Stanford – Vienna
Transatlantic Technology Law Forum

A joint initiative of
Stanford Law School and the University of Vienna School of Law

TTLF Working Papers

No. 158

Regulating Life Technologies at the Nanoscale: EU-US Comparative Perspectives on Dynamic Licensing under Deep Uncertainty

Umberto Nizza

2026

TTLF Working Papers

Editors: Siegfried Fina, Mark Lemley, and Roland Vogl

About the TTLF Working Papers

TTLF's Working Paper Series presents original research on technology, and business-related law and policy issues of the European Union and the US. The objective of TTLF's Working Paper Series is to share "work in progress". The authors of the papers are solely responsible for the content of their contributions and may use the citation standards of their home country. The TTLF Working Papers can be found at http://ttlf.stanford.edu. Please also visit this website to learn more about TTLF's mission and activities.

If you should have any questions regarding the TTLF's Working Paper Series, please contact Vienna Law Professor Siegfried Fina, Stanford Law Professor Mark Lemley or Stanford LST Executive Director Roland Vogl at the

Stanford-Vienna Transatlantic Technology Law Forum
http://ttlf.stanford.edu

Stanford Law School
Crown Quadrangle
559 Nathan Abbott Way
Stanford, CA 94305-8610

University of Vienna School of Law
Department of Business Law
Schottenbastei 10-16
1010 Vienna, Austria

About the Author

Umberto Nizza is an Adjunct Professor at the University of Bologna in Italy. He holds a Ph.D. in Law, Institutions, and Economics from Collegio Carlo Alberto and the University of Turin, Italy. Before entering academia, he worked as an attorney and an international consultant. His research lies at the intersection of law, economics, and emerging technology, with a particular focus on the regulation of artificial intelligence. His current research explores how regulatory frameworks can adapt to technologies whose risk profiles remain scientifically indeterminate at the point of authorization, a question this paper addresses formally through the case of nanotechnology, and which recurs across his work on artificial intelligence, algorithmic ESG oversight, and digital markets. A related strand of his work examines how emerging technologies are transforming the legal profession itself, from the adoption of legal tech to the liability implications of AI-generated advice. Another aspect of his work involves applying advanced econometric methods to the study of judicial institutions and examining how cultural and institutional factors shape the economic behavior of judicial actors, alongside related work on political productivity, party switching, and rent-seeking behavior in legislative and judicial settings. He welcomes opportunities to collaborate on future projects and to contribute his expertise at the intersection of law, economics, and technology, with the aim of informing evidence-based policy and governance.

General Note about the Content

The opinions expressed in this paper are those of the author and not necessarily those of the Transatlantic Technology Law Forum or any of its partner institutions, or the sponsors of this research project.

Suggested Citation

This TTLF Working Paper should be cited as:
Umberto Nizza, Regulating Life Technologies at the Nanoscale: EU-US Comparative Perspectives on Dynamic Licensing under Deep Uncertainty, Stanford-Vienna TTLF Working Paper No. 158, http://ttlf.stanford.edu.

Copyright

© 2026 Umberto Nizza

Abstract

Nanotechnology exposes a structural weakness common to the regulation of emerging technologies: legal authorization must often be granted before the scientific properties of the regulated object – and even its proper legal classification – are fully understood. This paper develops a formal law and economics model of dynamic regulatory licensing designed for precisely this setting. Departing from conventional regulatory theory, the model treats the regulator not as a Bayesian expected-utility maximizer but as a decision-maker operating under genuine ambiguity, facing a set of plausible probability distributions over harm rather than a single, known one. Framing the licensing decision as a dynamic optimal-stopping problem, the paper characterizes the welfare-maximizing policy as one in which authorization status, monitoring intensity, and regulatory obligations are continuously recalibrated in response to an ambiguity-adjusted harm index. Two extensions refine this baseline result: incorporating realistic administrative costs of granting and revoking authorization yields an optimal licensing rule in the form of a two-threshold hysteresis band rather than a single frictionless cutoff, and, in a setting where the innovator holds private information about the technology's risk type, the socially optimal incentive-compatible license menu requires monitoring the riskier type below its own first-best level, leaving an information rent to the safer type. The comparative analysis of European Union and United States nanotechnology governance developed in the paper shows that the European precautionary model and the American post-market model are each optimal only under conditions that nanotechnology systematically violates, and numerical simulations quantify the welfare gains of dynamic licensing relative to both benchmarks. The paper concludes by addressing the institutional implementation of the model – the allocation of regulatory functions between legislatures and specialized agencies, and the incentive-compatible production of regulatory information – and argues that the underlying framework, though developed through the empirical lens of nanotechnology, applies more broadly to any technological domain in which scientific understanding evolves throughout, rather than before, commercialization.

Table of Contents

1. Introduction

Few areas of regulation expose the limits of legal decision-making more clearly than emerging technologies, where legal choices must often precede scientific certainty. From the early governance of pharmaceutical compounds to the more recent challenges posed by artificial intelligence and genetic engineering, law has repeatedly confronted the problem of normative anticipation: how to govern what cannot yet be fully described, measured, or predicted.[1] Nanotechnology presents this problem in its most acute form. Indeed, at the nanoscale, matter behaves according to quantum mechanical principles that diverge from the classical physics underpinning most product safety law[2]. Substances that are inert in bulk form become reactive, conductive, or biologically active when reduced to dimensions between one and one hundred nanometers. A titanium dioxide particle used in sunscreen and a titanium dioxide nanoparticle may share an elemental identity while differing radically in their toxicological profile, their capacity to penetrate biological membranes, and their environmental persistence. The law, confronted with this duality, does not yet know what it is looking at.

Nanotechnology is, in this respect, best understood as a paradigmatic instance of a broader and increasingly common regulatory predicament, rather than as a self-contained object of legal analysis. The same structural features that make nanomaterials difficult to regulate – an object of regulation that resists stable classification, a

[1] The reference here is to the legal necessity of governing technological developments before their risks, applications, and social consequences can be fully characterized. See Kyle D Logue, 'Legal Transitions, Rational Expectations, and Legal Progress' (2003) 13 *Journal of Contemporary Legal Issues* 211.

[2] For an in-depth technical analysis of quantum mechanics and nanotechnology, which, for systemic reasons, cannot be accommodated in this context, see Vladimir V Mitin, Dmitry I Sementsov and Nizami Z Vagidov, *Quantum Mechanics for Nanostructures* (Cambridge University Press 2010) 52 ff.

probabilistic relationship between evidence and harm that is itself contested, and a rate of scientific learning that is neither negligible nor fast enough to be safely left to ex post correction – recur, in different empirical guises, across artificial intelligence, synthetic biology, and other emerging technological domains. The framework developed in this paper accordingly uses nanotechnology as a motivating and illustrative case, through which the general problem of regulatory licensing under ambiguity can be examined with analytical precision, rather than as a technology-specific model whose formal structure depends on the physics of the nanoscale.

