shocks. Suppose all agents face the same unknown mean punishment μ, so that a violating agent j's punishment is $Z_{jt} \sim N(\mu, \tau)$, and that luck $L_{jt} \sim N(0, \omega)$ is independent across agents and periods, with precisions τ and ω common to all agents.[30] Each violator's overall shock $\phi_{jt} = L_{jt} - Z_{jt} \sim N(-\mu, \nu)$ (with $\nu = (\tau^{-1} + \omega^{-1})^{-1}$ as before) is publicly observed. Non-violators' shocks may also be observed, but they are uninformative because they contain only luck, which is independent across agents and reveals nothing about punishment. Agents begin with a common, normal prior over the mean punishment, $\mu \sim N(m_0, \eta_0)$. Because all information is public, all agents share the same posterior thereafter. In every period, agents choose effort simultaneously, and shocks are then observed by all agents. If q_t agents violate in period t, beliefs update according to

$$m_{t+1} = m_t \left(\frac{\eta_t}{\eta_t + q_t \nu} \right) - \bar{\phi}_t \left(\frac{q_t \nu}{\eta_t + q_t \nu} \right) \tag{31}$$

$$\eta_{t+1} = \eta_t + q_t \nu \tag{32}$$

where $\bar{\phi}_t$ is the average shock among the q_t violators. These expressions reduce to Equations 19 and 20 when $q_t = 1$, and leave beliefs unchanged when $q_t = 0$.

How does peer learning change incentives under this new setup? On a universal compliance path, no one violates, so no shocks are generated, and beliefs remain constant, just as under private learning. But peer learning alters this calculus in two important ways. First, it changes the consequences of a violation that deviates from universal compliance, because the violator's shock is now observed by everyone. Second, it changes agent i's incentive to violate whenever she expects others to violate. These two observations give rise to three incentive margins that are not present in the private learning model, as well as a change in the dynamics of beliefs. We analyze them in turn.

4.4.1 Free-riding effect

All agents realize that others' violations provide valuable information about μ. That information is a non-rival and non-excludable public good, but its provision is privately costly in

[30] In the private learning model, mean punishment may differ across agents, but nothing depends on this. Under peer learning, a common mean is needed for others' shocks to be informative; with correlated agent-specific means, they would be partially informative.

expectation. Because Rāhui-compliant agents still learn from others' violations, the informational return to agent i from her own violation is only the value of one additional signal on top of those her peers provide. If this marginal value is smaller than the value of being the sole source of information (as one would expect), peer learning weakens the incentive to violate the Rāhui. Each agent would prefer to comply and let others bear the expected punishment from testing it. This is the free-rider problem common to many strategic experimentation settings, and it tends to strengthen compliance.

4.4.2 Encouragement effect

A subtler incentive for experimentation emerges under peer learning, which, following Bolton and Harris (1999), we call the encouragement effect. If i violates, her signal is public. If it suggests a low punishment, it will encourage others to violate the Rāhui, and i will glean valuable (also public) information from their subsequent violations. Anticipating this chain of events, i places a higher value on experimenting today.[31] Bolton and Harris (1999) show that the encouragement effect can offset the free-rider effect, so we can conjecture that peer learning could accelerate or suppress experimentation.

4.4.3 Resource abundance effect

Violations of the Rāhui also have consequences for the resource stock, which affect incentives. As described above, agent i's violation can reveal a low punishment, which induces others to violate. Their additional extraction reduces the resource stock available to all agents, including i, lowering her future payoff whether or not she complies in the future. This mechanism is not present under private learning, because in that setting others learn nothing from i's violation. Under peer learning, an agent who tests the Rāhui risks unraveling her neighbors' compliance and, if that happens, bears part of the cost through a reduced resource stock. This margin therefore makes experimentation less attractive. Unlike free-riding, this effect and the encouragement effect both apply to a deviation from universal compliance. Both arise when good news leads others to violate, but they work in opposite directions. And like the encouragement effect, this effect is one-sided when others are already complying. When others would otherwise violate, a severe shock deters them, which benefits i through a larger resource stock.

[31] If instead her signal suggests a severe punishment, others may be discouraged from experimenting, reducing the free information i receives. This offsetting force is limited. Starting from universal compliance, others are not experimenting to begin with, so only favorable signals matter. And when others would otherwise violate, their information is most valuable to i after a favorable signal; after an unfavorable one she is likely to comply anyway.

4.4.4 Rapid, synchronous learning

Public signals synchronize learning. Under private learning, even if multiple agents violate in the same period, each learns only from her own experience, so beliefs evolve separately and somewhat sluggishly. Under peer learning, every agent learns from i's violation. If it reveals a mild punishment, every agent's belief falls immediately; if it reveals a severe punishment, every agent's belief rises, reinforcing compliance across the community. If the true punishment is small or zero, mild shocks are more likely, and beliefs may erode rapidly. However, because extraction efforts are strategic substitutes, violations by some agents reduce the temptation for others to violate (Section 3.4). While synchronization alters the pace and correlation of learning, it does not, by itself, change any given agent's incentive to violate.

4.4.5 Net effect

On net, the free-riding and resource abundance effects favor compliance, the encouragement effect favors violation, and synchronous learning can lead to rapid unraveling if the true punishment is small or zero. Conversely, it can lead to entrenched compliance if the true punishment is large. Which force dominates depends on parameters and on the equilibrium being played. Characterizing the equilibria of the full game, in which experimentation decisions interact both through information and through the resource stock, is left for future work.

5 Robustness of the Rāhui

We now examine the robustness of the Rāhui institution to plausible changes. We begin by conducting comparative statics of the compliance decision on key model parameters. We then extend the model to the case of a selfish chief, who designs the parameters of the Rāhui to maximize his private rent capture.

5.1 Robustness to exogenous change

Here we summarize how the fundamental parameters of the model affect the stability of the Rāhui. In our model, sustainable exploitation in the commons depends entirely on the agent's decision to comply with the Rāhui, $\bar{E}$. This decision involves a tension between the immediate economic temptation to over-extract and the belief in supernatural punishment, mediated by the potential value of learning the truth about supernatural punishment.

Table 3 summarizes the comparative statics of the model, which we prove in Appendix B. We classify parameters into two categories: *Economic and Demographic* parameters, which affect the resource stock dynamics and payoffs, and *Information* parameters, which affect the learning process and belief structure. These comparative statics offer some insight into how forces exogenous to the model – such as immigration or changes in technology – may change key parameters that determine the stability of the Rāhui over time.

Table 3: Comparative statics: Effect of parameter increases on compliance with Rāhui

Param.	**Description**	**Compliance**	**Mechanism**
Economic and Demographic Parameters			
N	Number of agents	$\downarrow$	Smaller individual share of cooperative yield increases temptation to violate.
α	Inverse growth rate	$\downarrow$	Slower resource recovery increases scarcity rent, increasing the temptation to violate.
δ	Discount factor	$\downarrow$	Raises the static temptation to violate and raises the exploration premium, incentivizing more experimentation.
Information (Belief) Parameters			
m_{it}	Mean punishment belief	$\uparrow$	Directly increases the expected cost of violation.
η_{it}	Belief precision	$\uparrow$	Stronger conviction reduces the exploration premium from experimenting.
ν_i	Signal precision	$\downarrow$	Less background noise makes learning more efficient, raising the exploration premium and the incentive to violate the Rāhui.

Note: The Compliance column indicates whether an increase in the parameter makes compliance *more* likely ($\uparrow$) or *less* likely ($\downarrow$), holding all else constant.

5.1.1 Resource pressures

The stability of Rāhui is sensitive to demographic and ecological pressures. As the community size N grows, the share of the resource allocated to each compliant agent diminishes (because $\bar{E}$ is proportional to $1/N$). Because the marginal benefit of violation remains high, the "static" temptation to violate Rāhui increases. This implies that larger populations require stronger supernatural beliefs (higher m_{it}) to sustain the same level of cooperation as smaller populations. Forces such as declining mortality rates or immigration will tend to increase group sizes, making Rāhui more fragile. Similarly, a resource with lower intrinsic growth (higher α) increases the biological stakes of the extraction game, requiring stricter adherence to beliefs to prevent collapse.

