

Chakravarty, Surajeet; Kelsey, David; Teitelbaum, Joshua C.

Working Paper

Operationalizing reverse Bayesianism

CeDEx Discussion Paper Series, No. 2020-18

Provided in Cooperation with:

The University of Nottingham, Centre for Decision Research and Experimental Economics (CeDEx)

Suggested Citation: Chakravarty, Surajeet; Kelsey, David; Teitelbaum, Joshua C. (2020) : Operationalizing reverse Bayesianism, CeDEx Discussion Paper Series, No. 2020-18, The University of Nottingham, Centre for Decision Research and Experimental Economics (CeDEx), Nottingham

This Version is available at:

<https://hdl.handle.net/10419/248290>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



University of
Nottingham
UK | CHINA | MALAYSIA

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

Discussion Paper No. 2020-18

Surajeet Chakravarty
David Kelsey
Joshua C. Teitelbaum

November 2020

**Operationalizing Reverse
Bayesianism**

CeDEx Discussion Paper Series

ISSN 1749 - 3293



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit <http://www.nottingham.ac.uk/cedex> for more information about the Centre or contact

Suzanne Robey
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 14763
suzanne.robey@nottingham.ac.uk

The full list of CeDEX Discussion Papers is available at

<http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx>

Operationalizing Reverse Bayesianism*

Surajeet Chakravarty

University of Exeter

David Kelsey

University of Exeter and
University of Nottingham

Joshua C. Teitelbaum

Georgetown University

Draft: May 14, 2020

Abstract

Karni and Vierø (2013) propose a model of belief revision under growing awareness—*reverse Bayesianism*—which posits that as a person becomes aware of new acts, consequences, or act-consequence links, she revises her beliefs over an expanded state space in a way that preserves the relative likelihoods of events in the original state space. A key limitation of the model is that reverse Bayesianism alone does not fully determine the revised probability distribution. We provide an assumption—*act independence*—that imposes additional restrictions on reverse Bayesian belief revision. We show that under act independence, knowledge of the probabilities of new events in the expanded state space is sufficient to fully determine the revised probability distribution in each case of growing awareness. We thereby operationalize the reverse Bayesian model for applications. To illustrate how act independence operationalizes reverse Bayesianism, we consider the law and economics problem of optimal safety regulation.

JEL codes: D83, K23.

Keywords: act independence, reverse Bayesianism, safety regulation, unawareness.

*Corresponding author: Joshua C. Teitelbaum, Georgetown University, 600 New Jersey Avenue NW, Washington, DC 20001 (jct48@georgetown.edu). We thank the editor and the referees for their helpful comments. We also thank Jonah Gelbach, Tim Friehe, Simon Grant, John Mikhail, David Reinstein, Steve Salop, Steve Shavell, Larry Solum, Marie-Louise Vierø, and participants at the following conferences and seminars: the Law and Ambiguity Workshop at the University of Exeter, the Decisions, Markets, and Networks conference at Cornell University, the 34th Annual Conference of the European Association of Law and Economics, Bielefeld University, Towson University, Université Paris Nanterre, Universität Bonn, the 2018 Annual Meeting of the AALS Section on Law and Economics, the 28th Annual Meeting of the American Law and Economics Association, the 2018 Foundations of Utility and Risk Conference, the 8th Annual Law and Economics Theory Conference, Australian National University, and the University of Auckland.

1 Introduction

Overview. Economists traditionally model choice under uncertainty according to Savage’s theory of subjective expected utility (Savage, 1954). Savage’s theory posits a space of mutually exclusive and collectively exhaustive states of the world, representing all possible resolutions of uncertainty. It assumes that when a person chooses an act, although she is uncertain about the true state of the world and therefore about the consequences of her chosen act, she nevertheless has complete knowledge of the state space—she knows all the possible acts and all the possible consequences of each and every act.

In reality, however, a person often does not have complete knowledge of the state space. This is known as *unawareness*. A person may be unaware of some acts, some consequences, or that a known act can cause a known consequence. An extreme example of the latter is that no one was aware that supporting anti-Soviet fighters in Afghanistan in the 1980s could lead to the destruction of the World Trade Center in 2001. Unawareness creates the possibility of *growing awareness*—the expansion of the state space when a person discovers a new act, consequence, or act-consequence link. Examples include the discovery of a new product or technology (new act), the discovery of a new disease or injury (new consequence), or the discovery that a known product can cause a known injury (new act-consequence link).

"Unawareness refers to the lack of conception rather than the lack of information" (Schipper, 2014a,b). There is a fundamental difference between not knowing the state of the world (lack of information) and not knowing that a state of the world is possible (lack of conception). The Savage model allows the state space to contract with the arrival of information and is consistent with Bayesian updating of beliefs. It however does not admit unawareness and cannot accommodate growing awareness (Dekel et al., 1998a,b).

In a pioneering article, Karni and Vierø (2013) propose a model of belief revision under growing awareness called *reverse Bayesianism*. Reverse Bayesianism posits that as a person becomes aware of a new act, consequence, or act-consequence link, she revises her beliefs in a way that preserves the relative likelihoods of events in the original state space. More

specifically, the model postulates that (i) in the case of a new act or consequence, probability mass shifts proportionally away from the events in the original state space to the new events in the expanded state space, and (ii) in the case of a new act-consequence link, null events in the original state space become non-null, and probability mass shifts proportionally away from the original non-null events to the original null events that become non-null.

The reverse Bayesian model has (at least) two features that make it attractive to economists who wish to incorporate unawareness and growing awareness into applications. The first is transparency. Karni and Vierø (2013) provide an axiomatic foundation for the model, so one can judge the theory by the axioms.¹ The second attractive feature of the model is its accessibility. The model is built upon a choice-theoretic framework that is well known to economists (subjective expected utility theory), and the upshot is a belief revision theory that mirrors the familiar process of Bayesian updating.²

A key limitation of the reverse Bayesian model, however, is that reverse Bayesianism alone does not fully determine the revised probability distribution over the expanded state space. This is because reverse Bayesianism implies restrictions on the revised probabilities of non-null events in the original state space, but not on the probabilities of new events in the expanded state space. To borrow a term from the econometrics literature, reverse Bayesianism only partially identifies the revised probability distribution.³

In this paper, we provide an assumption—*act independence*—that implies additional restrictions on the revised probability distribution in the reverse Bayesian model. Essentially, act independence requires that acts are independent experiments. We show that under act independence, knowledge of the probabilities of new events in the expanded state space is

¹The key axioms of the model are the “consistency” axioms, which essentially require that preferences conditional on the original state of awareness are not altered by growing awareness.

²This feature prompts Dominiak and Tserenjigmid (2018, p. 3) to describe Karni and Vierø’s (2013) reverse Bayesian model as “elegant.”

³Karni and Vierø (2013, p. 2805) highlight this feature of reverse Bayesianism in their concluding remarks: “The model presented in this article predicts that, as awareness grows and the state space expands, the relative likelihoods of events in the original state space remain unchanged. The model is silent about the absolute levels of these probabilities. In other words, our theory does not predict the probability of the new events in the expanded state space.”

sufficient to fully determine the revised probability distribution over the expanded state space in each case of growing awareness (a new act, consequence, or act-consequence link). In this way, act independence makes the reverse Bayesian model operational for economic applications. This is our main contribution.

After introducing act independence into the reverse Bayesian model, we illustrate how it operationalizes reverse Bayesianism with an application in the field of law and economics. In our application, we study the problem faced by a safety regulator who is tasked with regulating the risky activities of a company. For instance, the regulator could be an environmental agency (such as the U.S. Environmental Protection Agency or the U.K. Environment Agency) or a health agency (such as the U.S. Centers for Disease Control and Prevention or Public Health England). We analyze how the regulator revises her beliefs about the risk of each activity, and in turn the safety standards for each activity, in the wake of growing awareness. We consider each case of growing awareness and then discuss the importance of the act independence assumption for the analysis of each case.

Related literature. The unawareness literature was pioneered by Fagin and Halpern (1988). Other early contributions include Modica and Rustichini (1994, 1999), Dekel et al. (1998b), Halpern (2001), Heifetz et al. (2006), and Halpern and Rêgo (2008). The early papers in the literature generally pursued an epistemic approach or a game-theoretic approach. Surveys of these papers are provided by Schipper (2014b) (which offers a “gentle introduction” to the literature) and Schipper (2015) (which provides an extended review).

Karni and Vierø (2013) are among the pioneers of the choice-theoretic approach (i.e., the state-space approach) to modeling unawareness. Subsequent papers build on their approach. For instance, Grant et al. (2019) invoke their approach to model learning by experimentation in a world with unawareness; Karni and Vierø (2015, 2017) extend their model to the cases where the decision maker is probabilistically sophisticated (but does not necessarily abide by expected utility theory) and where she anticipates her growing awareness; and Dominiak and Tserenjigmid (2018) generalize their model such that the decision maker perceives am-

biguity in the wake of growing awareness.⁴ Karni and Vierø (2013, 2017) and Dominiak and Tserenjigmid (2018) survey the papers that take a choice-theoretic approach.