The legal indeterminacy inherent in nanotechnology regulation is not a matter that will dissolve as scientific knowledge matures. It is a structural characteristic of the domain, reflecting a deeper epistemological condition that may be described as an ontological regulatory uncertainty. Conventional regulatory theory generally assumes that the object of regulation is identifiable even where its consequences remain uncertain.[3] The regulator may face, in this perspective, both risk and uncertainty, within a legislative framework in which the regulated entity, the harm it may cause, and the causal pathway connecting conduct to consequence are assumed to be legally articulable. Nanotechnologies challenge this assumption at its root, given that the regulated entity – the nanoparticle itself – resists stable legal classification precisely

[3] Cf., for an examination of the conventional law and economics approaches and schools of thought, in particular, Christopher J Webster, 'Public Choice, Pigouvian and Coasian Planning Theory' (1998) 35(1) *Urban Studies* 53; James M Buchanan, 'The Coase Theorem and the Theory of the State' (1973) 13 *Natural Resources Journal* 579; Hart Blanton, Anne E Stuart and Regina JJM van den Eijnden, 'An Introduction to Deviance-Regulation Theory: The Effect of Behavioral Norms on Message Framing' (2001) 27(7) *Personality and Social Psychology Bulletin* 848; Clifford Nowell and John Tschirhart, 'Testing Theories of Regulatory Behavior' (1993) 8(6) *Review of Industrial Organization* 653.

because its properties are not intrinsic but relational, varying with context, medium, and scale.[4]

The failure of existing regulatory frameworks to accommodate this challenge is well documented[5] yet relatively few attempt to explain why these shortcomings emerge from the epistemic structure of nanotechnology itself. The comparison between the European Union and the United States is particularly instructive because they embody the two most influential contemporary models of technology regulation. In the European Union, the precautionary principle – embedded in Article 191(2) of the Treaty on the Functioning of the European Union and elaborated through decades of jurisprudence from the Court of Justice[6] – provides the formal basis for regulatory intervention in the absence of certainty. Yet the precautionary principle, as currently operationalized, functions as a binary gate: it either permits deployment subject to pre-market authorization requirements or prohibits it pending further evaluation. This binary structure might be ill-suited to a domain where scientific understanding evolves continuously and where the relevant question is not "safe or unsafe" but "how safe, under what conditions of monitoring, and for how long." In the United States, the dominant approach relies on post-market correction, agency flexibility, and the corrective

[4] See, in this perspective, Mariana Córdoba, Fiorela Alassia and Alfio Zambon, 'Identity in the Nanoworld: Processes and Contextuality' (2024) 26(3) *Foundations of Chemistry* 413, 422.

[5] For concrete examples of the challenges currently posed by nanomaterials in the regulation of drugs, pesticides, and cosmetics, see Racliffe WS Lai, Katie WY Yeung, Mana MM Yung, Aleksandra B Djurišić, John P Giesy and Kenneth MY Leung, 'Regulation of Engineered Nanomaterials: Current Challenges, Insights and Future Directions' (2018) 25(4) *Environmental Science and Pollution Research* 3060.

[6] Cf., on this topic, Monika A Król, 'Prevention Principle and the Precautionary Principle' in Javier Cremades and Cristina Hermida (eds), *Encyclopedia of Contemporary Constitutionalism* (Springer International Publishing 2023) 1; Kristel De Smedt and Ellen Vos, 'The Application of the Precautionary Principle in the EU' in Harald Mieg (ed), *The Responsibility of Science* (Springer International Publishing 2022) 163; Tuomas Pöysti, 'The Precautionary Approach Design Pattern' (2024) 3(1) *Digital Society* 5.

potential of tort liability.[7] The Environmental Protection Agency, the Food and Drug Administration, and the Consumer Product Safety Commission have all asserted jurisdiction over nanotechnology applications. Nonetheless, in the absence of a coherent statutory mandate nanotechnology products regularly enter commerce under legal categories that were constructed for their bulk-form equivalents, trading short-term innovation incentives for long-term systemic vulnerability.[8]

The transatlantic divergence between these two approaches reflects two distinct, and in important respects opposed, theories of how law should relate to scientific knowledge. The European model, grounded in the precautionary tradition, insists that legal authorization must follow demonstrated safety, contending that the burden of proof lies on those who seek to introduce new substances into commerce, not on those who may be harmed by them.[9] The American model, informed by a pragmatist and market-oriented tradition, insists that legal prohibition must follow demonstrated harm and the costs of over-regulation, including foregone innovation and regulatory rigidity, are as real and as morally significant as the costs of under-protection.[10] Both positions have substantial intellectual and normative foundations. Yet both, as we demonstrate formally

[7] Cf. Catherine T Struve, 'The FDA and the Tort System: Postmarketing Surveillance, Compensation, and the Role of Litigation' (2005) 5 *Yale Journal of Health Policy, Law, and Ethics* 587 and Case Thomason, Larry Bucshon, Brian J Miller and Brian Yagi, 'The History of State Preemption and Medical Device Regulation: Lessons for Artificial Intelligence Oversight' (2026) 4(3) *Health Affairs Scholar* 1.

[8] See, in this perspective, Holly A Stretz, Andrew Kiss, Kathlyn N Mealio and Oluwaseyi Oluwasegun Ayeni, 'Regulatory and Environmental Issues of Nanotechnology Safety' in *Nanotechnology Safety* (Elsevier 2025) 59.

[9] This principle applies to many sectors, from food safety and air pollution to chemicals and it is strictly dependent on risk perception and the peculiar European preferences of policy makers. See, David Vogel, *The Politics of Precaution: Regulating Health, Safety, and Environmental Risks in Europe and the United States* (Princeton University Press 2012) 252 ff., in particular.

[10] See Jonathan B Wiener, 'The Real Pattern of Precaution' in James Hammit, Michael Rogers, Peter Sand and Jonathan B Wiener (eds), *The Reality of Precaution: Comparing Risk Regulation in the United States and Europe* (Routledge 2013).

in this paper, are optimal only under conditions that nanotechnology systematically violates. The European model is efficient when irreversibility of harm is high and the rate of information accumulation is low, according to a condition under which waiting costs are low relative to the option value of additional safety knowledge. The American model is efficient when harms are compensable and information accumulates rapidly through market experience, meaning that post-market correction should be feasible before damage becomes systemic. Nanotechnology sits in neither paradigm: potential harms may be both irreversible and slow to become scientifically observable, while information about risk accumulates neither through the pre-market review nor through tort litigation.

The central contribution of this paper is to develop a formal law and economics model of dynamic regulatory licensing that is suited to precisely this intermediate case. Our framework departs from existing models of regulation under uncertainty in three principal respects. First, we take seriously the distinction between situations in which a probability distribution over outcomes is available and situations in which no such distribution can be reliably specified. We model the regulator not as a Bayesian expected utility maximizer updating a prior over known states, but as an agent operating under ambiguity, facing a set of plausible probability distributions, none of which can be identified as the correct one on the basis of available evidence. Second, we model the regulatory decision as a dynamic optimal stopping problem in which the regulator must decide, at each point in time, whether to maintain, modify, or revoke a conditional license, based on the information accumulated through mandatory monitoring. We further show that once the realistic administrative costs of granting and revoking

authorization are incorporated into this problem, the optimal licensing rule takes the form of a two-threshold hysteresis band rather than a single, frictionless cutoff. Third, we model the contractual dimension of licensing, as an instrument that shapes the innovator's incentives to disclose information truthfully over the life of the license.