Perhaps counter-intuitively, greater patience (higher δ) undermines compliance with Rāhui. In standard repeated games, patience facilitates cooperation. Here, patience works against compliance through two channels. First, even without learning, a higher δ raises the static threshold μ^c. Part of the reason is that a more patient Rāhui designer (who shares a common discount factor with the rest of the community) sets a more restrictive Rāhui ($\bar{E}$ is decreasing in δ). On net, this raises the temptation to violate. Second, greater patience increases the present value of *future* knowledge. A patient agent is more willing to experiment now (and bear the expected punishment today) in order to resolve uncertainty for the future. This works in the same direction as the previous effect because it increases the incentive to violate. Increased lifespans or better access to credit may be associated with higher discount factors, which could undermine Rāhui compliance.

5.1.2 Beliefs and background noise

The information parameters also play an important role in sustaining beliefs. Compliance is reinforced by what we term "background noise" (low ν_i). When ν_i is low, perhaps due to high variability in luck (L_{it}) or vague definitions of supernatural punishment (Z_{it}), the signal-to-noise ratio of a violation is low. A violation yields little information, the exploration premium approaches zero, and the agent complies even with moderate beliefs. Pre-industrial societies, more dependent on agriculture and with greater disease exposure, presumably face higher background noise.

Conversely, "modernization" can be interpreted as an increase in ν_i (e.g., through scientific understanding of weather or disease that reduces the variance of L_{it}). As the background noise abates, the supernatural signal becomes potentially easier to isolate. This raises the exploration premium, incentivizing agents to violate the Rāhui to test their beliefs. Finally, the precision of the prior belief over punishment from violating Rāhui, η_{it}, is a measure of re-

ligious dogmatism. Agents who are uncertain about punishment (low η_{it}) are curious; agents who are dogmatic (high η_{it}) perceive little value in testing the Rāhui, because they expect a violation to provide little new information with which to update their beliefs. Dogmatism thus weakens the incentive to learn. With extreme religious dogmatism, compliance depends entirely on the static belief about punishment in Proposition 3. For this reason, dogmatic belief in a severe punishment can substitute for external enforcement. Interestingly, dogmatic belief in a mild punishment can have the opposite effect of locking in violations of Rāhui.

5.2 The political economy of Rāhui: Exploiting devotion

We now examine whether Rāhui compliance is robust to an alternative specification of the Rāhui design. Thus far, we have assumed a benevolent Rāhui planner who establishes norms that maximize social welfare (Section 3.3). Here we show that a more cynical take on Rāhui design could persist in equilibrium.

Rather than setting $\bar{E}$ to maximize social welfare, suppose that a self-interested Rāhui designer (the "selfish chief") sets $\tilde{E}$ to maximize rent capture. He accomplishes this by setting a more stringent Rāhui, so resource extraction by all other agents is lower, thus freeing up resource stock to be exploited by the chief himself. Naturally, the chief must balance the desire to extract rents against the participation by agents, which, of course, hinges on their beliefs. This tension generally gives rise to an optimal interior level of Rāhui stringency, which we derive below.

5.2.1 The selfish chief's problem

The selfish chief chooses the Rāhui limit $\tilde{E}$ and his own effort E_s to maximize his private present value returns. To be internally consistent, consider the case when the chief is either a non-believer or is exempt from punishment (for example, he is an intermediary between the supernatural and humans).[32] To ensure participation by all of his subjects, he must set these parameters to satisfy the other agents' compliance conditions. To conform to the analysis above, we focus on rules that sustain compliance by all $N-1$ devotees.[33]

With all devotees complying at $\tilde{E}$, the chief's value function is

$$V_s(X_t;\tilde{E},E_s)=\frac{\log(X_t)}{1-\alpha\delta}+\frac{1}{1-\delta}\left[\log(E_s)+\frac{\alpha\delta}{1-\alpha\delta}\log\left(1-E_s-(N-1)\tilde{E}\right)\right] \tag{33}$$

[32]If not, the chief would have to temper his own extraction effort to avoid being punished himself.

[33]The chief may, under some combinations of parameters, further increase rents with a rule that some devotees violate. Characterizing the chief's problem across partial-compliance regimes is beyond our scope.

(see Equation 13). Holding $\tilde{E}$ fixed, the chief's payoff is maximized at his unconstrained best response $E_s^u(\tilde{E}) \equiv (1-\alpha\delta)(1-(N-1)\tilde{E})$ (Equation 5), and it is strictly decreasing in $\tilde{E}$.

We now analyze arbitrary devotee j's compliance decision. Given the chief's effort (E_s) and the other devotees' compliance ($\tilde{E}$), devotee j's residual effort is $R(\tilde{E}, E_s) \equiv 1 - E_s - (N-2)\tilde{E}$. Her optimal violation is $(1-\alpha\delta)R$ (this was shown previously), and, repeating the analysis from Section 3.4 with $\bar{E}$ replaced by $\tilde{E}$, her compliance threshold (in the absence of learning) is

$$\mu^c(\tilde{E}, E_s) = \log\left(\frac{(1-\alpha\delta)R}{\tilde{E}}\right) + \frac{\alpha\delta}{1-\alpha\delta}\log\left(\frac{\alpha\delta R}{R-\tilde{E}}\right). \tag{34}$$

By Proposition 6, devotee j complies if and only if their mean beliefs are sufficiently large: $m_j \geq \mu^c(\tilde{E}, E_s) + B_j(\eta_j, \nu_j, \delta)$. Recall from Proposition 7(a) that B_j is an information term that does not depend on effort. Thus, we can rewrite the condition as $\mu^c(\tilde{E}, E_s) \leq \kappa_j$, where $\kappa_j \equiv m_j - B_j(\eta_j, \nu_j, \delta)$ is devotee j's *compliance slack*. Intuitively, κ_j measures how much temptation devotee j's beliefs can withstand; it is her expected punishment from violating, net of the value she places on learning. Because the Rāhui applies equally to all devotees, every devotee faces the same threshold $\mu^c(\tilde{E}, E_s)$, and the binding devotee, indexed by b, is the one with the smallest compliance slack, $\kappa_b \equiv \min_j \kappa_j$. Intuitively, this devotee may have the weakest belief, but this is not necessarily the case because a devotee with a stronger but less confident belief may have a larger exploration premium (this is essentially the intuition underpinning the bandit problem).

It is helpful to see how the threshold μ^c depends on $\tilde{E}$ and E_s. Whenever violation is tempting ($(1-\alpha\delta)R > \tilde{E}$),

$$\frac{d\mu^c}{d\tilde{E}} < 0 \qquad \text{and} \qquad \frac{\partial\mu^c}{\partial E_s} < 0. \tag{35}$$

The first result says that a tighter Rāhui raises the temptation to violate, because it raises the cost of compliance for every devotee. The second says that more extraction by the chief lowers the devotees' temptation to violate, because it shrinks each devotee's residual share and so the gain from violating.[34] We summarize the selfish chief results in the following proposition:

Proposition 8. *Assume the devotees' beliefs are strong enough to sustain the welfare-maximizing Rāhui, so $\kappa_b \geq \mu^c$, and focus on a Rāhui that sustains compliance by all devotees. The selfish chief's optimal choice of Rāhui and his own effort $(\tilde{E}^*, E_s^*)$ satisfies:*

[34] Differentiating Equation 34, $\partial\mu^c/\partial\tilde{E}\,|_R = \frac{\tilde{E}-(1-\alpha\delta)R}{(1-\alpha\delta)\tilde{E}(R-\tilde{E})} < 0$ and $\partial\mu^c/\partial R = \frac{(1-\alpha\delta)R-\tilde{E}}{(1-\alpha\delta)R(R-\tilde{E})} > 0$. Since $\partial R/\partial E_s = -1$ and $\partial R/\partial\tilde{E} = -(N-2)$, both inequalities in Equation 35 follow.