A handful of papers apply unawareness models to study legal topics. The bulk of these focus on contracts. For example, Board and Chung (2011) argue that asymmetric unawareness provides a justification for the *contra proferentem* doctrine of contract interpretation, which provides that ambiguous terms in a contract should be construed against the drafter; Zhao (2011) argues that unawareness may explain the existence of force majeure clauses in contracts; Grant et al. (2012) study aspects of differential awareness that give rise to contractual disputes; Filiz-Ozbay (2012) posits asymmetric awareness as a reason for the incompleteness of contracts; von Thadden and Zhao (2012, 2014) study the properties of optimal contracts under moral hazard when the agent may be partially unaware of her action space; and Auster (2013) introduces asymmetric unawareness into the canonical moral hazard model and analyzes the properties of the optimal contract.

Structure of the paper. Section 2 presents the reverse Bayesian model. Section 3 introduces the act independence assumption into the model and derives our main results. Section 4 contains our safety regulation application. Section 5 offers concluding remarks. The Appendix collects the proofs of all theorems.

2 Reverse Bayesian Model

The primitives of the reverse Bayesian model are a finite, non-empty set F of feasible *acts* and a finite, non-empty set Z of feasible *consequences*. States are functions from the set of acts to the set of consequences. A state assigns a consequence to each act. The set of all

⁴More specifically, Dominiak and Tserenjigmid (2018) provide a theory of choice under growing awareness in which subjective expected utility preferences (with unawareness) extend to maxmin expected utility preferences (without unawareness). They however leave unexplained how beliefs revise with further information. In our paper the decision maker has subjective expected utility preferences with possible unawareness before and after discovering a new act, consequence, or link and revising her beliefs. This framework could potentially be extended to multiple rounds of discovery and belief revision, if required.

possible states, Z^F , defines the *conceivable state space*. With m acts and n consequences, there are n^m conceivable states.

The decision maker originally conceives the set of acts to be $F = \{f_1, \dots, f_m\}$ and the set of consequences to be $Z = \{z_1, \dots, z_n\}$. The conceivable state space is $Z^F = \{s_1, \dots, s_{n^m}\}$, where each state $s \in Z^F$ is a vector of length m , the i th element of which, s^i , is the consequence $z_j \in Z$ produced by act $f_i \in F$ in that state of the world.

An act-consequence link, or *link*, is a causal relationship between an act and a consequence. The conceivable state space admits all conceivable links. However, the decision maker may perceive one or more links as infeasible, which brings her to nullify the states that admit such link. We refer to these as null states and denote them by $N \subset Z^F$. Taking only the non-null states defines the *feasible state space*, $S \equiv Z^F \setminus N$. There are $\prod_{i=1}^m (n - \nu_i)$ feasible states, where ν_i denotes the number of nullified links involving act f_i .

The decision maker's beliefs are represented by a probability measure p on the conceivable state space, Z^F . The support of p is the feasible state space, S . That is, $p(s) > 0$ for all $s \in S$ and $p(s) = 0$ for all $s \in N$.

The decision maker may initially fail to conceive one or more acts or consequences or to perceive as feasible one or more conceivable links. We refer to such failures of conception or perception as *unawareness*. However, the decision maker may later discover a new act or consequence, which expands both the feasible state space and the conceivable state space, or a new link, which expands the feasible state space but not the conceivable state space.⁵ We refer to such discoveries and expansions as *growing awareness*.

To illustrate, suppose $S = Z^F$ and the decision maker discovers a new consequence, z_{n+1} . Then the set of consequences becomes $\widehat{Z} = Z \cup \{z_{n+1}\}$ and the feasible and conceivable state spaces both expand to $\widehat{S} = \widehat{Z}^F = \{s_1, \dots, s_{(n+1)^m}\}$, where each state remains a vector of length m . Alternatively, suppose the decision maker discovers a new act, f_{m+1} . Then the set of acts becomes $\widehat{F} = F \cup \{f_{m+1}\}$ and the feasible and conceivable state spaces both

⁵To be clear, by “new” we mean “not previously conceived” in the case of acts and consequences, and “previously conceived but perceived as infeasible” in the case of links.

expand to $\widehat{S} = Z^{\widehat{F}} = \{s_1, \dots, s_{n(m+1)}\}$, where each state now is a vector of length $m + 1$. Lastly, suppose $S \subset Z^F$ because (and only because) the decision maker initially perceives as infeasible the link from f_1 to z_n . Discovery of the link from f_1 to z_n does not alter the conceivable state space, but the feasible state space expands to coincide with the conceivable state space: $\widehat{S} = Z^F$. Section 4 contains illustrative depictions of conceivable and feasible state spaces and their expansion due to the discovery of new acts, consequences, and links.

In the wake of growing awareness, the decision maker revises her beliefs in a way that preserves the relative likelihoods of the events in the original feasible state space (the non-null events in the original conceivable state space). In each case of growing awareness, probability mass shifts proportionally away from the events in the original feasible state space to the new events in the expanded feasible state space. In the case of a new act or consequence, the new events in the expanded feasible state space are also new events in the expanded conceivable state space. In the case of a new link, the new events in the expanded feasible state space are the null events in the original conceivable state space that become non-null.

Karni and Vierø (2013) refer to this belief revision process as *reverse Bayesianism*. Let \widehat{p} denote the decision maker's revised beliefs on the expanded feasible state space, \widehat{S} . Formally, reverse Bayesianism implies two restrictions on \widehat{p} : (i) in the case of a new consequence or link, $p(s)/p(t) = \widehat{p}(s)/\widehat{p}(t)$ for all $s, t \in S$; and (ii) in the case of a new act, $p(s)/p(t) = \widehat{p}(E(s))/\widehat{p}(E(t))$ for all $s, t \in S$, where $E(s)$ denotes the event in \widehat{S} that corresponds to state $s \in S$; that is, given a new act f_{m+1} , $E(s) \equiv \{t \in \widehat{S} : t^i = s^i \text{ for all } i \neq m + 1\}$.

3 Act Independence

We add an assumption to the reverse Bayesian model—*act independence*. Let $A_i(z_j) \subset \widehat{S}$ denote the event that f_i yields z_j ; that is, $A_i(z_j) \equiv \{t \in \widehat{S} : t^i = z_j\}$. We refer to events of this type as *act events*. We assume that act events are statistically independent.

Act independence. $A_i(z_j) \perp A_{i'}(z_{j'})$ for all i and i' where $i \neq i'$ and all j and j' .

Act independence implies additional restrictions on the decision maker's revised beliefs, \widehat{p} . Take any event $E \subseteq \widehat{S}$. We can express each state $s = (s^1, \dots, s^m) \in E$ as the intersection of a unique collection of act events in \widehat{S} : $s = \bigcap_i A_i(s^i)$. Act independence implies that $\widehat{p}(s) = \prod_i \widehat{p}(A_i(s^i))$ for all $s \in E$.

Growing awareness—whether it entails a new act, consequence, or link—gives rise to one or more new act events in \widehat{S} . In the remainder of this section we show that, under act independence, knowledge of the probabilities of the new act events in \widehat{S} is sufficient to fully determine \widehat{p} in each case of growing awareness. We start with the case of a new link.

3.1 New Link

Suppose $S \subset Z^F$ and the decision maker discovers a new link from f_l to z_k for some $l \in \{1, \dots, m\}$ and $k \in \{1, \dots, n\}$. Let \widehat{S} denote the expanded feasible state space and \widehat{p} denote the decision maker's revised beliefs on \widehat{S} . Observe that $\widehat{S} = S \cup \Delta$, where $\Delta = A_l(z_k)$ is the newly discovered event that f_l yields z_k . Intuitively, Δ is a copy of any one of the act events $A_l(z_j)$ in S , except that f_l yields z_k (instead of z_j) in every state in Δ . We assume that, by virtue of the discovery, the decision maker learns that f_l yields z_k with probability $\delta > 0$. By definition, $\delta = \widehat{p}(\Delta)$.

For each state $s \in \Delta$, let $L(s) \equiv \{t \in S : t^i = s^i, \forall i \neq l\}$ denote the event in S that corresponds to the state $s \in \Delta$. In other words, $L(s)$ comprises the states in S in which every act (other than f_l) yields the same consequence that it yields in state $s \in \Delta$.