We develop this dimension in two steps. We first characterize, in a minimal static extension, the optimal menu of conditional licenses that induces truthful revelation of a binary risk type, showing that the socially optimal mechanism involves a tradeoff between monitoring intensity and the innovator's net return from licensure that varies systematically across risk types, with low-risk technologies receiving undistorted, comparatively light-touch treatment and high-risk technologies monitored below their own first-best level in order to limit the information rent that truthful revelation requires. We then indicate how this static screening logic could in principle be embedded into the fully dynamic licensing problem – so that the innovator's report interacts with the continuously evolving ambiguity state over the life of the license – and leave the complete dynamic mechanism, with a richer type space and a fully specified revelation game, for future research building on the framework developed here. These results have direct implications for the design of regulatory processes in both the European and American contexts.

Beyond these individual contributions, the paper advances the broader theoretical claim that, in high-risk, high-potential technological domains characterized by deep uncertainty and learning-by-doing, the most efficient form of legal regulation is neither the most restrictive nor the most permissive, but the most responsive. A legal system that institutionalizes learning – that embeds revision into the legal rule itself – is not a

weaker system than one that imposes fixed rules, being, instead, a more welfare-maximizing one. The concept of provisional legality that we develop is not a concession to the limits of knowledge but a principled response to them, converting uncertainty into a regulatory parameter.

The remainder of the paper proceeds as follows. Section 2 provides an overview of the regulatory landscape for nanotechnology in the European Union and the United States, surveying the principal legal frameworks and identifying the specific failure modes that our model is designed to address. Section 3 develops the formal framework, beginning with a self-contained exposition of decision theory under ambiguity – accessible to legal scholars without a technical economics background – and proceeding to the dynamic licensing model, the optimal stopping analysis, and the mechanism design results. Section 4 presents a numerical simulation of the dynamic regulatory learning process, comparing the European, American, and dynamic licensing regimes. Section 5 discusses the institutional implementation of the model, addressing questions of legislative-administrative delegation and the incentive-compatible production of regulatory information. Section 6 concludes. Throughout, nanotechnology serves as the empirical anchor for the analysis; the formal apparatus developed in Section 3, however, is deliberately general, and Section 6 returns to the question of its applicability beyond the nanoscale.

2. The Regulatory Landscape: Nanotechnology Governance in the European Union and the United States

Any serious analysis of nanotechnology regulation must begin with what is arguably the most fundamental and underappreciated difficulty in the entire field: the problem of defining, with sufficient legal precision, the object that regulation is meant to govern. This problem is not merely terminological but reflects a deeper tension between the categorical demands of legal language and the relational, context-dependent nature of nanoscale phenomena. The scientific community, in fact, has not converged on a single definition of "nanomaterial"[11] and the definitions that have been proposed for regulatory purposes differ from one another in ways that have significant practical consequences. Size-based definitions – typically referencing the one to one hundred nanometer range[12] – offer an administratively convenient criterion because it is observable, measurable, and capable of supporting legal certainty. Scientifically, however, size alone does not determine whether a material poses novel risks.[13] The properties that make nanomaterials commercially valuable – enhanced reactivity, increased surface area to volume ratio, quantum confinement effects, altered optical and electrical

[11] In these terms, see Max Boholm and Rickard Arvidsson, 'A Definition Framework for the Terms Nanomaterial and Nanoparticle' (2016) 10(1) *NanoEthics* 25.

[12] For an overview and discussion of the different definitions adopted by European and U.S. regulatory bodies, see Wolfgang G Kreyling, Manuela Semmler-Behnke and Qasim Chaudhry, 'A Complementary Definition of Nanomaterial' (2010) 5(3) *Nano Today* 165.

[13] To make an example, a nanoparticle of gold may be biologically inert, while a nanoparticle of silver exhibits potent antimicrobial activity that raises questions about ecological toxicity and resistance development.

characteristics[14] – are the properties that may make them biologically and environmentally hazardous, but those properties depend on shape, surface chemistry, solubility, aggregation behavior, and the physicochemical properties of the medium in which the material is dispersed.

This definitional instability has direct consequences for regulatory design. If the definition is too narrow – limiting, for example, coverage to particles with all three dimensions in the nanoscale – then regulatory frameworks may exclude tubular or platelet-shaped nanomaterials, such as carbon nanotubes or nanoclay platelets, that present serious toxicological concerns.[15] If the definition is too broad – encompassing any material with some nanoscale feature – it may impose compliance obligations on industries that have safely used nano-structured materials, such as carbon black in rubber or nanocrystalline cellulose in paper, for generations.

The following subsections survey the principal regulatory frameworks existing in the European Union and the United States, focusing on their historical construction, systematic architecture, and the shared failure modes that neither system has resolved. The unresolved normative problems identified below will motivate, in this perspective, the foundations of the formal model developed in the Section 3.

[14] See, in this light, Aamir Sohail, 'Progress in Nanomaterials: Synthesis, Characterization, and Applications' (2025) 8 *Next Nanotechnology* 1.

[15] Cf., for an overview, Amit Kumar, Ramna Tripathi and Akhilesh Kumar, 'Nanoparticle and Nanofillers: A Review on Types, Properties, Application, and Toxicological Impact' in Shadpour Mallakpour and Chaudhery Mustansar Hussain (eds), *Handbook of Nanofillers* (Springer 2025) 3273.

2.1 The European Union

2.1.1 The Constitutional Foundations and Early Institutional Development

The European Union's approach to nanotechnology governance is structured, at its deepest level, by the precautionary principle, which Article 191(2) of the Treaty on the Functioning of the European Union establishes as one of the foundational principles of Union environmental policy, and which the Court of Justice has progressively extended, through its case law, to encompass food safety, public health, and consumer protection.[16] The precautionary principle, as interpreted in European law, authorizes – and in some circumstances requires[17] – preventive action in the face of scientifically uncertain risks, before the causal relationship between an activity and a potential harm has been definitively established.[18] The European Commission Communication on the Precautionary Principle of February 2000 established the conditions under which the principle may be invoked: the identification of potentially negative effects based on available scientific data, a scientific evaluation that is incomplete or inconclusive, and a risk assessment establishing that the chosen level of protection may be exceeded. Crucially, the Communication did not reduce these conditions to a mere judgment of plausibility; it required a structured scientific assessment establishing that the existing