(a) The chief ensures that the compliance constraint of the binding devotee just binds, $\mu^c(\tilde{E}^*, E_s^*) = \kappa_b$. The Rāhui is always more stringent than the welfare-maximizing rule, $\tilde{E}^* < \bar{E}$.[35]

(b) The chief extracts more than his own best response to the rule he sets, $E_s^* > E_s^u(\tilde{E}^*)$, and more than per-agent effort in the social optimum, $E_s^* > \bar{E}$.

(c) Aggregate extraction effort $E_S = E_s^* + (N-1)\tilde{E}^*$ exceeds the socially optimal level $1 - \alpha\delta$.

Proof. Let $f(\tilde{E}, E_s)$ denote the term in V_s in square brackets. This is strictly decreasing in $\tilde{E}$ and strictly concave in E_s with maximum at $E_s^u(\tilde{E})$. We prove each result in turn.

(a) If $\mu^c(\tilde{E}^*, E_s^*) < \kappa_b$, the chief could lower $\tilde{E}$ slightly while remaining feasible, strictly raising f. This implies $\mu^c(\tilde{E}^*, E_s^*) = \kappa_b$.

For any $\tilde{E} \geq \bar{E}$, the chief's payoff is at most $f(\bar{E}, E_s^u(\bar{E}))$, since f is decreasing in $\tilde{E}$. Note that $E_s^u(\bar{E}) = \hat{E} > \bar{E}$. At $(\bar{E}, \hat{E})$, the constraint does not bind (because μ^c is decreasing in E_s, $\mu^c(\bar{E}, \hat{E}) < \mu^c(\bar{E}, \bar{E}) = \mu^c \leq \kappa_b$). The chief can therefore lower $\tilde{E}$ slightly below $\bar{E}$, holding $E_s = \hat{E}$, remain feasible, and strictly exceed $f(\bar{E}, E_s^u(\bar{E}))$. Thus $\tilde{E}^* < \bar{E}$.

(b) To show that the chief over-extracts, suppose $E_s^* \leq E_s^u(\tilde{E}^*)$. Raising E_s slightly does not lower f to first order, because f is weakly increasing in E_s below its peak, but it strictly relaxes the binding constraint, because $\partial\mu^c/\partial E_s < 0$. The chief could then lower $\tilde{E}$ and strictly raise f, a contradiction. Thus $E_s^* > E_s^u(\tilde{E}^*)$. Moreover,

$$E_s^u(\tilde{E}^*) - \bar{E} = (1-\alpha\delta)(N-1)\left(\frac{1}{N} - \tilde{E}^*\right) > 0,$$

since $\tilde{E}^* < \bar{E} < \frac{1}{N}$. This proves (b).

(c) To see that aggregate effort is higher than welfare-maximizing, use (b), $E_S > E_s^u(\tilde{E}^*) + (N-1)\tilde{E}^* = (1-\alpha\delta) + \alpha\delta(N-1)\tilde{E}^* > 1 - \alpha\delta$, proving (c).

A larger κ_b grows the chief's feasible effort set. Because the constraint binds at the optimum, the previous optimum becomes slack, so by the first step the chief's payoff strictly increases. Finally, $\kappa_b = m_b - B_b(\eta_b, \nu_b, \delta)$ is increasing in m_b, increasing in η_b, and nonincreasing in ν_b, by Proposition 7(b)–(c).

[35] The chief's payoff is strictly increasing in κ_b, and therefore weakly increasing in the binding devotee's mean belief m_b and belief precision η_b, and nonincreasing in her signal precision ν_b.

□

Proposition 8 shows that the selfish chief has two reinforcing incentives to over-extract. First, and most directly, more extraction raises his returns. Second, because extraction efforts are strategic substitutes, his extraction reduces each devotee's residual share and with it her temptation to violate. This allows the selfish chief to impose a stricter Rāhui. This second channel operates through the devotees' share of the resource, not through the level of the stock, which, by Proposition 4, does not affect compliance decisions.

These results suggest that societies in which beliefs are strong and confidently held (high κ_b) are highly susceptible to exploitation and rent capture by elites. Rent capture by a selfish chief is maximized when learning about superstition is noisy, since the information value of testing the Rāhui is low, allowing the chief to limit devotees' access to resources and extract them himself.

In addition, the selfish chief extension illustrates that persistence need not imply welfare-improving coordination. A belief system can be robust to falsification, and therefore survive to be documented, while still serving the interests of an elite at the expense of the community.

6 Discussion and conclusion

Social scientists have noted the ability of small scale societies to sustain cooperation in a common pool resource setting using only informal institutions, norms and repeated interaction (e.g., Ostrom, 1990). While social interactions may be sufficient to overcome incentives for overextraction, the prevalence of beliefs in supernatural punishment in many settings suggests that another force may be at work. We model the role of beliefs in supernatural punishment for cooperation in the commons. Based on a new database of beliefs around the world, the model shows that cooperation (and beliefs) can be sustained even if violation is accompanied by information about punishment. What is key in the model – and common to nearly all of the cases in the database – is that the punishment signal is obscured by other background shocks that cannot be separated from punishment outcomes. This impedes learning, lowers the information value of violation, and allows beliefs to be sustained in conditions that would unequivocally lead to violation in the absence of superstitious beliefs.

At the same time, the model's comparative statics highlight a number of reasons why traditional beliefs have eroded over time in so many cases. Modernization exerts pressure in the form of population growth that increases the number of resource users, and their incentive to violate. New sources of information – whether through greater human mobility or from information technology – weakens the precision of beliefs, increasing the information

value of violation.[36] As the economy grows and the shocks that impede learning, such as crop failure and disease, become less common, learning becomes more efficient. All of these forces may improve people's well-being but at some cost to beliefs that otherwise have the potential to support cooperative outcomes in a common pool resource setting.

A minority of cases in our database involve punishments that are directly related to the resource stock. For example, the Sto:lo of British Columbia believe that failure to perform salmon rituals reduces future salmon runs. In these cases, punishment is tied directly to the resource stock rather than being an idiosyncratic shock. While this formally departs from the assumptions in our model, we can use our model to speculate on its consequences. Stock-dependent punishment generates a new kind of uncertainty for potentially violating agents. If others' extraction is not perfectly observed, an agent who violates the Rāhui and then observes a low future stock cannot distinguish supernatural punishment (arising from her own violation) from the biological consequences of extraction by others. There are two reasons this may not fundamentally alter our conclusions about the robustness of Rāhui. First, uncertainty about others' extraction acts as an additional source of background noise about punishment, and by Proposition 7(c), noisier signals lower the exploration premium and favor compliance. Second, when agent i observes a low stock, she may interpret it as evidence that others are extracting more. Because efforts are strategic substitutes, this would lower her temptation to violate. What reduces her temptation here is the extraction she expects from others. The low stock itself plays no role, since in our baseline model the level of the stock does not affect compliance decisions (Proposition 4). Formalizing these channels beyond the speculation provided here would involve a substantial departure from our formal framework because it breaks the additive separability between resource dynamics and belief updating that underlies the bandit problem, so we leave a full treatment to future work.

Our empirical analysis in Section 2.3 is primarily descriptive. At the same time, the beliefs we observe are a selected sample of those robust enough to persist until documented. If true supernatural punishment is weak relative to the beliefs that sustain compliance, and if belief systems whose compliance unravels are less likely to be documented, our model implies that documented belief systems should tend to feature vague punishment. The near-universal prevalence of vague punishment in our database is consistent with this selection mechanism. Because the database contains only relatively long-lived belief systems, however, it cannot by itself distinguish this mechanism from other determinants of persistence or documentation. A more direct empirical test of the model could take one of at least three

[36]Other forces that increase the value of violation, such as a lower extraction cost through technological improvement, will also tend to decrease compliance.

directions, all of which face challenges. First, the beliefs observed around the world today could be taken as a steady state outcome, where the conditions associated with cooperation or unraveling (including plausibly exogenous variation in these conditions) can be compared against the strength of current beliefs.[37] Our database serves as the most comprehensive (to our knowledge) inventory of beliefs about supernatural beliefs and natural resources – but it is not up to the task: sources date from over two centuries of writing, beginning with early anthropological records from the time of European arrival. This makes the steady state interpretation more challenging. Second, we could have taken a field experimental approach, but this would require either manipulating incentives for violation, or beliefs themselves, both of which pose ethical challenges. Third, we could have moved into the lab, and attempted to create a setting that mimics the set up of the model. This, however, would have been unable to capture actual beliefs in supernatural punishment. Instead, we rely on existing literature to ground the model and show that documented belief systems adhere to the model's parameters and results.