By reverse Bayesianism, the relative likelihoods of the states in S are preserved: $p(s)/p(t) = \widehat{p}(s)/\widehat{p}(t)$ for all $s, t \in S$. By act independence, the probability of each state in Δ equals the product of the probabilities of the act events in \widehat{S} whose intersection defines such state: $\widehat{p}(s) = \prod_{i=1}^m \widehat{p}(A_i(s^i))$ for all $s = (s^1, \dots, s^m) \in \Delta$. It follows that:

Theorem 1. *In the case of a new link involving f_l :*

(i) $\widehat{p}(s) = (1 - \delta)p(s)$ for all $s \in S$; and

(ii) $\widehat{p}(s) = \delta p(L(s))$ for all $s \in \Delta$.

Theorem 1 says that (i) the fraction δ of the probability mass of each state in S is taken away, and that (ii) the total probability mass δ taken away from the states in S is distributed among the states in Δ in proportion to the probability masses of their corresponding events in S . Reverse Bayesianism dictates the first result (how probability mass is shifted away from the states in S), while act independence dictates the second result (how the shifted probability mass is apportioned among the states in Δ). Together, reverse Bayesianism and act independence fully determine the revised probability distribution \hat{p} on \hat{S} .⁶

3.2 New Act

Next, suppose $S \subseteq Z^F$ and the decision maker discovers a new act, f_{m+1} . Again, let \hat{S} denote the expanded feasible state space and \hat{p} denote the decision maker's revised beliefs on \hat{S} . Observe that $\hat{S} = \bigcup_{j=1}^n \Delta_j$, where $\Delta_j = A_{m+1}(z_j)$ is the newly discovered event that f_{m+1} yields z_j . Intuitively, each Δ_j is an augmented copy of S in which f_{m+1} yields z_j in every state. We assume that, by virtue of the discovery, the decision maker learns that f_{m+1} yields z_j with probability $\delta_j > 0$ for all $j = 1, \dots, n$.⁷ Note that $\delta_j = \hat{p}(\Delta_j)$ and $\sum_{j=1}^n \delta_j = 1$.

For each state $s \in S$, let $E(s) \equiv \{t \in \hat{S} : t^i = s^i, \forall i \neq m+1\}$ denote the event in \hat{S} that corresponds to the state $s \in S$. In other words, $E(s)$ comprises the states in \hat{S} in which every act (other than f_{m+1}) yields the same consequence that it yields in state $s \in S$. Observe that $\hat{S} = \bigcup_{s \in S} E(s)$, where $E(s)$ comprises n states, one in which f_{m+1} yields z_1 , one in which f_{m+1} yields z_2 , and so forth. Index the states in each $E(s)$ by $j = 1, \dots, n$, such that $s_j \in E(s)$ is the state in $E(s)$ in which f_{m+1} yields z_j . The connection between the sets of events $\{E(s) : s \in S\}$ and $\{\Delta_j : j = 1, \dots, n\}$, both of which partition \hat{S} , is that Δ_j collects the j th state from each $E(s)$.

By reverse Bayesianism, $p(s)/p(t) = \hat{p}(E(s))/\hat{p}(E(t))$ for all $s, t \in S$. By act independence, $\hat{p}(s) = \prod_{i=1}^{m+1} \hat{p}(A_i(s^i))$ for all $s = (s^1, \dots, s^{m+1}) \in \hat{S}$. It follows that:

⁶Note that p is the Bayesian update of \hat{p} conditional on the event S ; hence the term reverse Bayesianism.

⁷Assuming $\delta_j > 0$ for all $j = 1, \dots, n$ is without loss of generality. We can deal with the case where $\delta_j = 0$ for some j by assuming $\delta_j > 0$ for the first $k < n$ and changing n to k as necessary in the statements below.

Theorem 2. *In the case of a new act f_{m+1} , for all $s \in S$ and corresponding $E(s) \subset \widehat{S}$, $\widehat{p}(s_j) = \delta_j p(s)$ for all $s_j \in E(s)$, $j = 1, \dots, n$.*

Here is the intuition behind Theorem 2. After the discovery of f_{m+1} , each state $s \in S$ is split into n states $s_j \in \widehat{S}$, one for each consequence z_j , $j = 1, \dots, n$. (In state s_j , f_{m+1} yields z_j .) These n states comprise the event $E(s) \subset \widehat{S}$ that corresponds to the state $s \in S$. For each state $s \in S$, reverse Bayesianism dictates that its probability mass is shifted to the corresponding event $E(s) \subset \widehat{S}$, while act independence dictates that the fraction δ_j of the shifted probability mass is apportioned to state $s_j \in E(s)$ for all $j = 1, \dots, n$.

3.3 New Consequence

Last, suppose $S \subseteq Z^F$ and the decision maker discovers a new consequence, z_{n+1} . Once again, let \widehat{S} denote the expanded feasible state space and \widehat{p} denote the decision maker's revised beliefs on \widehat{S} . Observe that $\widehat{S} = S \cup \Delta$, where $\Delta = \bigcup_{i=1}^m A_i(z_{n+1})$ is the union of the newly discovered events that f_i yields z_{n+1} for all $i = 1, \dots, m$. We assume that, by virtue of the discovery, the decision maker learns that f_i yields z_{n+1} with probability $\alpha_i > 0$ for all $i = 1, \dots, m$.⁸ That is, $\alpha_i = \widehat{p}(A_i(z_{n+1}))$. Let $\delta = \widehat{p}(\Delta)$ and note that $1 - \delta = \prod_{i=1}^m (1 - \alpha_i)$.

For each state $s \in \Delta$, let $I(s) \equiv \{i \in \{1, \dots, m\} : s^i = z_{n+1}\}$ denote the indices of the acts that yield z_{n+1} in that state of the world, and let $\bar{I}(s) \equiv \{i \in \{1, \dots, m\} : s^i \neq z_{n+1}\}$ denote the indices of the acts that do not yield z_{n+1} in that state of the world.⁹ In addition, for each $s \in \Delta$, let $C(s) \equiv \{t \in S : t^i = s^i, \forall i \in \bar{I}(s)\}$ denote the event in S that corresponds to $s \in \Delta$ on $\bar{I}(s)$. In other words, $C(s)$ comprises the states in S in which every act (other than the acts that yield z_{n+1}) yields the same consequence that it yields in state $s \in \Delta$.

By reverse Bayesianism, $p(s)/p(t) = \widehat{p}(s)/\widehat{p}(t)$ for all $s, t \in S$. By act independence, $\widehat{p}(s) = \prod_{i=1}^m \widehat{p}(A_i(s^i))$ for all $s = (s^1, \dots, s^m) \in \Delta$. It follows that:

⁸Assuming $\alpha_i > 0$ for all i is without loss of generality. We can deal with the case where $\alpha_i = 0$ for some i by assuming $\alpha_i > 0$ for the first $l < m$ and changing m to l as necessary in the statements below.

⁹Section 4.4 contains examples of the sets $I(s)$ and $\bar{I}(s)$.

Theorem 3. *In the case of a new consequence z_{n+1} :*

(i) $\hat{p}(s) = \left(\prod_{i=1}^m (1 - \alpha_i)\right) p(s)$ for all $s \in S$;

(ii) $\hat{p}(s) = \left(\prod_{i \in I(s)} \alpha_i\right) \left(\prod_{i \in \bar{I}(s)} (1 - \alpha_i)\right) p(C(s))$ for all $s \in \Delta$ such that $I(s) \subset \{1, \dots, m\}$;

(iii) $\hat{p}(s) = \prod_{i=1}^m \alpha_i$ for the $s \in \Delta$ such that $I(s) = \{1, \dots, m\}$.

Theorem 3 is similar to Theorem 1. The first result says that the fraction δ of the probability mass of each state in S is taken away. (Recall that $1 - \delta = \prod_{i=1}^m (1 - \alpha_i)$.) This result is dictated by reverse Bayesianism. The second and third results say how the total probability mass δ taken away from the states in S is distributed among the states in Δ . These results are dictated by act independence. Specifically, the third result says that probability mass $\prod_{i=1}^m \alpha_i$ is apportioned to the one state in which every act results in z_{n+1} (this is a clear implication of act independence), while the second result says that the remaining probability mass, $\delta - \prod_{i=1}^m \alpha_i$, is distributed among the other states in Δ in proportion to the probability masses of their corresponding events in S .

4 Application: Safety Regulation

In this section, we illustrate how act independence operationalizes reverse Bayesianism with an application in the field of law and economics. Unawareness plays an important role in many legal contexts. A prime example is that parties may write incomplete contracts due to unforeseen contingencies.¹⁰ In our application, we study the implications of unawareness and growing awareness for the problem of setting legal standards of conduct in safety regulation.¹¹

Examples of growing awareness that are relevant to safety regulation include, just to name a few, the development of modern day hydraulic fracturing, or “fracking,” in the late 1990s (Gold, 2014) (new act); the discovery of HIV/AIDS in the early 1980s (U.S. Centers for Disease Control and Prevention, 2011) and bovine spongiform encephalopathy, or “mad cow

¹⁰See, e.g., Filiz-Ozbay (2012).