[16] See, for a prospective review, Robert Black, 'Risk, Precaution and Uncertainty' in *Legislating for Risk and Precaution: Bridging the Divide between Science and Law for Biosecurity* (CRC Press 2025)
[17] This is particularly true for the governance of technological progress, as shown in Andy Stirling, 'Precaution in the Governance of Technology' in Roger Brownsword, Eloise Scotford and Karen Yeung (eds), *The Oxford Handbook of Law, Regulation and Technology* (Oxford University Press 2017) 645.
[18] See, in this perspective, Noah M Sachs, 'Rescuing the Strong Precautionary Principle from Its Critics' (2011) *University of Illinois Law Review* 1285.

evidence base, however incomplete, supports a finding that protective intervention is warranted.[19]

The institutional translation of these principles into nanotechnology-specific policy began gradually. The European Commission's 2004 Communication "*Towards a European Strategy for Nanotechnology*" acknowledged the transformative potential of nanoscience and nanotechnology while calling for an integrated, responsible approach that incorporated safety, ethical, and environmental considerations from the earliest stages of research and development.[20] The subsequent *Action Plan for Nanosciences and Nanotechnologies of 2005-2009* committed to reviewing existing regulatory frameworks for their adequacy in addressing nanomaterials, developing new standardized testing and characterization methodologies, promoting additional international coordination, while the 2008 Communication on "*Regulatory Aspects of Nanomaterials*," acknowledged significant gaps in implementation capacity and scientific knowledge that limited the effectiveness of that legislation in practice.[21] From the perspective of the present paper, the importance of these definitional debates lies less in the precise numerical thresholds than in the fact that they reveal a deeper institutional inability to identify the relevant object of regulation in a stable manner.

---

[19] For further discussion, see Katherine Blyth Allan, 'Commentary on the European Commission's Communication on the Precautionary Principle' (2000) 34 *The Law Teacher* 215.

[20] For a useful overview, see Henk ten Have, 'Nanotechnology and Ethics – European Public Policies' in Bert Gordijn and Anthony Mark Cutter (eds), *In Pursuit of Nanoethics* (Springer 2014) 193.

[21] For a detailed examination, cf. Camilo Fautz, Torsten Fleischer, Ying Ma, Miao Liao and Amit Kumar, 'Discourses on Nanotechnology in Europe, China and India' in Miltos Ladikas, Sachin Chaturvedi, Yandong Zhao and Dirk Stemerding (eds), *Science and Technology Governance and Ethics* (Springer 2015) 125 and Diana M Bowman, 'More than a Decade On: Mapping Today's Regulatory and Policy Landscapes Following the Publication of Nanoscience and Nanotechnologies: Opportunities and Uncertainties' (2017) 11(2) *NanoEthics* 169.

2.1.2 The Definition of Problem: From the 2011 Recommendation to the 2022 Revision

The most consequential European Union-level action on the definitional problem was Commission Recommendation 2011/696/EU, which defined a "nanomaterial" as a "*natural, incidental or manufactured material consisting of solid particles that are present, either on their own or as identifiable constituent particles in aggregates or agglomerates*, *where 50 % or more of these particles in the number-based size distribution"* and one or more external dimensions is in the size range one to one hundred nanometers[22]. The Recommendation represented the first attempt to provide a harmonized, cross-sectoral definition of the regulated object, but included nanostructures in addition to nanoparticles, creating confusion and extensive critical commentary.[23] The fifty-percent threshold was criticized as arbitrary, unsupported by any identified toxicological threshold effect at that percentage level, and potentially both over- and under-inclusive depending on the material and application in question. More fundamentally, the Recommendation's purely dimensional approach failed to incorporate functional criteria related to the novel properties that nanoscale dimensions confer – the very properties that generate both commercial value and regulatory concern.

[22] The original definition mentioned that "*(a) one or more external dimensions of the particle are in the size range 1 nm to 100 nm; (b) the particle has an elongated shape, such as a rod, fibre or tube, where two external dimensions are smaller than 1 nm and the other dimension is larger than 100 nm; (c) the particle has a plate-like shape, where one external dimension is smaller than 1 nm and the other dimensions are larger than 100 nm*."

[23] See, in this regard, Christian Micheletti and Marco Venturini, 'Regulation and Safety: Criticism of Application to Nanomaterials' (2015) 14(3) *Nutrafoods* 169.

Following an extensive review process and stakeholder consultation, the Commission replaced the 2011 Recommendation with Recommendation 2022/C 229/01. The revision introduced several substantive modifications. The fifty-percent number-based threshold was retained, but the flexibility clause that had permitted sector-specific instruments to lower that threshold to as little as one percent was eliminated in the interest of cross-sectoral harmonization. The treatment of particles with non-standard geometries was clarified: elongated particles with two external dimensions below one nanometer and a length above one hundred nanometers, and plate-like particles with one external dimension below one nanometre and lateral dimensions above one hundred nanometres, are now explicitly covered, addressing the ambiguity that had previously surrounded carbon nanotubes and graphene-related materials. The volume-specific surface area criterion was restructured: whereas the 2011 Recommendation had used a VSSA above sixty square metres per cubic centimetre as a positive proxy for nanomaterial classification, the 2022 Recommendation replaced this with a VSSA below six square metres per cubic centimetre as an exclusion criterion, providing a simpler and less technically demanding screening tool for materials that clearly do not exhibit nanoscale behaviour regardless of their particle size distribution. Finally, an upper size limit was introduced, providing that particles with at least two orthogonal external dimensions larger than one hundred micrometres need not be counted for the purposes of the size distribution analysis. [24]

[24] For a comparison between the two recommendations see Hubert Rauscher, Andrej Kobe, Vikram Kestens and Kirsten Rasmussen, 'Is It a Nanomaterial in the EU? Three Essential Elements to Work It Out' (2023) 49 *Nano Today* 1.

2.1.3 Sector-Specific Instruments

The most operationally developed European responses to nanotechnology have emerged through progressive adaptation of sector-specific regimes governing products associated with heightened exposure risk and direct implications for human health. Three instruments merit particular attention for the purposes of the present analysis.

Regulation (EC) No 1223/2009 on cosmetic products, applicable since July 2013, represents the most comprehensive sectoral instrument specifically addressing nanomaterials at Union level. Its definition of nanomaterial is narrower than the Commission Recommendation, being limited to insoluble or biopersistent and intentionally manufactured materials with one or more external dimensions on the scale from one to one hundred nanometres,[25] with a formulation that deliberately focuses regulatory scrutiny on the materials considered most likely to raise toxicological concerns. [26] Article 16 establishes a pre-market notification mechanism requiring manufacturers to notify the Commission at least six months before placing nanomaterial-containing products on the market, providing detailed technical information on the nanomaterial's identity, specifications, toxicological profile, and conditions of safe use. Nanomaterials used as colorants, preservatives, or UV-filters remain subject to the stricter positive-list authorization procedures under the relevant Annexes, with the Scientific Committee on Consumer Safety having issued opinions on

[25] See Diana Bowman, Joel D'Silva and Geert Van Calster, 'Defining Nanomaterials for the Purpose of Regulation within the European Union' (2010) 1(2) *European Journal of Risk Regulation* 115.