We are largely silent on the question of where belief systems come from in the first place, how they evolve and whether natural resource management is, in fact, a central motivation in their creation. These are topics of active debate in the anthropology, sociology and psychology literatures. We leave the question of "beliefs design" and the associated question of its efficacy in preserving natural resources for other work.

[37] This is similar in spirit to the empirical test in Espín-Sánchez, Gil-Guirado, and Ryan (2026); a more direct analog of their insight that praying for rain should be more common where the rainfall hazard function is increasing would focus on the subset of Rāhui that specify punishment related to the stock of the resource. These should be less common in, for example, fisheries for pelagic species, where the feedback between fishing effort and fish stock is weak.

References

Abatayo, Anna Lou and John Lynham (2016). "Endogenous vs. exogenous regulations in the commons". *Journal of Environmental Economics and Management* 76, pp. 51–66.

Acemoglu, Daron (2009). *Introduction to modern economic growth*. Princeton University Press.

Acemoglu, Daron et al. (2011). "Bayesian learning in social networks". *The Review of Economic Studies* 78.4, pp. 1201–1236.

Aguiar, Leonardo Matheus Pereira, Eraldo Medeiros Costa Neto, and Dídac Santos Fita (2023). "Caipora and the conservation of natural resources in tropical forests in the South Recôncavo region, Bahia State, Northeast Brazil". *Ethnobiology and Conservation* 12.

Amer Meziane, Mohamed (2021). *Des empires sous la terre. Histoire écologique et raciale de la sécularisation*. Paris: Éditions La Découverte.

— (2023). *Au Bord des Mondes: Vers une anthropologie métaphysique*. Bruxelles: Vues de L'esprit.

Andréfouët, Serge and Mehdi Adjeroud (2019). "French polynesia". *World Seas: an Environmental Evaluation*. Elsevier, pp. 827–854.

Baland, Jean-Marie and Jean-Philippe Platteau (1996). *Halting degradation of natural resources: Is there a role for rural communities?* Food & Agriculture Org.

Bambridge, Tamatoa, ed. (2016). *The Rahui: Legal pluralism in Polynesian traditional management of resources and territories*. Pacific series. Acton, ACT, Australia: Australian National University Press.

Banerjee, Abhijit V (1992). "A simple model of herd behavior". *The Quarterly Journal of Economics* 107.3, pp. 797–817.

Bardhan, Pranab (1993). "Symposium on management of local commons". *Journal of Economic Perspectives* 7.4, pp. 87–92.

Basnet, Janga Bdr. and Tashi Dendup (2020). "Role of traditional beliefs and cultural myths in environmental conservation: A case study from Kabjisa, Punakha, Bhutan". *SSRN Electronic Journal*.

Bénabou, Roland and Jean Tirole (2006). "Belief in a just world and redistributive politics". *The Quarterly Journal of Economics* 121.2, pp. 699–746.

— (2011). "Identity, morals, and taboos: Beliefs as assets". *The Quarterly Journal of Economics* 126.2, pp. 805–855.

Berry, Donald A. and Bert Fristedt (1985). *Bandit problems: Sequential allocation of experiments*. Vol. 5. Monographs on Statistics and Applied Probability. London: Chapman and Hall.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992). "A theory of fads, fashion, custom, and cultural change as informational cascades". *Journal of Political Economy* 100.5, pp. 992–1026.

Bolton, Patrick and Christopher Harris (1999). "Strategic experimentation". *Econometrica* 67.2, pp. 349–374.

Brock, William A and Leonard J Mirman (1972). "Optimal economic growth and uncertainty: The discounted case". *Journal of Economic Theory* 4.3, pp. 479–513.

Buccione, Giulia (2023). "Religious messaging and adaptation to water scarcity: Evidence from jordan". *Working Paper*.

Bursztyn, Leonardo, Georgy Egorov, and Stefano Fiorin (2020). "From extreme to mainstream: The erosion of social norms". *American Economic Review* 110.11, pp. 3522–3548.

Casari, Marco and Charles R Plott (2003). "Decentralized management of common property resources: Experiments with a centuries-old institution". *Journal of Economic Behavior & Organization* 51.2, pp. 217–247.

Colding, Johan and Carl Folke (2001). "Social taboos:"invisible" systems of local resource management and biological conservation". *Ecological Applications* 11.2, pp. 584–600.

Cox, Michael, Sergio Villamayor-Tomas, and Yasha Hartberg (Feb. 2014). "The role of religion in community-based natural resource management". *World Development* 54, pp. 46–55.

Deopa, Neha and Daniele Rinaldo (2025). "Sacred ecology". *Working Paper*.

Easley, David and Nicholas M. Kiefer (1988). "Controlling a Stochastic Process with Unknown Parameters". *Econometrica* 56.5, pp. 1045–1064.

Elster, Jon (1989). "Social norms and economic theory". *Journal of Economic Perspectives* 3.4, pp. 99–117.

Espín-Sánchez, José-Antonio, Salvador Gil-Guirado, and Nicholas Ryan (2026). "Praying for rain". *The Quarterly Journal of Economics*.

Fehr, Ernst and Simon Gächter (2000). "Cooperation and punishment in public goods experiments". *American Economic Review* 90.4, pp. 980–994.

Fehr, Ernst and Andreas Leibbrandt (2011). "A field study on cooperativeness and impatience in the tragedy of the commons". *Journal of Public Economics* 95.9-10, pp. 1144–1155.

Fudenberg, Drew and David K Levine (1993). "Self-confirming equilibrium". *Econometrica* 61.3, pp. 523–545.

— (2006). "Superstition and rational learning". *American Economic Review* 96.3, pp. 630–651.

Garcia, Maximiliano and Jose Belmar (2026). "Governing environmental markets: Evidence from irrigation in water markets". *Working Paper*.

Gittins, John, Kevin Glazebrook, and Richard Weber (2011). *Multi-armed Bandit Allocation Indices*. 2nd ed. Chichester: Wiley.

Gittins, John C and David M Jones (1979). "A dynamic allocation index for the discounted multiarmed bandit problem". *Biometrika* 66.3, pp. 561–565.

Hackett, Steven, Edella Schlager, and James Walker (1994). "The role of communication in resolving commons dilemmas: Experimental evidence with heterogeneous appropriators". *Journal of Environmental Economics and Management* 27.2, pp. 99–126.

Haseeb, Muhammad (2024). "Resource scarcity and cooperation". *Working Paper*.

Henrich, Joseph (2016). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton, NJ: Princeton University Press.

Iannaccone, Laurence R (1998). "Introduction to the Economics of Religion". *Journal of economic literature* 36.3, pp. 1465–1495.

Iyer, Sriya (2016). "The new economics of religion". *Journal of Economic Literature*, p. 54.

Johnson, Dominic and Oliver Krüger (2004). "The good of wrath: Supernatural punishment and the evolution of cooperation". *Political Theology* 5.2, pp. 159–176.

Keller, Godfrey, Sven Rady, and Martin Cripps (2005). "Strategic experimentation with exponential bandits". *Econometrica* 73.1, pp. 39–68.

Kosfeld, Michael and Devesh Rustagi (2015). "Leader punishment and cooperation in groups: Experimental field evidence from commons management in ethiopia". *American Economic Review* 105.2, pp. 747–783.

Le Rossignol, Etienne, Sara Lowes, and Nathan Nunn (2022). "Traditional supernatural beliefs and prosocial behavior". *Working Paper*.

Levhari, David and Leonard Mirman (1980). "The great fish war: An example using a dynamic Cournot-Nash solution". *The Bell Journal of Economics* 11.1, pp. 322–334.