¹¹Although we focus on the context of safety regulation, we note that the problem of setting legal standards of conduct arises in numerous other legal contexts, such as criminal law and tort law, and that our analysis in this section can be adapted to such contexts.

disease,” in the late 1980s (Collee and Bradley, 1997) (new consequences); and the discovery of links between Agent Orange and cancer after the Vietnam War (National Academies of Sciences, Engineering, and Medicine, 2018) and between American football and chronic traumatic encephalopathy in the late 2000s (Lindsley, 2017) (new links).

4.1 Model of Safety Regulation

Consider the problem faced by a safety regulator who is tasked with regulating the activities of a company. For instance, the regulator could be an environmental agency regulating the fracking operations of an energy company or a health agency regulating the screening practices of a blood bank. Suppose that, as far as the regulator is aware, the situation is as follows.¹² The company can engage in two activities, f_1 and f_2 . Each activity has the potential to cause harm to others, though the outcomes of the activities are independent. This is the act independence assumption.¹³ There are two potential degrees of harm, $z_1 = 0$ and $z_2 > 0$. Activity f_i yields harm z_j with probability π_{ij} , where $\pi_{i1} + \pi_{i2} = 1$ for $i = 1, 2$. Thus, activity f_i ’s expected harm is $\pi_{i1}z_1 + \pi_{i2}z_2 = \pi_{i2}z_2$.

Given that $F = \{f_1, f_2\}$ and $Z = \{0, z_2\}$, the conceivable state space, Z^F , comprises four states: $s_1 = (0, 0)$, $s_2 = (0, z_2)$, $s_3 = (z_2, 0)$, and $s_4 = (z_2, z_2)$. Suppose, for the moment, that the regulator perceives both activities as risky (i.e., $\pi_{ij} > 0$ for all i, j). Then the feasible state space is $S = Z^F$. Let $p_k \equiv p(s_k)$, $k = 1, \dots, 4$, denote the regulator’s beliefs on S . We can depict S and p as follows:

p	p_1	p_2	p_3	p_4
$F \setminus S$	s_1	s_2	s_3	s_4
f_1	0	0	z_2	z_2
f_2	0	z_2	0	z_2

¹²Throughout the application, we consider the perspective of the regulator. We assume that the company’s awareness and beliefs always coincide with those of the regulator’s.

¹³While act independence is a reasonable assumption in many settings, there undoubtedly are settings in which it is not. We explore the implications of relaxing the act independence assumption in Section 4.5.

Observe that $\pi_{11} = p_1 + p_2$, $\pi_{12} = p_3 + p_4$, $\pi_{21} = p_1 + p_3$, and $\pi_{22} = p_2 + p_4$. We assume that (i) when the regulator is fully aware, she has correct beliefs about each harm probability, π_{ij} , and (ii) when the regulator is unaware of an act, consequence, or link, her beliefs, although incorrect with respect to the absolute likelihoods of events, are nevertheless correct with respect to the relative likelihoods of non-null events.¹⁴

For each activity f_i , the company can take safety precautions, or *care*, to reduce the activity's expected harm. The cost to the company of taking level of care $x_i \in [0, 1]$ is $c(x_i) = (x_i)^2$. Taking care reduces the activity's expected harm at a constant rate: $h_i(x_i) = \pi_{i2}z_2\tau(x_i)$, where $\tau(x_i) = (1 - x_i)$. We assume that $c(\cdot)$ and $\tau(\cdot)$ are known to the regulator and are the same for all activities.¹⁵

The regulator's problem is to set a standard of care, $\bar{x} = (\bar{x}_1, \bar{x}_2)$ that minimizes the social costs of the company's activities (the costs of care plus the expected harms):¹⁶

$$\begin{aligned} & \underset{x_1, x_2}{\text{minimize}} && [c(x_1) + h_1(x_1)] + [c(x_2) + h_2(x_2)] \\ & \text{such that} && x_1 \in [0, 1] \text{ and } x_2 \in [0, 1]. \end{aligned}$$

The solution $\tilde{x} = (\tilde{x}_1, \tilde{x}_2)$ is given implicitly by the first order conditions

$$c'(\tilde{x}_i) = -h'_i(\tilde{x}_i), \quad i = 1, 2,$$

and is given explicitly by

$$\tilde{x}_i = \frac{\pi_{i2}z_2}{2}, \quad i = 1, 2.$$

We refer to \tilde{x}_i as the efficient level of care for activity f_i . It is the level of care at which the marginal cost of care equals the marginal benefit (the marginal reduction in expected harm).

¹⁴Without the second assumption, the regulator could not have correct beliefs when she becomes fully aware, which would violate the first assumption.

¹⁵We make the latter assumption for simplicity; it is without loss of generality given the former.

¹⁶We leave aside the problem of optimal enforcement (Becker, 1968) and assume that the company always complies with the standard of care set by the regulator.

4.2 New Link

We start with the case of a new link. To illustrate this case, we assume that the regulator initially perceives activity f_1 as safe and activity f_2 as risky. That is, we assume the regulator initially perceives the event $\Delta = \{s_3, s_4\}$ as infeasible (null). This implies $p_3 = p_4 = 0$. We can depict the original feasible state space, $S \subset Z^F$, as follows:

p	p_1	p_2
$F \setminus S$	s_1	s_2
f_1	0	0
f_2	0	z_2

Given S and p , the efficient levels of care are

$$\tilde{x}_1 = 0 \quad \text{and} \quad \tilde{x}_2 = \frac{p_2 z_2}{2},$$

and the regulator sets $\bar{x}_1 = \tilde{x}_1$ and $\bar{x}_2 = \tilde{x}_2$ as the standards of care for f_1 and f_2 , respectively.

Suppose the regulator discovers that activity f_1 is risky. For instance, suppose the company engages in f_1 and it result in harm z_2 . The feasible state space expands to $\widehat{S} = S \cup \Delta$ and the regulator revises her beliefs from p to \widehat{p} :

\widehat{p}	\widehat{p}_1	\widehat{p}_2	\widehat{p}_3	\widehat{p}_4
$F \setminus \widehat{S}$	s_1	s_2	s_3	s_4
f_1	0	0	z_2	z_2
f_2	0	z_2	0	z_2

Observe that for each state s in Δ there is an event $L(s)$ in S that corresponds with s on activity f_2 . Specifically, $L(s_3) = \{s_1\}$ and $L(s_4) = \{s_2\}$.

We assume that, by virtue of the discovery, the regulator learns that f_1 yields harm z_2 with probability $\delta > 0$. By definition, $\delta = \widehat{p}(\Delta) = \widehat{p}_3 + \widehat{p}_4$. It follows from Theorem 1 that the revised probability distribution \widehat{p} is given by:

Proposition 1. $\widehat{p}_1 = (1 - \delta)p_1$, $\widehat{p}_2 = (1 - \delta)p_2$, $\widehat{p}_3 = \delta p_1$, and $\widehat{p}_4 = \delta p_2$.

Given \widehat{S} and \widehat{p} , the efficient levels of care are

$$\widehat{x}_1 = \frac{(\widehat{p}_3 + \widehat{p}_4)z_2}{2} = \frac{\delta z_2}{2} \quad \text{and} \quad \widehat{x}_2 = \frac{(\widehat{p}_2 + \widehat{p}_4)z_2}{2} = \frac{p_2 z_2}{2}.$$

Note that $\widehat{x}_1 > \widetilde{x}_1$ but $\widehat{x}_2 = \widetilde{x}_2$. Thus, the discovery that f_1 is risky necessitates the stipulation of a new standard of care for f_1 but not for f_2 .

4.3 New Act

We next consider the case of a new act. We assume that $S = Z^F$:

p	p_1	p_2	p_3	p_4
$F \setminus S$	s_1	s_2	s_3	s_4
f_1	0	0	z_2	z_2
f_2	0	z_2	0	z_2

Given S and p , the efficient levels of care are

$$\widetilde{x}_1 = \frac{(p_3 + p_4)z_2}{2} \quad \text{and} \quad \widetilde{x}_2 = \frac{(p_2 + p_4)z_2}{2},$$

and the regulator sets $\bar{x}_1 = \widetilde{x}_1$ and $\bar{x}_2 = \widetilde{x}_2$ as the standards of care for f_1 and f_2 , respectively.