[26] This criticism is reported, *inter alia*, in Umberto M Musazzi, Valentina Marini, Antonella Casiraghi and Paola Minghetti, 'Is the European Regulatory Framework Sufficient to Assure the Safety of Citizens Using Health Products Containing Nanomaterials?' (2017) 22(6) *Drug Discovery Today* 870.

nano-titanium dioxide, nano-zinc oxide, nano-silica, carbon black, and hydroxyapatite, among others.

Regulation (EC) No 1907/2006 (REACH) represents the foundational instrument of European chemicals governance and, in principle, the broadest framework for nanoscale substance regulation. In its original formulation, nonetheless, REACH contained no nano-specific provisions, and the standard registration requirements – designed for conventional chemicals and largely dependent on bulk-form data – were widely regarded as inadequate for capturing the distinctive properties of nanomaterials.[27] Scientific advisory bodies associated with the European Chemicals Agency repeatedly concluded that registration dossiers submitted for nanomaterials systematically lacked the data necessary for meaningful risk assessment.[28] These criticisms produced Commission Regulation (EU) 2018/1881, which amended the REACH Annexes to introduce nano-specific requirements applicable since January 2020, including detailed characterization data on particle size distribution, surface area, surface treatment, crystalline structure, and dustiness, as well as separate chemical safety assessments and exposure scenarios for each registered nanoform.[29] In this framework, while the European Chemicals Agency (ECHA) has developed extensive supporting guidance – e.g., the Guidance on the Registration of Nanomaterials and the nanoform-specific Appendix to the Guidance

[27] See, in particular, Elen Stokes, 'Regulating Nanotechnologies: Sizing Up the Options' (2009) 29(2) *Legal Studies* 281.

[28] See Christina Rudén and Sven Ove Hansson, 'Registration, Evaluation, and Authorization of Chemicals (REACH) Is but the First Step–How Far Will It Take Us? Six Further Steps to Improve the European Chemicals Legislation' (2010) 118(1) *Environmental Health Perspectives* 6.

[29] For further analysis, see Jan Labuda, Jiří Barek, Zuzana Gajdosechova, Heidi Goenaga-Infante, Linda J Johnston, Zoltan Mester and Sergei Shtykov, 'Analytical Chemistry of Engineered Nanomaterials: Part 1. Scope, Regulation, Legislation, and Metrology (IUPAC Technical Report)' (2023) 95(2) *Pure and Applied Chemistry* 133.

on Information Requirements and Chemical Safety Assessment – the practical implementation of reformed obligations continues to reveal difficulties, given that required characterization methodologies remain technically demanding and only partially standardized, generating substantial compliance burdens for registrants.[30]

Regulation (EU) No 10/2011 on food contact plastics, as amended by Regulation (EU) 2020/1245, applies the traditional positive-list model to substances used in food contact applications while introducing the analytically significant principle that authorization of a conventional form does not automatically extend to the nanoform of the same substance.[31] This explicit rejection of assumed equivalence between bulk and nanoscale forms represents a normative commitment to the position that nanoscale transformation is legally material. The European Food Safety Authority has developed detailed guidance on the risk assessment of food contact nanomaterials, acknowledging substantial methodological difficulties in assessing nanomaterial migration – in light of recognition that the toxicological and physicochemical properties of a substance may change substantially when produced at nanoscale dimensions[32] – and the limitations of

[30] On these aspects, see Shahriar Sharifi, Hossein Hosseinkhani, Alireza Salihi, Farhad Hosseinkhani, Hamed Jouyban-Gharamaleki, Javad Bazaz, Volker Mailänder, Twan Lammers, Morteza Mahmoudi and Omid Akhavan, 'Importance of Standardizing Analytical Characterization Methodology for Improved Reliability of the Nanomedicine Literature' (2022) 14(1) *Nano-Micro Letters* 172, 178.

[31] See, for further reference, Nicola Barrett, 'European Union Regulation of Nanotechnology in the Food Industry' (2011) 8 *Nanotechnology Law & Business* 252.

[32] Cf., in particular, Mohmmad Younus Wani, Mohd Ali Hashim, Firdosa Nabi and Maqsood Ahmad Malik, 'Nanotoxicity: Dimensional and Morphological Concerns' (2011) *Advances in Physical Chemistry* 1; Peter P Fu, Qingsu Xia, Huey-Min Hwang, Paresh C Ray and Hongtao Yu, 'Mechanisms of Nanotoxicity: Generation of Reactive Oxygen Species' (2014) 22(1) *Journal of Food and Drug Analysis* 64; Eşref Demir, 'A Review on Nanotoxicity and Nanogenotoxicity of Different Shapes of Nanomaterials' (2021) 41(1) *Journal of Applied Toxicology* 118.

existing test methods when applied to nanoscale materials under physiological conditions.[33]

Several structural problems in the European framework remain unresolved and are analytically significant for the model which will be developed in Section 3. The first, and most fundamental, is the coexistence of multiple, divergent sectoral definitions, creating a normative framework in which the same material may simultaneously qualify as a nanomaterial for one regulatory purpose and not for another. This incoherence generates genuine regulatory arbitrage, unequal treatment of substantially equivalent risks across product categories, and obstacles to the mutual recognition of safety assessments. The second problem concerns the static character of authorization decisions, which are calibrated to the state of scientific knowledge at the time of initial evaluation and do not automatically adapt as that knowledge develops. This static quality might be particularly costly in a domain characterized by rapid knowledge accumulation, where the appropriate regulatory response may change substantially between initial authorization and the steady-state deployment of a technology at scale.

A third and significant problem is the systematic underinvestment in post-market surveillance: the collection of real-world exposure data, the monitoring of adverse effects in deployed populations, and the systematic updating of risk assessments as a condition of continued authorization are either absent or insufficiently operationalized

[33] See, in this light, Iwona Wrześniewska-Wal, Beata Dominiak, Agnieszka Wal and Marian A Gralak, 'Regulatory Challenges of Nanomaterials in Food and Feed' (2025) 54(4) *Acta Alimentaria* 576.

across all regulatory instruments.[34] Despite their institutional differences, all regulatory instruments attempt to incorporate nanotechnology into pre-existing legal structures by progressively adapting definitions, testing requirements, and authorization procedures. None, however, fundamentally modifies the temporal architecture of regulation. Authorization remains largely static, whereas scientific knowledge evolves dynamically.