Lowes, Sara (2021). "Ethnographic and field data in historical economics". *The Handbook of Historical Economics*. Elsevier, pp. 147–177.

Moreman, Christopher M (2017). *Beyond the threshold: Afterlife beliefs and experiences in world religions*. 2nd ed. Rowman & Littlefield.

Noack, Frederik and Christopher Costello (2024). "Credit markets, property rights, and the commons". *Journal of Political Economy* 132.7.

Norenzayan, Ara (2013). *Big gods: How religion transformed cooperation and conflict*. Princeton University Press.

North, Douglass C (1990). "Institutions, institutional change and economic performance". *Cambridge Univ Pr*.

Nunn, Nathan and Raul Sanchez de la Sierra (May 2017). "Why being wrong can be right: Magical warfare technologies and the persistence of false beliefs". *American Economic Review* 107.5.

Ostrom, Elinor (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.

— (2000). "Collective action and the evolution of social norms". *Journal of Economic Perspectives* 14.3, pp. 137–158.

Ostrom, Elinor, Roy Gardner, and James Walker (1994). *Rules, games, and common-pool resources*. University of Michigan Press.

Ostrom, Elinor, James Walker, and Roy Gardner (1992). "Covenants with and without a sword: Self-governance is possible". *American Political Science Review* 86.2, pp. 404–417.

Purzycki, Benjamin Grant et al. (2016). "Moralistic gods, supernatural punishment and the expansion of human sociality". *Nature* 530.7590, pp. 327–330.

Rigo, Bernard (Mar. 2016). "Political power and rahui in ancient polynesian society". *The Rahui: Legal Pluralism in Polynesian Traditional Management of Resources and Territories*. 1st. ANU Press.

Rothschild, Michael (1974). "A Two-Armed Bandit Theory of Market Pricing". *Journal of Economic Theory* 9.2, pp. 185–202.

Sargent, Thomas J (2008). "Evolution and intelligent design". *American Economic Review* 98.1, pp. 55–84.

Sethi, Rajiv and Eswaran Somanathan (1996). "The evolution of social norms in common property resource use". *The American Economic Review* 86.4, pp. 766–788.

Shariff, Azim F. and Ara Norenzayan (Mar. 2011). "Mean gods make good people: Different views of God predict cheating behavior". *International Journal for the Psychology of Religion* 21.2, pp. 85–96.

Signor, P. W. I. and J. H. Lipps (1982). "Sampling bias, gradual extinction patterns and catastrophes in the fossil record". *Geological Implications of Impacts of Large Asteroids and Comets on the Earth*. Ed. by L. T. Silver and P. H. Schultz. Boulder, CO: Geological Society of America, pp. 291–296.

Teuira, Henry (1928). *Tahiti aux temps anciens*. Bishop Museum D'Honolulu.

Tylor, Edward Burnett (1871). *Primitive culture: Researches into the development of mythology, philosophy, religion, art, and custom*. Vol. 2. J. Murray.

Villeval, Marie Claire (2020). "Public goods, norms and cooperation". *Handbook of Experimental Game Theory*. Edward Elgar Publishing, pp. 184–212.

Walker, James M et al. (2000). "Collective choice in the commons: Experimental results on proposed allocation rules and votes". *The Economic Journal* 110.460, pp. 212–234.

Weitzman, Martin L. (1979). "Optimal search for the best alternative". *Econometrica* 47.3, pp. 641–654.

Yao, Yi-Ching (2006). "Some results on the Gittins index for a normal reward process". *IMS Lecture Notes Monograph Series* 50, pp. 284–293.

A Database of traditional beliefs

A.1 Literature search

The following search tools were utilized to identify cases for inclusion in the database:

- Google Scholar
- AnthroSource
- Historical Abstracts
- America History and Life
- Ethnic News Watch
- Latin American Newspaper
- Latin American Digital Beat
- SAGE Research Methods
- Atla religion database
- Mongabay

Google Scholar provided the most results.

Next, keyword strategies were developed. The following list of keywords was paired across keyword types (beliefs and resources). Keywords were added as cases were identified.

- Beliefs keywords:
  - Traditional beliefs
  - Traditional knowledge
  - Indigenous knowledge
  - Superstition
  - Indigenous people
  - Natural beliefs
  - Ethnoecology
  - Spirit

  - Taboo

- Resource keywords:
  - Conservation
  - Resource management
  - Natural resources
  - Habitat
  - Hunting
  - Wildlife
  - Land use
  - Management
  - Environmental

When a paper was identified, its list of references was searched for additional relevant papers. If the full references cited were unavailable, direct searches were done for the referenced community, region, or superstition.

Papers were selected for inclusion if they contained a rule regarding a natural resource tied to cultural/spiritual beliefs (i.e., management decisions did not come from formal or informal governance mechanisms). Papers with direct or indirect references to a supernatural punishment for rule defection were eligible. Papers that discussed resource management but had rules and/or punishments defined through formal (laws or regulations) or community-based (shunning, exclusion from community activities) institutions were not included in the review unless they also included supernatural punishment. Papers that included a rule connected to a spiritually important resource and were governed by supernatural belief but did not include discussion of punishments were also included in the final set of papers, though punishment information is missing in these cases.

The review did not restrict to any geographies ex ante. A clear pattern of geography emerged, with particular countries / regions returning more case studies than others. These include Ghana, Nigeria, and the countries and territories that make up Polynesia. Major gaps are likely to persist, such as Aboriginal beliefs in Australia, which were not returned in the literature search.

Given this process for constructing the cases around which the database is built, errors of omission are likely. Consequently, the frequencies reported in the main text should be interpreted as conditional on the assembled cases, not as representative of all relevant belief systems.

A.2 Creation of database

The database is organized such that each observation represents a communally held belief, defined as a unique combination of community, resource, and rule. A single case study may contribute multiple observations if it documents beliefs across several communities or describes distinct rules applied to the same resource. The belief-level dataset includes (1) the continent, country, region, and community in which the belief was documented, (2) the resource involved, (3) the rule or prohibition, (4) the rule category (tax, restriction, or prohibition), (5) the type of punishment, (6) whether defection affects the resource stock, (7) whether punishment is related to the stock, and (8) the citation of the source from which the entry was drawn. An excerpt from the database is provided in Table A.5. The full database is available for download here.

The belief-level dataset was constructed by coding qualitative details from each case. For example, a case documenting that the Hupa people of Northwestern California believed that failure to perform the first salmon ritual would result in illness or death is coded as: country = United States, resource = fish/game, rule = ritual (tax), punishment = illness/death.

Simplified variables for resource rule, rule category, and type of punishment were used to classify case studies by their shared traits. The definitions in Table 1 were used for classification.

A.3 Coding rules for categorical variables

Resource type. Each belief is assigned to one of three resource categories based on the primary target of the rule:

- *Habitat*: rules applying to a place rather than a species — sacred forests, protected zones, bodies of water, or other spatially defined areas. This is the appropriate category when the restriction governs entry into or exploitation of a designated area, regardless of what is extracted there.
- *Wildlife*: rules applying to specific animal or plant species that are regarded as sacred, totemic, or otherwise spiritually significant, and are not primarily targeted for consumption. Includes totem animals, sacred birds, and species associated with ancestral spirits.
- *Fish/Game*: rules applying to species or resources that are targeted for consumption or use, including fish, game animals, timber, and cultivated plants. When a rule applies to both a habitat and a harvestable resource, classification follows the primary focus of the restriction.

Rule category. Rules are assigned to one of three broad categories that map onto modern resource management instruments:

- *Tax*: rules that prescribe a required action before or after resource use, rather than limiting use itself. Includes rituals, offerings, and worship requirements. Analogous to a tax in that use is permitted but carries an additional cost.
- *Restriction*: rules that limit but do not prohibit resource use, typically by constraining the timing, quantity, method, or identity of who may extract. Includes seasonal closures, harvest limits, and role-based access rules.
- *Prohibition*: rules that categorically forbid extraction, entry, or use. Includes sacred species that may not be killed, zones that may not be entered, and resources that may not be harvested under any circumstances.