Suppose the regulator discovers a new activity, f_3 , which can cause harm. For instance, suppose the company invents and engages in f_3 and it results in harm z_2 . The expanded

feasible state space is $\widehat{S} = \Delta_1 \cup \Delta_2$, where $\Delta_1 = \{s_1, s_2, s_3, s_4\}$ and $\Delta_2 = \{s_5, s_6, s_7, s_8\}$:

\widehat{p}	\widehat{p}_1	\widehat{p}_2	\widehat{p}_3	\widehat{p}_4	\widehat{p}_5	\widehat{p}_6	\widehat{p}_7	\widehat{p}_8
$F \setminus \widehat{S}$	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8
f_1	0	0	z_2	z_2	0	0	z_2	z_2
f_2	0	z_2	0	z_2	0	z_2	0	z_2
f_3	0	0	0	0	z_2	z_2	z_2	z_2

Observe that Δ_1 is an augmented copy of S in which f_3 yields no harm in every state, and that Δ_2 is an augmented copy of S in which f_3 yields harm z_2 in every state. Stated differently, each state in S is split into two depending on whether f_3 yields no harm or harm z_2 . Thus, for each state s in S there is a corresponding event $E(s)$ in \widehat{S} . Specifically, $E(s_1) = \{s_1, s_5\}$, $E(s_2) = \{s_2, s_6\}$, $E(s_3) = \{s_3, s_7\}$, and $E(s_4) = \{s_4, s_8\}$.

We assume that, by virtue of the discovery, the regulator learns that f_3 yields harm z_2 with probability $\delta > 0$. Thus, $1 - \delta = \widehat{p}(\Delta_1) = \widehat{p}_1 + \widehat{p}_2 + \widehat{p}_3 + \widehat{p}_4$ and $\delta = \widehat{p}(\Delta_2) = \widehat{p}_5 + \widehat{p}_6 + \widehat{p}_7 + \widehat{p}_8$. It follows from Theorem 2 that the revised probability distribution \widehat{p} is given by:

Proposition 2. $\widehat{p}_1 = (1 - \delta)p_1$, $\widehat{p}_2 = (1 - \delta)p_2$, $\widehat{p}_3 = (1 - \delta)p_3$, $\widehat{p}_4 = (1 - \delta)p_4$,
 $\widehat{p}_5 = \delta p_1$, $\widehat{p}_6 = \delta p_2$, $\widehat{p}_7 = \delta p_3$, and $\widehat{p}_8 = \delta p_4$.

Given \widehat{S} and \widehat{p} , the efficient levels of care are

$$\begin{aligned} \widehat{x}_1 &= \frac{(\widehat{p}_3 + \widehat{p}_4 + \widehat{p}_7 + \widehat{p}_8) z_2}{2} = \frac{(p_3 + p_4) z_2}{2}, \\ \widehat{x}_2 &= \frac{(\widehat{p}_2 + \widehat{p}_4 + \widehat{p}_6 + \widehat{p}_8) z_2}{2} = \frac{(p_2 + p_4) z_2}{2}, \\ \text{and } \widehat{x}_3 &= \frac{(\widehat{p}_5 + \widehat{p}_6 + \widehat{p}_7 + \widehat{p}_8) z_2}{2} = \frac{\delta z_2}{2}. \end{aligned}$$

Thus, the discovery of f_3 necessitates the stipulation of a new standard of care, \widehat{x}_3 , but it does not necessitate the stipulation of a new standards of care for f_1 or f_2 .

4.4 New Consequence

We last consider the case of a new consequence. As above, we assume that $S = Z^F$:

p	p_1	p_2	p_3	p_4
$F \setminus S$	s_1	s_2	s_3	s_4
f_1	0	0	z_2	z_2
f_2	0	z_2	0	z_2

Given S and p , the efficient levels of care are

$$\tilde{x}_1 = \frac{(p_3 + p_4)z_2}{2} \quad \text{and} \quad \tilde{x}_2 = \frac{(p_2 + p_4)z_2}{2},$$

and the regulator sets $\bar{x}_1 = \tilde{x}_1$ and $\bar{x}_2 = \tilde{x}_2$ as the standards of care for f_1 and f_2 , respectively.

Suppose the regulator discovers a new consequence, $z_3 > z_2$, which she links to f_1 and f_2 . For instance, suppose the company engages in f_1 and f_2 and each results in harm z_3 . The expanded feasible state space is $\widehat{S} = S \cup \Delta$, where $\Delta = \{s_5, s_6, s_7, s_8, s_9\}$:

\widehat{p}	\widehat{p}_1	\widehat{p}_2	\widehat{p}_3	\widehat{p}_4	\widehat{p}_5	\widehat{p}_6	\widehat{p}_7	\widehat{p}_8	\widehat{p}_9
$F \setminus \widehat{S}$	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9
f_1	0	0	z_2	z_2	z_3	z_3	0	z_2	z_3
f_2	0	z_2	0	z_2	0	z_2	z_3	z_3	z_3

Observe that for each state s in Δ there is an event $C(s)$ in S that corresponds with s on the activity that does not yield harm z_3 . Specifically, $C(s_5) = \{s_1, s_3\}$, $C(s_6) = \{s_2, s_4\}$, $C(s_7) = \{s_1, s_2\}$, $C(s_8) = \{s_3, s_4\}$, and $C(s_9) = \{\emptyset\}$.¹⁷

We assume that, by virtue of the discovery, the regulator learns that activity f_1 yields harm z_3 with probability $\alpha > 0$ and that activity f_2 yields harm z_3 with probability $\beta > 0$.

¹⁷In this example, $I(s_5) = I(s_6) = \{1\}$, $I(s_7) = I(s_8) = \{2\}$, and $I(s_9) = \{1, 2\}$. Accordingly, $\bar{I}(s_5) = \bar{I}(s_6) = \{2\}$, $\bar{I}(s_7) = \bar{I}(s_8) = \{1\}$, and $\bar{I}(s_9) = \{\emptyset\}$.

By definition, $\alpha = \hat{p}_5 + \hat{p}_6 + \hat{p}_9$ and $\beta = \hat{p}_7 + \hat{p}_8 + \hat{p}_9$. It follows from Theorem 3 that the revised probability distribution \hat{p} is given by:

Proposition 3. $\hat{p}_1 = (1 - \alpha)(1 - \beta)p_1$, $\hat{p}_2 = (1 - \alpha)(1 - \beta)p_2$, $\hat{p}_3 = (1 - \alpha)(1 - \beta)p_3$,
 $\hat{p}_4 = (1 - \alpha)(1 - \beta)p_4$, $\hat{p}_5 = \alpha(1 - \beta)(p_1 + p_3)$, $\hat{p}_6 = \alpha(1 - \beta)(p_2 + p_4)$,
 $\hat{p}_7 = \beta(1 - \alpha)(p_1 + p_2)$, $\hat{p}_8 = \beta(1 - \alpha)(p_3 + p_4)$, and $\hat{p}_9 = \alpha\beta$.

Let $\delta = \hat{p}(\Delta) = \hat{p}_5 + \hat{p}_6 + \hat{p}_7 + \hat{p}_8 + \hat{p}_9$. Note that $\delta = \alpha + \beta - \alpha\beta$ and $1 - \delta = (1 - \alpha)(1 - \beta)$.

We can rewrite \hat{p} in terms of δ as follows:

Corollary 1. $\hat{p}_1 = (1 - \delta)p_1$, $\hat{p}_2 = (1 - \delta)p_2$, $\hat{p}_3 = (1 - \delta)p_3$, $\hat{p}_4 = (1 - \delta)p_4$,
 $\hat{p}_5 = \frac{\alpha}{1 - \alpha}(1 - \delta)(p_1 + p_3) = (\delta - \beta)(p_1 + p_3)$, $\hat{p}_6 = \frac{\alpha}{1 - \alpha}(1 - \delta)(p_2 + p_4) = (\delta - \beta)(p_2 + p_4)$,
 $\hat{p}_7 = \frac{\beta}{1 - \beta}(1 - \delta)(p_1 + p_2) = (\delta - \alpha)(p_1 + p_2)$, $\hat{p}_8 = \frac{\beta}{1 - \beta}(1 - \delta)(p_3 + p_4) = (\delta - \alpha)(p_3 + p_4)$,
and $\hat{p}_9 = \alpha + \beta - \delta$.

Given \hat{S} and \hat{p} , the efficient levels of care are

$$\begin{aligned} \hat{x}_1 &= \frac{(\hat{p}_3 + \hat{p}_4 + \hat{p}_8)z_2 + (\hat{p}_5 + \hat{p}_6 + \hat{p}_9)z_3}{2} = \frac{(1 - \alpha)(p_3 + p_4)z_2 + \alpha z_3}{2} \\ \text{and } \hat{x}_2 &= \frac{(\hat{p}_2 + \hat{p}_4 + \hat{p}_6)z_2 + (\hat{p}_7 + \hat{p}_8 + \hat{p}_9)z_3}{2} = \frac{(1 - \beta)(p_2 + p_4)z_2 + \beta z_3}{2}. \end{aligned}$$

Note that $\hat{x}_1 > \tilde{x}_1$ and $\hat{x}_2 > \tilde{x}_2$. Thus, the discovery of z_3 necessitates the stipulation of new standards of care for both f_1 and f_2 .