2.2 The United States

2.2.1 Institutional Architecture and the Preference for Regulatory Adaptation

The United States has not enacted any nanotechnology-specific legislation. The National Nanotechnology Initiative (NNI), established in 2000, coordinates research and development across dozens of participating agencies and academia, with substantial scientific progresses that still need to find their place in public awareness of nanotechnologies[35]. The NNI does not constitute a regulatory authority nor does it creates binding requirements for nanomaterials, which remain under-regulated by few bills that have been concretely passed in the US Congress, despite large amount of proposal that have been discussed by this legislative body.[36] The deliberate choice not

[34] See, in this regard, among others, Bharti Mangla, Pankaj Kumar, Shamama Javed, Tabish Pathan, Waquar Ahsan and Geeta Aggarwal, 'Regulating Nanomedicines: Challenges, Opportunities, and the Path Forward' (2025) 20(15) *Nanomedicine* 1911.
[35] For a comprehensive discussion, see Mihail C Roco, 'National Nanotechnology Initiative at 20 Years: Enabling New Horizons' (2023) 25(10) *Journal of Nanoparticle Research* 197.
[36] See, for further references, Nor Akhmal Hasmin, Juan Matmin and Sayidah Asma Basir, 'National Nanotechnology Policies in Malaysia, European Union and United States: A Comparative Analysis' (2022) 3(3) *Quantum Journal of Social Sciences and Humanities* 72.

to enact technology-specific legislation might reflect the dominant American regulatory philosophy of adapting existing frameworks to novel technologies, preserving institutional continuity and avoiding the rigidity associated with premature statutory classification.[37] While this approach has genuine theoretical advantages in conditions of rapidly evolving scientific understanding, it presents significant disadvantages in the nanotechnology context because existing frameworks might be designed for substances and properties that nanomaterials may not share.

2.2.2 The Toxic Substances Control Act and the Environmental Protection Agency

The Toxic Substances Control Act (TSCA), enacted in 1976 and reformed in 2016, provides the primary federal statutory authority for the regulation of industrial chemicals, including nanomaterials, requiring manufacturers and importers of new chemical substances – those not already present in the TSCA *Chemical Substance Inventory* to submit a Pre-Manufacture Notice to the United States Environmental Protection Agency (EPA) at least ninety days before commencing manufacture or importation, allowing the Agency to evaluate whether the substance presents an unreasonable risk to health or the environment.[38] The statutory framework of the TSCA

[37] Cf., *inter alia*, Nathan Cortez, 'Regulating Disruptive Innovation' (2014) 29(1) *Berkeley Technology Law Journal* 175; Ryan Hagemann, Jennifer Huddleston Skees and Adam Thierer, 'Soft Law for Hard Problems: The Governance of Emerging Technologies in an Uncertain Future' (2018) 17(1) *Colorado Technology Law Journal* 37; Sidney A Shapiro and Thomas O McGarity, 'Not So Paradoxical: The Rationale for Technology-Based Regulation' (1991) 1991(3) *Duke Law Journal* 729.

[38] See, in this regard, J Vincent Nabholz, 'Environmental Hazard and Risk Assessment under the United States Toxic Substances Control Act' (1991) 109–110 *Science of the Total Environment* 649.

was widely considered “excessively burdensome”[39] and the major difficulty in the application to nanomaterials has been whether they constitute a "new chemical substance" for statutory purposes, given the difficulty in assessing this criterion when the “nano” version is the same substance as the corresponding bulk form already listed on the Inventory.[40] A second inherent limitation of TSCA is that it discourages manufacturers from voluntarily sharing information, since the statute effectively treated the absence of data as evidence of no risk.[41]

The subsequent 2016 reform of TSCA – the so-called Lautenberg Amendment or Lautenberg Act[42] – required manufacturers and processors of nanomaterials to provide advance notification before commencing new uses, introducing a mandatory risk evaluation process for existing chemicals on the Inventory.[43] Implementation of the reformed TSCA framework for nanomaterials remains, nevertheless, incomplete. The pre-manufacture notification process provides an opportunity for prospective evaluation, but the EPA's capacity to require testing or impose restrictions under the law is constrained by the statutory standard of unreasonable risk, a standard that, given the current state of nano-toxicological knowledge, is difficult to satisfy for many novel

[39] Cit. Charles W Schmidt, 'TSCA 2.0: A New Era in Chemical Risk Management' (2016) 124(10) *Environmental Health Perspectives* 183, 183.
[40] For a detailed discussion, cf. Christopher J Preston, Michael Y Sheinin, David J Sproat and Vijay P Swarup, 'The Novelty of Nano and the Regulatory Challenge of Newness' (2010) 4(1) *NanoEthics* 13 and John C Monica Jr and John C Monica, 'Examples of Recent EPA Regulation of Nanoscale Materials under the Toxic Substances Control Act' (2009) 6 *Nanotechnology Law & Business* 388.
[41] This issue is raised, among others, in George A Kimbrell, 'Governance of Nanotechnology and Nanomaterials: Principles, Regulation, and Renegotiating the Social Contract' (2009) 37(4) *Journal of Law, Medicine & Ethics* 706.
[42] See, in this perspective, Sarah E Mische, 'The Cost of PFAS Cleanup in Waterways: Who Pays and How' (2025) 96(3) *University of Colorado Law Review* 906.
[43] or an in-depth discussion of these requirements, see Patricia D Koman, Veena Singla, Juleen Lam and Tracey J Woodruff, 'Population Susceptibility: A Vital Consideration in Chemical Risk Evaluation under the Lautenberg Toxic Substances Control Act' (2019) 17(8) *PLOS Biology* 1.

nanomaterials before they have generated market-based exposure data.[44] The practical result is a regulatory regime that operates as a *de facto* permissive system in which most nanomaterials reach commerce and working environments subject only to voluntary risk management commitments.[45]

2.2.3 The Food and Drug Administration: Multi-Framework Governance

The *Food and Drug Administration* (FDA) regulates nanotechnology applications across multiple categories under a corresponding multiplicity of authorities and significant inter-categorical variation in stringency, according to guidance documents with nano-specific considerations on food, cosmetics, animal food, and, with some limitations, nanomedicine.[46] In cosmetics, the FDA's approach rests entirely on voluntary industry engagement and post-market enforcement.[47] The 2014 Guidance for Industry on Safety of Nanomaterials in Cosmetic Products acknowledged the absence of specific regulations governing nanomaterials in cosmetics and encouraged manufacturers to conduct appropriate safety testing and to consult the Agency voluntarily before marketing. The *Modernization of Cosmetics Regulation Act of 2022*

[44] See, *ex multis*, Martha E Marrapese, Sara Beth Watson and Sarah E Amick, 'What the "New TSCA" Has Accomplished–and What Comes Next' (2026) 40(4) *Natural Resources & Environment* 11.
[45] On these aspects, see Preetika Biswas and Ashutosh Yadav, 'Regulations for Safety Assessment of Nanomaterials' in Vineet Kumar, Nandita Dasgupta and Shivendu Ranjan (eds), *Environmental Toxicity of Nanomaterials* (CRC Press 2018) 447.
[46] This point is raised in Rachel Foulkes, Ernest Man, Jasmine Thind, Suet Yeung, Abigail Joy and Clare Hoskins, 'The Regulation of Nanomaterials and Nanomedicines for Clinical Application: Current and Future Perspectives' (2020) 8(17) *Biomaterials Science* 4653.
[47] See, for further details, Renee Stover, Prashiela Manga, Elizabeth H Anderson and Linda M Katz, 'The Food and Drug Administration and Modernization of US Cosmetic Regulations' in Peter Elsner, Shari R Lipner and Howard Maibach (eds), *Cosmetics: Controlled Efficacy Studies and Regulation* (Springer Nature 2026).