Rule. Within each category, the specific rule is coded from the following list, selected to capture the primary behavioral constraint:

- Tax: ritual, offering, worship
- Restriction: exploitation restriction, hunting restriction, entry restriction
- Prohibition: no exploitation, no killing, no entry

When a case describes multiple rules, each is coded as a separate observation if they apply to distinct resources or communities. When a single rule description contains elements of multiple categories (e.g., a restriction that becomes a prohibition during certain seasons), the dominant or more restrictive characterization is used.

Punishment type. Punishments are assigned to the following categories based on the primary described consequence of violation:

- *Illness/death (self or other)*: physical harm to the defector, their family, or others, including disease, injury, and death.
- *Loss of divine protection*: withdrawal of supernatural protection that the belief system otherwise provides, leaving the community or individual vulnerable to harm.
- *Loss of crops/livestock*: damage to agricultural production or domestic animals unrelated to the primary resource.

- *Loss of hunt*: failure of future hunting or fishing, including animals avoiding the defector or being unavailable.
- *Natural disasters*: severe weather, drought, flood, famine, or other large-scale environmental disruptions.
- *Bad luck*: unspecified misfortune described as punishment but without further detail.
- *Psychological harm*: mental illness, family conflict, disorientation, or spirit possession.
- *Other*: punishments not fitting the above categories.

When the source does not describe the punishment or provides insufficient detail for classification, the entry is coded as *no info*. Punishment is coded as *unclear* when the source is ambiguous despite sufficient detail being present.

A.4 Punishment characteristics

The punishment was further classified by its severity, specificity and reason why learning is hard.

The severity score is ranked from high to low, as follows:

1. Low: punishment causes minor physical harm (a rash on the body); punishment creates conflict that is unrelated to physical harm of anyone / anything (bad luck, quarrels among family).
2. Medium: punishment causes more severe physical harm (illness); punishment interferes with access to food sources related or unrelated to the resource (loss of hunt, destruction of harvest, rivers drying up); punishment may put defector in potential harm (getting lost in the forest, encounter with spirits).
3. High: death of any kind; major calamities (drought, famine, war); major disruptions in food source (resource never returning again).

The variable 'why learning is hard' has five categories, which are not mutually exclusive:

1. Time: If it is unclear when the punishment will occur after the defector has broken a rule, then the rule is unspecific in time.
2. Kind: if it is unclear in what the punishment will be (e.g. bad luck, illness, mental disorientation), then the rule is unspecific in kind.

3. Who: If it is unclear to whom the punishment will happen to (the defector, loved ones of the defectors, the community, unrelated people), then the rule is not specific in who will be punished.

4. Random: Unclear if the defection would lead to a punishment at all (gods may be angered and if they are angered, punishment occurs)

5. Punishment rule does not make learning difficult: This applies to punishment rules that are specific in all the areas listed above.

These are further summarized into a specificity score. A high score means that the punishment rule does not make learning difficult. A medium score means that the punishment has one specificity problem. A low score means that the punishment has two or more specificity problem.

A.5 Database excerpt

Table A.1: Database Overview

Country	Region	Resource	Rule	Rule Category	Punishment	Citation
Ethiopia	Gamo Highland	habitat	exploitation restriction	restriction	loss of crops/livestock	(Assefa, 2017)
Ghana	Tallensi-Nabdam	fish / game	hunting restriction	restriction	illness/death (self or other)	(Barre et al., 2009)
Madagascar	Androy region	habitat	offering	tax	illness/death (self or other)	(Tengö et al., 2007)
Mozambique	Choa Highlands, Barue District	habitat	exploitation restriction	restriction	other	(Virtanen, 2002)
Nigeria	Ndokwa East Local Government Area	fish / game	hunting restriction	restriction	illness/death (self or other)	(Rim-Rukeh et al., 2013)
South Africa	Vhembe District	habitat	entry restriction	restriction	bad luck	(Sinthumule, 2023)
Zimbabwe	Makonde District	wildlife	no killing	prohibition	illness/death (self or other)	(Sigauke et al. 2015)
Bhutan	Unakha Dzongkhag	habitat	offering	tax	natural disasters	(Basnet & Dendup, 2020)
India	Maharashtra	wildlife	worship	tax	loss of crops/livestock	(Nair et al., 2021)
Indonesia	West Sumba	habitat	no exploitation	prohibition	illness/death (self or other)	(Fowler, 2003)
Taiwan	Orchid Island	fish / game	ritual	tax	no info	(Tang & Tang, 2009)
Thailand	Lower Songkhram River Basin	habitat	no entry	prohibition	psychological harm	(Chunhabunyatip et al., 2018)
Canada	British Columbia	fish / game	ritual	tax	loss of hunt	(Mohs, 1985)
Mexico	Yucatan Peninsula	fish / game	offering	tax	loss of hunt	(Villa Rojas, 1945)
United States	Northwestern California	fish / game	ritual	tax	illness/death (self or other)	(Swezey & Heizer, 1977)
Cook Islands	Mangaia	habitat	exploitation restriction	restriction	other	(Dixon, 2016)
French Polynesia	Maiao	wildlife	no killing	prohibition	illness/death (self or other)	(Teuira Henry, 1928 (2004))
Wallis	Uvea	habitat	entry restriction	restriction	natural disasters	(Hélène Guiot, 1998)
Brazil	Brazillian Amazon	fish / game	hunting restriction	restriction	illness/death (self or other)	(Costa Neto et al., 2023)
Peru	Peruvian Amazon	habitat	no exploitation	prohibition	illness/death (self or other)	(Fabiano et al., 2021)

Table A.2: Comparison with existing datasets

Database Name	Description	Geographic Coverage	Coded?	Resource Management	Natural Resource	Rule	Punishment
Pulotu	Cross-cultural database focused on traditional religious beliefs and practices among Austronesian cultures, with coded variables across three time periods.	Austronesian	Yes	Yes	No	No	No
eHRAF World Cultures	Digital ethnographic database providing full-text, key word searchable paragraphs of case studies on global cultures, also tagged by subject codes.	Global	No	Yes	Partial*	Partial*	Partial*
Mtyhology and Religion Database	Collection of myths and folktales organized by motif codes with narrative summaries and maps	Global	No	No	No	No	No
D-PLACE	Database of environmental, linguistic, and cultural data, coded with binary or ordinal variables from individual case studies.	Global	Yes	No	No	No	No
Hunter-gather Religion Database	Compiled data on early religious systems among 84 hunter-gatherer societies, including coded variables and qualitative summaries.	Global	Yes	No	No	No	No

Continued on next page

Database Name	Description	Geographic Coverage	Coded?	Resource Management	Natural Resource	Rule	Punishment
World Religion and Spiriuality Project	Online directory of religious movements and spiritual traditions, with narrative case study profiles covering beliefs, practices, and institutional structures.	Global	No	No	No	No	No
The Database of Religious History	Database containing qualitative data on religious beliefs and practices across world history.	Global	No	Partial*	Partial*	Partial*	No
Sheshat Global History Databank	Historical database combining coded variables and narrative summary on governance, religion, and social structure across world societies.	Global	Yes	No	No	No	No

** Partial means that not all case studies have the same information available regarding the variable.*

Database references

Aguiar, Leonardo Matheus Pereira, Eraldo Medeiros Costa Neto, and Dídac Santos Fita (Nov. 2023). "Caipora and the conservation of natural resources in tropical forests in the South Recôncavo region, Bahia State, Northeast Brazil". *Ethnobiology and Conservation* 12.

Akpabio, Emmanuel M. (Apr. 2011). "Water and people: Perception and management practices in Akwa Ibom State, Nigeria". *Society & Natural Resources* 24.6, pp. 584–596.

Aniah, Philip (2014). "Behind the myth: Indigenous knowledge and belief systems in natural resource conservation in northeast Ghana". *International Journal of Environmental Protection and Policy* 2.3, p. 104.

Asante, Eric Appau, Stephen Ababio, and Kwadwo Boakye Boadu (Jan. 2017). "The use of indigenous cultural practices by the Ashantis for the conservation of forests in Ghana". *SAGE Open* 7.1.