4.5 Act Independence

We conclude our application with a few remarks about the importance of act independence. As previously noted, reverse Bayesianism alone is not sufficient to fully identify the revised probability distribution \hat{p} . The reason is that reverse Bayesianism implies restrictions on the revised probabilities of non-null states in the original state space (or, in the case of a new act, their corresponding events in the expanded state space), but not on the probabilities of new states in the expanded state space. In other words, reverse Bayesianism prescribes

how probability mass shifts away from non-null states in the original state space to the corresponding states or events in the expanded state space, but it does not dictate how the shifted probability mass is distributed among the new states in the expanded state space. This is where act independence comes in. It determines how the shifted probability mass is apportioned among the new states. Together, reverse Bayesianism and act independence fully identify the revised probability distribution \hat{p} .

How realistic is act independence? The answer depends on the nature of the specific activities in question. For instance, the risk that fracking for natural gas results in groundwater contamination is likely to be independent of the risk that importing liquefied natural gas results in a fire or explosion. By contrast, the risk of contracting HIV from sharing drug injection needles is likely to be correlated with the risk of contracting HIV from having unprotected sex, since both depend on the prevalence of HIV in the population.

Because there exist activities whose outcomes are not independent, it is useful to investigate the importance of the act independence assumption for our results.

New link. In the case of a new link, reverse Bayesianism alone implies $\hat{p}_1 = (1 - \delta)p_1$, $\hat{p}_2 = (1 - \delta)p_2$, and $\hat{p}_3 + \hat{p}_4 = \delta$. Importantly, reverse Bayesianism alone is not sufficient to separately identify \hat{p}_3 and \hat{p}_4 . This leaves a set of posteriors \hat{p} . Effectively, unawareness has been turned into ambiguity.¹⁸

As it turns out, this does not create an issue with respect to activity f_1 . Recall that, by assumption, the regulator learns δ (the probability that f_1 yields z_2). Because the efficient level of care for f_1 is a function of the sum $\hat{p}_3 + \hat{p}_4$, the regulator can stipulate a new standard of care for f_1 in terms of δ .

Relaxing act independence, however, creates ambiguity with respect to the revised risk of activity f_2 . Because the efficient level of care for f_2 is a function of the sum $\hat{p}_2 + \hat{p}_4$, without act independence (or another assumption that separately identifies \hat{p}_3 and \hat{p}_4), the

¹⁸Dominiak and Tserenjigmid (2018) have a similar result. In their model, growing awareness extends subjective expected utility preferences to maxmin expected utility preferences, and newly discovered events can be ambiguous. Thus, reverse Bayesian belief revision can result in a set of posteriors.

regulator cannot stipulate a precise new standard of care for f_2 . The best the regulator can do is specify lower and upper bounds, using the knowledge that $\widehat{p}_4 \in (0, \delta)$.

Of course, the ambiguity can be resolved if, by virtue of the discovery, the regulator learns more about \widehat{p} . For instance, if the regulator learns not only δ but also either $\widehat{p}_2 + \widehat{p}_4$ (the revised probability that f_2 yields z_2) or \widehat{p}_4 (the joint probability that f_1 and f_2 yield z_2), this is sufficient to separately identify \widehat{p}_3 and \widehat{p}_4 . With this, the regulator can stipulate a precise new standard of care for f_2 .¹⁹

New act. In the case of a new act, reverse Bayesianism alone implies $\widehat{p}_1 + \widehat{p}_5 = p_1$, $\widehat{p}_2 + \widehat{p}_6 = p_2$, $\widehat{p}_3 + \widehat{p}_7 = p_3$, $\widehat{p}_4 + \widehat{p}_8 = p_4$, and $\widehat{p}_5 + \widehat{p}_6 + \widehat{p}_7 + \widehat{p}_8 = \delta$. Again, this leaves a set of posteriors \widehat{p} . Recall that the efficient level of care for f_1 is a function of the sum $\widehat{p}_3 + \widehat{p}_4 + \widehat{p}_7 + \widehat{p}_8$, which equals $p_3 + p_4$; the efficient level of care for f_2 is a function of the sum $\widehat{p}_2 + \widehat{p}_4 + \widehat{p}_6 + \widehat{p}_8$, which equals $p_2 + p_4$; and the efficient level of care for f_3 is a function of the sum $\widehat{p}_5 + \widehat{p}_6 + \widehat{p}_7 + \widehat{p}_8$, which equals δ . Hence, even without act independence, the regulator's information is sufficiently precise (i) to know that she need not stipulate new standards of care for activities f_1 and f_2 and (ii) to stipulate a standard of care for the new activity f_3 .

New consequence. In the case of a new consequence, reverse Bayesianism alone implies $\widehat{p}_1 = (1-\delta)p_1$, $\widehat{p}_2 = (1-\delta)p_2$, $\widehat{p}_3 = (1-\delta)p_3$, $\widehat{p}_4 = (1-\delta)p_4$, and $\widehat{p}_5 + \widehat{p}_6 + \widehat{p}_7 + \widehat{p}_8 + \widehat{p}_9 = \delta$. Once again, this leaves a set of posteriors \widehat{p} . By assumption, the regulator learns $\widehat{p}_5 + \widehat{p}_6 + \widehat{p}_9 = \alpha$ (the probability that f_1 yields z_3) and $\widehat{p}_7 + \widehat{p}_8 + \widehat{p}_9 = \beta$ (the probability that f_2 yields z_3). Assume the regulator also learns \widehat{p}_9 (the joint probability that f_1 and f_2 yield z_3), and let $\widehat{p}_9 = \gamma$. Note that $\delta = \alpha + \beta - \gamma$.

Recall that the efficient level of care for activity f_1 is a function of α and the sum $\widehat{p}_3 + \widehat{p}_4 + \widehat{p}_8$ (the revised probability that f_1 yields z_2), and the efficient level of care for activity f_2 is a function of β and the sum $\widehat{p}_2 + \widehat{p}_4 + \widehat{p}_6$ (the revised probability that f_2 yields

¹⁹We note that, in the case of a new link, reverse Bayesianism fully identifies \widehat{p} without act independence if the probability that activity f_2 yields harm z_2 is unchanged by the discovery of the new link between f_1 and z_2 (i.e., $\pi_{22} = p_2 + p_4 = \widehat{p}_2 + \widehat{p}_4 + \widehat{p}_4 = \widehat{\pi}_{22}$).

z_2). Without act independence, the sums $\hat{p}_3 + \hat{p}_4 + \hat{p}_8$ and $\hat{p}_2 + \hat{p}_4 + \hat{p}_6$ are only partially identified (because \hat{p}_6 and \hat{p}_8 are not separately identified), creating ambiguity with respect to the revised risks of both activities. As a result, the regulator cannot stipulate precise new standards of care for activities f_1 and f_2 . The best the regulator can do is specify bounds:

$$\begin{aligned} \hat{x}_1 &\in \left(\frac{(1-\delta)(p_3+p_4)z_2 + \alpha z_3}{2}, \frac{((1-\delta)(p_3+p_4) + \delta)z_2 + \alpha z_3}{2} \right) \\ \text{and } \hat{x}_2 &\in \left(\frac{(1-\delta)(p_2+p_4)z_2 + \beta z_3}{2}, \frac{((1-\delta)(p_2+p_4) + \delta)z_2 + \beta z_3}{2} \right). \end{aligned}$$

As before, the ambiguity can be resolved if, by virtue of the discovery, the regulator learns more about \hat{p} . For instance, if the regulator learns not only δ and γ but also either $\hat{p}_3 + \hat{p}_4 + \hat{p}_8$ or $\hat{p}_2 + \hat{p}_4 + \hat{p}_6$, this is sufficient to separately identify \hat{p}_5 , \hat{p}_6 , \hat{p}_7 , and \hat{p}_8 . With this, the regulator can stipulate precise new standards of care for f_1 and f_2 .

In summary, without act independence, reverse Bayesianism only partially identifies \hat{p} . This does not create an issue in the case of a new act—the regulator’s information is sufficiently precise to stipulate a standard of care with respect to each activity. In the case of a new link or consequence, by contrast, the partial identification of \hat{p} creates ambiguity with respect to the revised risk of one or both activities, leading to imprecise standards. This ambiguity, however, can be resolved if the regulator learns more about \hat{p} . In other words, the more the regulator learns about the new probability of harm, the less important is the act independence assumption.