strengthened FDA authorities in several respects – introducing adverse event reporting requirements, facility registration, and safety substantiation obligations – but did not establish pre-market disclosure or safety evaluation requirements for nanomaterial-containing cosmetics.[48]

In food and food ingredients, the "generally recognized as safe" (GRAS) doctrine implies that a substance is intended “safe” when there is a reasonable certainty in the minds of competent scientists that the substance is not harmful under the intended conditions of use.[49] The problem that has been raised with the GRAS doctrine is that manufacturers may self-determine GRAS status and are not legally required to notify the FDA of such determinations, and nanomaterials, despite being widely regarded as potentially risky due to concerns about bodily accumulation and immune reactions and lacking scientific consensus on their safety, appear to have entered the food supply through undisclosed GRAS self-determinations used in both food products and packaging.[50] Among the many issues determined by this regulatory posture, is that nanotechnology has become pervasive throughout the food chain, with engineered nanomaterials incorporated into agricultural inputs, food packaging, storage,

[48] This issue has been pinpointed, among others, in Simone Swafford, 'Cosmetics' in Eunjoo Pacifici and Susan Bain (eds), *An Overview of FDA Regulated Products: From Drugs and Cosmetics to Food and Tobacco* (2nd edn, Academic Press 2025) 257.

[49] See, on this matter, Laurie J Beyranevand, 'Generally Recognized as Safe?: Analyzing Flaws in the FDA’s Approach to GRAS Additives' (2013) 37 *Vermont Law Review* 887.

[50] See, in this perspective, Sheetal R D'Mello, Celia N Cruz, Mei-Ling Chen, Mamta Kapoor, Sau L Lee and Katherine M Tyner, 'The Evolving Landscape of Drug Products Containing Nanomaterials in the United States' (2017) 12(6) *Nature Nanotechnology* 523.

transportation, and processing, resulting in widespread and often unavoidable consumer exposure.[51]

Finally, in medicine, the regulatory picture includes pre-market requirements applicable to pharmaceutical products. The FDA advocates a risk-based regulatory approach requiring comprehensive characterization of nanomaterials and assessment of factors such as their physicochemical properties, biological behavior, manufacturing processes, stability, and intended use to ensure product quality, safety, and efficacy.[52]

The FDA complements premarket evaluation with post-market surveillance and risk management, while collaborating through the NNI to develop standardized assessment methods, advance safety research, and support the safe and effective translation of nanomedicines into clinical practice.[53] Nonetheless, FDA oversight of nanotechnology remains product-based rather than technology-based, allowing some nano-enabled products to rely on safety assessments conducted on their conventional counterparts instead of undergoing dedicated premarket evaluation.[54]

Having explored the American regulatory framework on nanomaterials, we might say that the US governance systems presents a set of structural limitations that, while distinct

[51] For additional details, see Ilise L Feitshans, 'Law and Policy of Nanotechnology in Food: Global Commerce Promoting Global Health from Field to Fork' (2025) 9(1) *Georgetown Law Technology Review* 123.

[52] See, in this regard, Nimeet Desai, Dhwani Rana, Mitali Patel, Neha Bajwa, Rajendra Prasad and Lalitkumar K Vora, 'Nanoparticle Therapeutics in Clinical Perspective: Classification, Marketed Products, and Regulatory Landscape' (2025) 21(29) *Small* 1.

[53] Cf. Neha Raghuwanshi, Vinita Isad, R K Shanmathi and Chaitanya Kale, 'Bridging the Gap: Regulatory Pathway and Clinical Translation of Nanomaterials in Medicines' in *Smart Nanomaterials in Biomedical Applications* (Springer 2025) 305; Snežana Đorđević, María Medel Gonzalez, Inmaculada Conejos-Sánchez, Barbara Carreira, Sabina Pozzi, Rita C Acúrcio, Ronit Satchi-Fainaro, Helena F Florindo and María J Vicent, 'Current Hurdles to the Translation of Nanomedicines from Bench to the Clinic' (2022) 12(3) *Drug Delivery and Translational Research* 500.

[54] This issue is raised, *ex plurimis*, in Raj Bawa, 'Regulating Nanomedicine–Can the FDA Handle It?' (2011) 8(3) *Current Drug Delivery* 227.

in form, mirror in important respects the weaknesses observed in the European system. Central among these is the reliance on voluntary guidance in product domains where exposure risks are non-trivial, coupled with the absence of mandatory pre-market evaluation, which precludes the systematic generation of safety data at the stage where it would be most consequential, before large-scale market diffusion entrenches economic and institutional resistance to later intervention. A second and closely related limitation lies in the application of the GRAS framework to nanoforms, where self-determined safety classifications, optional notification, and limited agency review produce a persistent informational deficit; although the FDA maintains that conventional GRAS determinations do not automatically extend to nanomaterials, this position lacks practical enforceability without compulsory disclosure mechanisms capable of ensuring regulatory visibility over such determinations. Finally, the system's reliance on post-market correction shifts the burden of evidence generation onto market experience, an approach that is poorly suited to nanomaterial-related harms that are often chronic, delayed, or difficult to causally attribute – such as long-term respiratory effects, carcinogenic risks, or ecological accumulation – limiting the effectiveness of retrospective enforcement even where it is formally available.

3. A Dynamic Theory of Regulatory Licensing under Uncertainty

Having outlined the structural limitations of the European and American regulatory framework, it becomes evident that the combination of scientific ambiguity, delayed

harm manifestation, and imperfect enforceability implies that regulatory decisions are taken under conditions in which probabilities are difficult to specify and future states of knowledge remain fundamentally open. This motivates a shift in analytical perspective from descriptive institutional critique to formal modelling. Accordingly, this Section develops a dynamic theory of regulatory licensing under uncertainty, aimed at capturing how regulators optimally update approval and restriction decisions when both risk magnitudes and informational quality evolve over time in a non-quantifiable manner.

3.1 The Economic Environment

The comparative analysis developed in the previous Section revealed a structural limitation shared by the two dominant regulatory paradigms governing emerging technologies. The European precautionary approach and the American post-market approach are usually presented as reflecting different normative commitments regarding innovation, public health, and the role of scientific uncertainty in regulatory decision-making.[55] From a law and economic perspective, however, both approaches may be interpreted as alternative solutions to the same underlying optimization problem, seeking to maximize social welfare under incomplete scientific knowledge. We formalize this optimization problem and derive the welfare-maximizing regulatory mechanism under uncertainty, demonstrating that neither the traditional precautionary model nor the conventional post-market model might be generally optimal. Instead, an optimal institutional arrangement would take the form of a dynamic licensing regime in

[55] Reshma Leslie, 'Strict Liability or Fault-Based Regimes for AI-Caused Harm? A Doctrinal Analysis Across Common Law and Civil Law Systems' (2024) 52(4) *Power System Protection and Control* 146.

which regulatory authorization evolves continuously together with the regulator's information.