Assefa, Engdawork and Bork Hans-Rudolf (Sept. 2017). "Indigenous resource management practices in the Gamo Highland of Ethiopia: Challenges and prospects for sustainable resource management". *Sustainability Science* 12.5, pp. 695–709.

Baaweh, Louis, Issah Baddianaah, and Bernard Nuoleyeng Baatuuwie (Mar. 2022). "Traditional knowledge and practices in natural resource conservation: A study of the Zukpiri Community Resource Management Area, Ghana". *International Journal of Rural Management*.

Bambridge, Tamatoa (Mar. 2016). "The law of rahui in the Society Islands". *The Rahui: Legal Pluralism in Polynesian Traditional Management of Resources and Territories*. 1st. ANU Press.

Barre, Rita Yembilah, Miriam Grant, and Dianne Draper (Feb. 2009). "The role of taboos in conservation of sacred groves in Ghana's Tallensi-Nabdam district". *Social & Cultural Geography* 10.1, pp. 25–39.

Basnet, Janga Bdr. and Tashi Dendup (2020). "Role of traditional beliefs and cultural myths in environmental conservation: A case study from Kabjisa, Punakha, Bhutan". *SSRN Electronic Journal*.

Bhagwat, Shonil A. and Claudia Rutte (Dec. 2006). "Sacred groves: Potential for biodiversity management". *Frontiers in Ecology and the Environment* 4.10, pp. 519–524.

Chambers, Charlotte (Mar. 2016). "European contact and systems of governance on Tongareva". *The Rahui: Legal Pluralism in Polynesian Traditional Management of Resources and Territories*. 1st. ANU Press.

Chunhabunyatip, Prompassorn et al. (Aug. 2018). "Influence of indigenous spiritual beliefs on natural resource management and ecological conservation in Thailand". *Sustainability* 10.8, p. 2842.

Conte, Eric (Mar. 2016). "Technical exploitation and 'ritual' management of resources in Napuka and Tepoto (Tuamotu Archipelago)". *The Rahui: Legal Pluralism in Polynesian Traditional Management of Resources and Territories*. 1st. ANU Press.

Costa Neto, Eraldo Medeiros, Dídac Santos-Fita, and Leonardo Matheus Pereira Aguiar (Mar. 2023). "Curupira e Caipora: O papel dos seres elementais como guardiões da natureza". *Boletim Do Museu Paraense Emílio Goeldi. Ciências Humanas* 18, e20210095.

Cox, Michael, Sergio Villamayor-Tomas, and Yasha Hartberg (Feb. 2014). "The role of religion in community-based natural resource management". *World Development* 54, pp. 46–55.

Devi Khumbongmayum, Ashalata, M. L. Khan, and R. S. Tripathi (June 2005). "Sacred groves of Manipur, northeast India:biodiversity value, status and strategies for their conservation". *Biodiversity and Conservation* 14.7, pp. 1541–1582.

Dixon, Rod (Mar. 2016). "I uta i tai — a preliminary account of ra'ui on Mangaia, Cook Islands". *The Rahui: Legal Pluralism in Polynesian Traditional Management of Resources and Territories*. 1st. ANU Press.

Fabiano, Emanuele, Christopher Schulz, and Manuel Martín Brañas (Jan. 2021). "Wetland spirits and indigenous knowledge: Implications for the conservation of wetlands in the Peruvian Amazon". *Current Research in Environmental Sustainability* 3, p. 100107.

Ferronato, Bruno O. and Giselle Cruzado (Dec. 2013). "Uses, beliefs, and conservation of turtles by Ashaninka indigenous people, central Peru". *Chelonian Conservation and Biology* 12.2, pp. 308–313.

Filous, Alexander et al. (Jan. 2021). "Fisheries science and marine education catalyze the renaissance of traditional management (rahui) to improve an artisanal fishery in French Polynesia". *Marine Policy* 123, p. 104291.

Fowler, Cynthia T. (2003). *The ecological implications of ancestral religion and reciprocal exchange in a sacred forest in Karendi*.

Ghasarian, Christian (2016). "Protection of natural resources through a sacred prohibition: The rahui on Rapa Iti". *The Rahui: Legal Pluralism in Polynesian Traditional Management of Resources and Territories*. ANU Press, pp. 139–156.

Golden, Christopher D. and Jean Comaroff (2015). "The human health and conservation relevance of food taboos in northeastern Madagascar". *Ecology and Society* 20.2, art42.

Gottesfeld, Leslie M. Johnson (1994). "Conservation, territory, and traditional beliefs: An analysis of Gitksan and Wet'suwet'en subsistence, northwest British Columbia, Canada". *Human Ecology* 22.4, pp. 443–465.

Guiot, Hélène (1998). "Forêt taboue et représentations de l'environnement à'Uvea (Wallis). Approche ethno-archéologique". *Journal de la Société des Océanistes* 107.2, pp. 179–198.

Hajdú, Péter (1968). *The Samoyed peoples and languages*. 2nd. Vol. 14. Uralic and Altaic Series. Indiana University Publications.

Kandari, Laxman Singh et al. (Dec. 2014). "Conservation and management of sacred groves, myths and beliefs of tribal communities: A case study from north-India". *Environmental Systems Research* 3.1, p. 16.

Kondo, Masaki (Jan. 1991). "The formation of sacred places as a factor of environmental preservation: The case of Setonaikai (Inland Sea) in Japan". *Marine Pollution Bulletin*. Environmental Management and Appropriate Use of Enclosed Coastal Seas 23, pp. 649–652.

Luzar, Jeffrey B., Kirsten M. Silvius, and Jose M. V. Fragoso (Dec. 2012). "Church affiliation and meat taboos in indigenous communities of Guyanese Amazonia". *Human Ecology* 40.6, pp. 833–845.

Mair, Lucy Philip (1934). *An African people in the twentieth century*. Routledge & Sons.

Masalu, D.C.P, M.S. Shalli, and R.A. Kitula (2010). *Customs and taboos: The role of indigenous knowledge in the management of fish stocks and coral reefs in Tanzania*. Tech. rep. Coral Reef Targeted Research and Capacity Building for Management Program, p. 64.

Mmari, Mercy W. et al. (Nov. 2017). "Traditions, beliefs and indigenous technologies in connection with the edible longhorn grasshopper Ruspolia differens (Serville 1838) in Tanzania". *Journal of Ethnobiology and Ethnomedicine* 13.1, p. 60.

Mohs, Gordon (1985). *Spiritual sites, ethnic significance and native spirituality: The heritage and heritage sites of the Sto:lo Indians of British Columbia*.

Mukul, Sharif Ahmed, A.Z.M. Manzoor Rashid, and Mohammad Belal Uddin (June 2012). "The role of spiritual beliefs in conserving wildlife species in religious shrines of Bangladesh". *Biodiversity* 13.2, pp. 108–114.

Nair, Ramya et al. (June 2021). "Sharing spaces and entanglements with big cats: The Warli and their Waghoba in Maharashtra, India". *Frontiers in Conservation Science* 2.

Negi, Chandra Singh (May 2010a). "The institution of taboo and local resource management and conservation surrounding sacred natural sites in Uttarakhand, central Himalaya". *International Journal of Biodiversity and Conservation*.

— (Aug. 2010b). "Traditional culture and biodiversity conservation: Examples from Uttarakhand, central Himalaya". *Mountain Research and Development* 30.3, pp. 259–265.

Nijhawan, Sahil and Achili Mihu (July 2020). "Relations of blood: Hunting taboos and wildlife conservation in the Idu Mishmi of northeast India". *Journal of Ethnobiology* 40.2, pp. 149–166.

Ottino-Garanger, Pierre et al. (Mar. 2016). "Tapu and kahui in the Marquesas". *The Rahui: Legal Pluralism in Polynesian Traditional Management of Resources and Territories*. 1st. ANU Press.

Rim-Rukeh, A, G Irerhievwie, and I E Agbozu (2013). "Traditional beliefs and conservation of natural resources: Evidences from selected communities in Delta State, Nigeria". *International Journal of Biodiversity and Conservation* 5.7, pp. 426–432.