5 Conclusion

For economists who wish to incorporate unawareness and growing awareness into applications, reverse Bayesianism offers an elegant choice-theoretic belief revision theory that mirrors the familiar process of Bayesian updating. An important limitation of Karni and Vierø’s (2013) model, however, is that reverse Bayesianism alone does not fully determine the

revised probability distribution over the expanded state space. We overcome this limitation in a relatively simple way, by assuming that acts are statistically independent. We show that with act independence, and knowledge of the probabilities of new act events in the expanded state space, reverse Bayesianism fully determines the revised probability distribution over the expanded state space in each case of growing awareness. In this way, we make a contribution to the reverse Bayesian model and operationalize it for economic applications.²⁰

To illustrate how act independence operationalizes reverse Bayesianism, we consider the law and economics problem of safety regulation. We analyze how a safety regulator, in the wake of growing awareness about the risky activities within her purview, revises her beliefs about the risk of each activity and resets the safety standards for each activity. Of course, unawareness and growing awareness—via technological progress, scientific discovery, or otherwise—play an important role in many legal contexts. Accordingly, we believe that the reverse Bayesian model could be fruitfully applied to study the implications of unawareness and growing awareness for the economic analysis of numerous other legal subjects, including contract remedies, criminal law, litigation and settlement, and tort law.

²⁰At the same time, the model has other limitations that we do not address. For instance, Chambers and Hayashi (2018) criticize its empirical content from a revealed preference perspective. They show that, in the case of a new consequence, the model does not make singular predictions about observable choices over feasible acts. Another limitation of the model is that it assumes a naive or myopic unawareness—people are unaware that they are unaware. A sophisticated unawareness, where people are aware that they are unaware, may be more realistic. Aware of this limitation, Karni and Vierø (2017) extend their model to the case of sophisticated unawareness. The end result is a generalization that maintains the flavor of reverse Bayesianism and nests the naive model as a special case.

Appendix

Proof of Theorem 1

(i) Take any $s \in S$. By reverse Bayesianism, we have $|S| - 1$ linearly independent equations:

$$\widehat{p}(t) = \frac{p(t)}{p(s)}\widehat{p}(s), \quad \forall t \in S, t \neq s. \quad (1.1)$$

By the definition of δ and $\sum_{t \in \widehat{S}} \widehat{p}(t) = 1$, we have

$$\sum_{t \in S} \widehat{p}(t) = 1 - \delta. \quad (1.2)$$

Substituting (1.1) into (1.2), we have

$$\widehat{p}(s) + \sum_{t \in S: t \neq s} \frac{p(t)}{p(s)}\widehat{p}(s) = 1 - \delta,$$

which implies

$$\widehat{p}(s) = \frac{(1 - \delta)p(s)}{\sum_{t \in S} p(t)} = (1 - \delta)p(s), \quad (1.3)$$

where the last equality follows from $\sum_{t \in S} p(t) = 1$.

(ii) Take any $s \in \Delta$. By act independence,

$$\widehat{p}(s) = \prod_{i=1}^m \widehat{p}(A_i(s^i)).$$

Observe that $\widehat{p}(A_l(s^l)) = \widehat{p}(A_l(z_k)) = \delta$ and $\bigcap_{i \neq l} A_i(s^i) = L(s) \cup \{s\}$. It follows that

$$\begin{aligned} \widehat{p}(s) &= \delta \prod_{i \neq l} \widehat{p}(A_i(s^i)) = \delta \widehat{p}\left(\bigcap_{i \neq l} A_i(s^i)\right) \\ &= \delta \widehat{p}(L(s) \cup \{s\}) = \delta [\widehat{p}(L(s)) + \widehat{p}(s)], \end{aligned}$$

which implies

$$\widehat{p}(s) = \frac{\delta}{1 - \delta} \widehat{p}(L(s)). \quad (1.4)$$

Observe that $L(s)$ is the union of all $t \in S$ such that $t^i = s^i$ for all $i \neq l$. It follows that

$$\widehat{p}(L(s)) = \sum_{t \in L(s)} \widehat{p}(t) = \sum_{t \in L(s)} (1 - \delta)p(t) = (1 - \delta)p(L(s)), \quad (1.5)$$

where the second equality follows from (1.3). Substituting (1.5) back into (1.4), we have $\widehat{p}(s) = \delta p(L(s))$.

Proof of Theorem 2

Take any $s \in S$. By reverse Bayesianism, we have $|S| - 1$ linearly independent equations:

$$p(t)\widehat{p}(E(s)) = p(s)\widehat{p}(E(t)), \quad \forall t \in S, t \neq s.$$

Summing the left- and right-hand sides, and adding $p(s)\widehat{p}(E(s))$ to each side, yields

$$\widehat{p}(E(s)) \sum_{t \in S} p(t) = p(s) \sum_{t \in S} \widehat{p}(E(t)).$$

Because $\sum_{t \in S} p(t) = 1$ and $\sum_{t \in S} \widehat{p}(E(t)) = 1$, we have

$$\widehat{p}(E(s)) = p(s). \tag{2.1}$$

Take any $s_j \in E(s)$, $j \in \{1, \dots, n\}$. By act independence,

$$\widehat{p}(s_j) = \prod_{i=1}^{m+1} \widehat{p}(A_i(s_j^i)).$$

Observe that $\widehat{p}(A_{m+1}(s_j^{m+1})) = \widehat{p}(A_{m+1}(z_j)) = \delta_j$ and $\bigcap_{i=1}^m A_i(s_j^i) = E(s)$. It follows that

$$\widehat{p}(s_j) = \delta_j \prod_{i=1}^m \widehat{p}(A_i(s_j^i)) = \delta_j \widehat{p}\left(\bigcap_{i=1}^m A_i(s_j^i)\right) = \delta_j \widehat{p}(E(s)). \tag{2.2}$$

Substituting (2.1) into (2.2), we have $\widehat{p}(s_j) = \delta_j p(s)$.

Proof of Theorem 3

(i) Take any $s \in S$. By reverse Bayesianism, we have $|S| - 1$ linearly independent equations:

$$p(t)\widehat{p}(s) = p(s)\widehat{p}(t), \quad \forall t \in S, t \neq s.$$

Summing the left- and right-hand sides, and adding $p(s)\widehat{p}(s)$ to each side, yields

$$\widehat{p}(s) \sum_{t \in S} p(t) = p(s) \sum_{t \in S} \widehat{p}(t).$$

Observe that $\sum_{t \in S} p(t) = 1$ and $\sum_{t \in S} \widehat{p}(t) = 1 - \delta = \prod_{i=1}^m (1 - \alpha_i)$. Thus,

$$\widehat{p}(s) = (1 - \delta)p(s) = \left(\prod_{i=1}^m (1 - \alpha_i)\right) p(s). \tag{3.1}$$

(ii) Take any $s \in \Delta$ such that $I(s) = \{k\}$ for any $k \in \{1, \dots, m\}$. By act independence,

$$\widehat{p}(s) = \prod_{i=1}^m \widehat{p}(A_i(s^i)).$$

Observe that $\widehat{p}(A_k(s^k)) = \widehat{p}(A_k(z_{n+1})) = \alpha_k$. Thus,

$$\widehat{p}(s) = \alpha_k \prod_{i \in \bar{I}(s)} \widehat{p}(A_i(s^i)).$$

Observe that $I(s) = \{k\}$ implies $\bigcap_{i \in \bar{I}(s)} A_i(s^i) = C(s) \cup \{s\}$. Hence,

$$\begin{aligned} \widehat{p}(s) &= \alpha_k \prod_{i \in \bar{I}(s)} \widehat{p}(A_i(s^i)) = \alpha_k \widehat{p}\left(\bigcap_{i \in \bar{I}(s)} A_i(s^i)\right) \\ &= \alpha_k \widehat{p}(C(s) \cup \{s\}) = \alpha_k (\widehat{p}(C(s)) + \widehat{p}(s)), \end{aligned}$$

which implies

$$\widehat{p}(s) = \frac{\alpha_k}{1 - \alpha_k} \widehat{p}(C(s)). \quad (3.2)$$

Observe that $C(s)$ is the union of all $t \in S$ such that $t^i = s^i$ for all $i \in \bar{I}(s)$. It follows that

$$\widehat{p}(C(s)) = \sum_{t \in C(s)} \widehat{p}(t) = \sum_{t \in C(s)} (1 - \delta)p(t) = (1 - \delta)p(C(s)), \quad (3.3)$$

where the second equality follows from (3.1). Substituting (3.3) back into (3.2), we have

$$\widehat{p}(s) = \frac{\alpha_k}{1 - \alpha_k} (1 - \delta)p(C(s)) = \alpha_k \prod_{i \in \bar{I}(s)} (1 - \alpha_i)p(C(s)),$$

where the last equality follows from $1 - \delta = \prod_{i=1}^m (1 - \alpha_i)$.

Next take any $s \in \Delta$ such that $I(s) = \{k, l\}$ for any $\{k, l\} \subset \{1, \dots, m\}$. By act independence,

$$\widehat{p}(s) = \prod_{i=1}^m \widehat{p}(A_i(s^i)).$$

Observe that $\widehat{p}(A_k(s^k)) = \widehat{p}(A_k(z_{n+1})) = \alpha_k$. Thus,

$$\widehat{p}(s) = \alpha_k \prod_{i \in \{\bar{I}(s) \cup \{l\}\}} \widehat{p}(A_i(s^i)).$$

Observe that $I(s) = \{k, l\}$ implies $\bigcap_{i \in \{\bar{I}(s) \cup \{l\}\}} A_i(s^i) = D(s) \cup \{s\}$, where $D(s) \equiv \{r \in \Delta : r^i = s^i, \forall i \in \{\bar{I}(s) \cup \{l\}\}\}$. Hence,

$$\begin{aligned} \widehat{p}(s) &= \alpha_k \prod_{i \in \{\bar{I}(s) \cup \{l\}\}} \widehat{p}(A_i(s^i)) = \alpha_k \widehat{p}\left(\bigcap_{i \in \{\bar{I}(s) \cup \{l\}\}} A_i(s^i)\right) \\ &= \alpha_k \widehat{p}(D(s) \cup \{s\}) = \alpha_k (\widehat{p}(D(s)) + \widehat{p}(s)), \end{aligned}$$

which implies

$$\widehat{p}(s) = \frac{\alpha_k}{1 - \alpha_k} \widehat{p}(D(s)). \quad (3.4)$$

Observe further that $I(r) = \{l\}$ for all $r \in D(s)$. It follows that

$$\begin{aligned} \widehat{p}(D(s)) &= \sum_{t \in D(s)} \widehat{p}(t) \\ &= \sum_{t \in D(s)} \frac{\alpha_l}{1 - \alpha_l} (1 - \delta) p(C(t)) \\ &= \frac{\alpha_l}{1 - \alpha_l} (1 - \delta) p(C(s)). \end{aligned} \quad (3.5)$$

Substituting (3.5) back into (3.4), we have

$$\begin{aligned} \widehat{p}(s) &= \frac{\alpha_k}{1 - \alpha_k} \frac{\alpha_l}{1 - \alpha_l} (1 - \delta) p(C(s)). \\ &= \alpha_k \alpha_l \prod_{i \in \bar{I}(s)} (1 - \alpha_i) p(C(s)). \end{aligned}$$

Proceeding in this fashion to consider $s \in \Delta$ such that $I(s)$ is an ι -element subset of $\{1, \dots, m\}$ for all $\iota = 3, \dots, m - 1$, we establish that

$$\widehat{p}(s) = \left(\prod_{i \in I(s)} \alpha_i \right) \left(\prod_{i \in \bar{I}(s)} (1 - \alpha_i) \right) p(C(s))$$

for all $s \in \Delta$ such that $I(s) \subset \{1, \dots, m\}$.

(iii) Take the $s \in \Delta$ such that $I(s) = \{1, \dots, m\}$. By act independence, $\widehat{p}(s) = \prod_{i=1}^m \widehat{p}(A_i(s^i))$. Observe that $\widehat{p}(A_i(s^i)) = \widehat{p}(A_i(z_{n+1})) = \alpha_i$ for all $i \in I(s)$. Because $I(s) = \{1, \dots, m\}$, we have $\widehat{p}(s) = \prod_{i=1}^m \alpha_i$.

References

- Auster, S. “Asymmetric Awareness and Moral Hazard.” *Games and Economic Behavior*, Vol. 82 (2013), pp. 503–521.
- Becker, G. S. “Crime and Punishment: An Economic Approach.” *Journal of Political Economy*, Vol. 76 (1968), pp. 169–217.
- Board, O. and Chung, K.-S. “Object-Based Unawareness: Theory and Applications.” Working Paper, School of Law, New York University, 2011.
- Chambers, C. P. and Hayashi, T. “Reverse Bayesianism: A Comment.” *American Economic Journal: Microeconomics*, Vol. 10 (2018), pp. 315–324.
- Collee, J. G. and Bradley, R. “BSE: A Decade On—Part I.” *The Lancet*, Vol. 349 (1997), pp. 636–641.
- Dekel, E., Lipman, B. L., and Rustichini, A. “Recent Developments in Modeling Unforeseen Contingencies.” *European Economic Review*, Vol. 42 (1998a), pp. 523–542.
- Dekel, E., Lipman, B. L., and Rustichini, A. “Standard State-Space Models Preclude Unawareness.” *Econometrica*, Vol. 66 (1998b), pp. 159–173.
- Dominiak, A. and Tserenjigmid, G. “Ambiguity under Growing Unawareness.” Working Paper, Department of Economics, Virginia Tech, 2018.
- Fagin, R. and Halpern, J. Y. “Belief, Awareness, and Limited Reasoning.” *Artificial Intelligence*, Vol. 34 (1988), pp. 39–76.
- Filiz-Ozbay, E. “Incorporating Unawareness into Contract Theory.” *Games and Economic Behavior*, Vol. 76 (2012), pp. 181–194.
- Gold, R. *The Boom: How Fracking Ignited the American Energy Revolution and Changed the World*. New York: Simon & Schuster, 2014.
- Grant, S. J., Kline, J., and Quiggin, J. “Differential Awareness, Ambiguity, and Incomplete Contracts: A Model of Contractual Disputes.” *Journal of Economic Behavior and Organization*, Vol. 82 (2012), pp. 494–504.
- Grant, S., Meneghel, I., and Tourky, R. “Learning under Unawareness.” Working Paper, Research School of Economics, Australian National University, 2019.
- Halpern, J. Y. “Alternative Semantics for Unawareness.” *Games and Economic Behavior*, Vol. 37 (2001), pp. 321–339.
- Halpern, J. Y. and Rêgo, L. C. “Interactive Unawareness Revisited.” *Games and Economic Behavior*, Vol. 62 (2008), pp. 232–262.
- Heifetz, A., Meier, M., and Schipper, B. C. “Interactive Unawareness.” *Games and Economic Behavior*, Vol. 130 (2006), pp. 78–94.

- Karni, E. and Vierø, M.-L. “Reverse Bayesianism: A Choice-Based Theory of Growing Awareness.” *American Economic Review*, Vol. 103 (2013), pp. 2790–2810.
- Karni, E. and Vierø, M.-L. “Probabilistic Sophistication and Reverse Bayesianism.” *Journal of Risk and Uncertainty*, Vol. 50 (2015), pp. 189–208.
- Karni, E. and Vierø, M.-L. “Awareness of Unawareness: A Theory of Decision Making in the Face of Ignorance.” *Journal of Economic Theory*, Vol. 168 (2017), pp. 301–328.
- Lindsley, C. W. “Chronic Traumatic Encephalopathy (CTE): A Brief Historical Overview and Recent Focus on NFL Players.” *ACS Chemical Neuroscience*, Vol. 8 (2017), pp. 1629–1631.
- Modica, S. and Rustichini, A. “Awareness and Partitional Information Structures.” *Theory and Decision*, Vol. 27 (1994), pp. 265–298.
- Modica, S. and Rustichini, A. “Unawareness and Partitional Information Structures.” *Games and Economic Behavior*, Vol. 37 (1999), pp. 107–124.
- National Academies of Sciences, Engineering, and Medicine. “Veterans and Agent Orange: Update 11.” Washington, DC: National Academic Press, 2018.
- Savage, L. J. *The Foundations of Statistics*. New York: Wiley, 1954.
- Schipper, B. C. “Preference-Based Unawareness.” *Mathematical Social Sciences*, Vol. 70 (2014a), pp. 1–9.
- Schipper, B. C. “Unawareness—A Gentle Introduction to Both the Literature and the Special Issue.” *Mathematical Social Sciences*, Vol. 70 (2014b), pp. 34–41.
- Schipper, B. C. “Awareness.” In H. van Ditmarsh, J. Y. Halpern, W. van der Hock, and B. Kooi, eds., *Handbook of Epistemic Logic*. London: College Publications, 2015.
- U.S. Centers for Disease Control and Prevention. “Thirty Years of HIV — 1981–2011.” *Morbidity and Mortality Weekly Report*, Vol. 60, No. 21 (2011), p. 689.
- von Thadden, E.-L. and Zhao, X. “Incentives for Unaware Agents.” *Review of Economic Studies*, Vol. 77 (2012), pp. 197–222.
- von Thadden, E.-L. and Zhao, X. “Multi-task Agency with Unawareness.” *Theory and Decision*, Vol. 79 (2014), pp. 1151–1174.
- Zhao, X. “Framing Contingencies in Contracts.” *Mathematical Social Sciences*, Vol. 61 (2011), pp. 31–40.