Ruru, Jacinta and Nicola Wheen (Mar. 2016). "Providing for rāhui in the law of Aotearoa New Zealand". *The Rahui: Legal Pluralism in Polynesian Traditional Management of Resources and Territories*. 1st. ANU Press.

Ryley, Kathleen (2013). *Fêtes traditionnelles et festivals glocalisés aux Marquises: Utilisation des systèmes alimentaires syncrétiques pour forger des identités hybrides*.

Santos-Fita, Dídac et al. (Sept. 2015). "Symbolism and ritual practices related to hunting in Maya communities from central Quintana Roo, Mexico". *Journal of Ethnobiology and Ethnomedicine* 11.1, p. 71.

Sigauke, E. et al. (2015). "African traditional cultural values and beliefs: A driving force to natural resource management: A study of Makonde District, Mashonaland West Province, Zimbabwe". *Journal of Culture, Society and Development* 9.0, p. 78.

Sinthumule, Ndidzulafhi Innocent (Aug. 2024). "Traditional taboos: Informal and invisible protection of remaining patches of forest in Vhembe District in Limpopo, South Africa". *Frontiers in Conservation Science* 5.

Swezey, Sean L. and Robert F. Heizer (July 1977). "Ritual management of salmonid fish resources in California". *The Journal of California Anthropology* 4.1.

Tengö, Maria et al. (Dec. 2007). "Taboos and forest governance: Informal protection of a hotspot dry forest in southern Madagascar". *AMBIO: a Journal of the Human Environment* 36.8, pp. 683–691.

Teuira, Henry (1928). *Tahiti aux temps anciens*. Bishop Museum D'Honolulu.

Thorax, Lorin (Mar. 2016). "From traditional to modern management in Fakarava". *The Rahui: Legal Pluralism in Polynesian Traditional Management of Resources and Territories*. 1st. ANU Press.

Torrente, Frédéric (Mar. 2016). "Ancient magic and religious trends of the rāhui on the atoll of Anaa, Tuamotu". *The Rahui: Legal Pluralism in Polynesian Traditional Management of Resources and Territories*. 1st. ANU Press.

Villa Rojas, Alfonso (1945). "The maya of east central Quintana Roo". *Carnegie Institution of Washington*.

Virtanen, P. (May 2002). "The role of customary institutions in the conservation of biodiversity: Sacred forests in Mozambique". *Environmental Values* 11.2, pp. 227–241.

Watts, Joseph et al. (Sept. 2015). "Pulotu: Database of Austronesian supernatural beliefs and practices". *PLOS ONE* 10.9. Ed. by Michiel van Elk.

Westerman, Diedrich H. and Frieda Schütze (1921). *The Kpelle: A negro tribe in Liberia*. Vandenhoeck & Ruprecht ; J. C. Hinrichs'sche Buchhandlung.

B Proofs of comparative statics in Table 3

By Proposition 6, when all other agents comply, agent i complies if and only if

$$m_{it} \geq T_i \equiv \mu^c + B_i(\eta_{it}, \nu_i, \delta). \tag{B.1}$$

The threshold T_i has two components. The static threshold μ^c depends only on N, α and δ; in computing its derivatives we allow $\bar{E}$, $\hat{E}$ and E_{-i} to adjust with the parameters, so that the Rāhui is always set at the cooperative level of Proposition 2. The exploration premium B_i depends only on η_{it}, ν_i and δ (we showed this in Proposition 7(a)). For any parameter θ, we say that *compliance is increasing in* θ if the compliance threshold T_i is decreasing in θ. This would imply that the set of beliefs for which agent i complies expands. Compliance is *decreasing in* θ if T_i is increasing in θ.

It is convenient to define $a \equiv \alpha\delta \in (0, 1)$. Because $1 - \frac{N-1}{N}(1-a) = \frac{1+a(N-1)}{N}$, Equation 16 simplifies to

$$\mu^c = \frac{\log\left(1 + a(N-1)\right) - a \log N}{1-a}. \tag{B.2}$$

1. Number of Agents (N)

Claim: Compliance decreases in N ($\downarrow$).

Proof. The parameter N enters T_i only through μ^c. Taking the derivative of Equation B.2, we obtain:

$$\frac{d\mu^c}{dN} = \frac{a}{1-a}\left(\frac{1}{1+a(N-1)} - \frac{1}{N}\right) = \frac{a(N-1)}{N\left(1+a(N-1)\right)} > 0 \quad \text{for } N > 1.$$

T_i is increasing in N, and thus, compliance decreases in N. Intuitively, as N increases, the cooperative effort for each agent $\bar{E} = \frac{1-\alpha\delta}{N}$ shrinks, while the payoff from violating

remains comparatively large. This drives the increase in μ^c because it takes beliefs in higher punishment to ensure compliance with Rāhui. □

2. Inverse Growth Rate (α)

Claim: Compliance decreases in α ($\downarrow$).

Proof. The parameter α enters T_i only through μ^c, and only through $a = \alpha\delta$. Differentiating Equation B.2 with respect to a and multiplying by $\partial a/\partial\alpha = \delta$ gives

$$\frac{d\mu^c}{d\alpha} = \frac{\delta\, g(a)}{(1-a)^2}, \quad \text{where} \quad g(a) \equiv \frac{(1-a)(N-1)}{1+a(N-1)} + \log\left(\frac{1+a(N-1)}{N}\right).$$

To sign g, note that $g(1) = 0$ and

$$g'(a) = -\frac{(N-1)^2(1-a)}{\left(1+a(N-1)\right)^2} < 0 \quad \text{for } a \in (0,1),\ N > 1.$$

Thus $g(a) > 0$ for all $a \in (0,1)$, so $d\mu^c/d\alpha > 0$, and T_i is increasing in α, so compliance decreases in α. Intuitively, a higher α implies slower resource recovery, which raises the scarcity rent and thus the temptation to violate. □

3. Discount Factor (δ)

Claim: Compliance decreases in δ ($\downarrow$).

Proof. The discount factor enters T_i through both μ^c and B_i so care must be taken to analyze both margins.

We begin with μ^c. Changes in δ affect $\bar{E}$, $\hat{E}$ and E_{-i}, and therefore the payoffs K^C and K^V of both arms. But the Normal Gittins index has a location invariance property that implies that these changes affect the compliance threshold only through $\mu^c = K^V - K^C$. By Equation B.2, μ^c depends on δ only through $a = \alpha\delta$, so by the α comparative static argument above,

$$\frac{d\mu^c}{d\delta} = \frac{\alpha\, g(a)}{(1-a)^2} > 0.$$

Turning to B_i, the exploration incentives are also affected by δ. $B_i(\eta_{it}, \nu_i, \delta)$ is the Gittins index of a normal arm with zero prior mean, whose prior and signal precisions do not depend on δ. It depends on δ only through patience, and by Proposition 7(d) it is increasing in δ.

Since both components of T_i are increasing in δ, compliance decreases in δ. □

4. Mean Punishment Belief (m_{it})

Claim: Compliance increases in m_{it} ($\uparrow$).

Proof. The mean belief m_{it} enters only the left-hand side of Equation B.1 (T_i is independent of m_{it} by Proposition 7(a)). Thus if agent i complies at m_{it}, she complies at any $m'_{it} > m_{it}$. Intuitively, a larger m_{it} reflects the agent's belief in stronger punishment, which increases compliance. □

5. Belief Precision (η_{it})

Claim: Compliance increases in η_{it} ($\uparrow$).

Proof. The precision of the prior enters T_i only through B_i, which is decreasing in η_{it} by Proposition 7(b). Thus, T_i is decreasing in η_{it}. Intuitively, a more confident agent expects to learn less from a violation, so is less likely to violate to learn. □

6. Signal Precision (ν_i)

Claim: Compliance (weakly) decreases in ν_i ($\downarrow$).

Proof. The signal precision enters T_i only through B_i, which is nondecreasing in ν_i by Proposition 7(c). Thus T_i is nondecreasing in ν_i. Intuitively, when background noise is low, a single violation yields a clear signal about the true punishment, which raises the value of experimenting by violating the Rāhui. □