3.1.1 Agents and States of Nature

The model developed below is deliberately formulated at a level of generality that abstracts from the specific physical or chemical mechanisms through which nanoscale materials generate risk. This choice is not a simplification of convenience but reflects the paper's central claim: that the regulatory problem posed by nanotechnology is, at its core, an informational and institutional problem common to any technology characterized by ambiguity and endogenous learning, and that its formal treatment need not, and should not, be tied to nanotechnology-specific parameters in order to be useful in that domain. Consider an economy populated by three categories of agents. The first is a *social regulator*, representing society as a whole, assumed to maximize intertemporal social welfare. We treat the regulator as a benevolent social planner, to focus on the informational structure of the regulatory problem rather than on questions of political economy, bureaucratic incentives, or regulatory capture. The second agent is the *innovator*, alias a profit-maximizing firm that has developed a commercial application based upon engineered nanomaterials and seeks regulatory authorization to introduce the technology into the market. The innovator receives private benefits from commercialization but bears the costs imposed by regulatory compliance, including monitoring obligations, reporting requirements, and possible restrictions on market access. The third category consists of *society*, understood as the collection of individuals

whose welfare is affected by the deployment of the technology.[56] Society includes consumers who benefit from innovation, workers exposed during production processes, patients receiving medical applications, downstream firms incorporating nanomaterials into production chains, and members of the public who may be indirectly exposed through environmental pathways. Throughout the model these heterogeneous interests are aggregated into a single social welfare function. This simplification allows the analysis to concentrate on the informational problem created by scientific uncertainty without introducing distributive considerations that are orthogonal to the principal argument.

The regulator must choose a licensing policy before the scientific consequences of commercial deployment are fully understood. To formalize this situation, let $(\Omega, \mathcal{F})$ denote a measurable state space, in which Ω is a space of fundamental states of nature and each ω encodes all payoff-relevant features, although the mapping between ω and observable signals may itself be subject to model uncertainty. Each element $\omega \in \Omega$ represents a complete description of the physical world relevant to the regulatory decision.[57] Once the true state is realized, every uncertainty relevant to the regulator is resolved and each state determines, simultaneously: the biological mechanisms through which adverse effects may arise; the exposure

[56] These include, among others, consumers benefiting from innovation, workers exposed during production processes, patients receiving medical applications, firms incorporating nanomaterials into production chains, and members of the public who may be indirectly exposed through environmental pathways. Throughout the model heterogeneous interests are aggregated into a single social welfare function, concentrating the attention on the informational problem created by scientific uncertainty without introducing distributive considerations that are orthogonal to the principal argument.

[57] Importantly, a state of nature should be interpreted, in this context, as a complete scientific description of reality and not a single observable variable such as toxicity or environmental persistence.

pathways affecting workers, consumers, and ecosystems; the effectiveness of available mitigation measures; the long-term evolution of environmental accumulation; other scientific characteristic capable of influencing the social consequences of commercial deployment. In this context, uncertainty arises because the regulator does not know which element of Ω corresponds to that structure at the moment the licensing decision must be taken, while $\mathcal{F}$ ensures that probabilities can eventually be assigned consistently to relevant events whenever a probabilistic representation becomes appropriate.

3.1.2 Regulatory Actions and Social Consequences

The regulator chooses a regulatory mechanism whose consequences depend upon the state of nature that eventually materializes. Let $\mathcal{A}$ denote the set of admissible regulatory actions. An element $a \in \mathcal{A}$ specifies an entire regulatory policy rather than a single administrative decision. In the context of nanotechnology, a regulatory action should include at least four distinct dimensions. First, the regulator determines whether commercial deployment is authorized. Second, if authorization is granted, the regulator specifies the monitoring obligations imposed upon the innovator, including data collection requirements, reporting frequency, post-market studies, and surveillance protocols. Third, the regulator determines the conditions under which the authorization will subsequently be reviewed, modified, suspended, or revoked. Finally, the regulator determines the information that the innovator must disclose throughout the life of the license.

Every regulatory action generates consequences that depend upon the true state of nature. Formally, let X denote the set of possible social outcomes. A regulatory action induces an outcome through a mapping $x : \mathcal{A} \times \Omega \rightarrow X$. The outcome produced by regulatory action a under state ω is therefore $x(a, \omega)$. The aggregate consequences of deployment can be represented through the welfare function $W : X \rightarrow \mathbb{R}$. The composition $W(x(a, \omega)$ hence measures the social value generated when regulatory policy a is implemented and the true state of nature is ω. Expected utility, in this light, assumes that the regulator possesses a unique probability distribution over Ω.

3.2 Information, Scientific Knowledge, and the Nature of Regulatory Uncertainty

The economic environment described before identifies the agents, the available regulatory actions, and the social consequences associated with each possible state of nature. It does not yet specify, however, what the regulator knows at the moment the licensing decision must be taken. The central proposition of this paper is that nanotechnology regulation differs from most traditional regulatory problems not because the potential harms are necessarily greater, but because the regulator's information possesses a qualitatively different structure. The regulatory problem, in this perspective, is epistemic and requires separating three concepts that are frequently conflated in both legal scholarship and policy discourse: the state of the world, scientific evidence, and probabilistic beliefs.

3.2.1 The State of the World and the Regulator's Information

We have anticipated that, at every point in time there exists a unique but unknown state $\omega \in \Omega$ and that the regulator never observes this state directly. Instead, regulatory decisions are based upon scientific evidence that provides only partial information regarding the underlying reality. Formally, let $(\mathcal{F}_t)_{t\geq 0}$ be a filtration on $(\Omega, \mathcal{F})$, where $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$ for every $t \geq 0$. This filtration represents the evolution of scientific knowledge through time. At each period t, $\mathcal{F}_t$ includes all information that has become available to the regulator by that point. Information may derive, among others, from laboratory experiments, toxicological studies, epidemiological investigations, environmental monitoring, post-market surveillance, computational modelling, or any other scientifically relevant source. The filtration captures the cumulative stock of knowledge available to the regulator. Given that scientific research is cumulative,[58] the information available tomorrow includes the information available today together with any newly generated evidence. Formally, $\mathcal{F}_t$ is increasing over time. Nevertheless, an expanding information set does not imply that uncertainty disappears. Additional information may clarify existing hypotheses, contradict earlier findings, reveal previously unknown mechanisms of harm, or even generate entirely new scientific questions, and learning should not be understood as a monotonic reduction in ignorance.

3.2.2 Signals and Scientific Evidence

[58] See Gideon D Markman, 'Will your Study Make the World A Better Place?' (2022) 59(6) *Journal of Management Studies* 1597, cit. 1599.

The regulator does not observe the state of nature directly. Instead, information arrives through observable signals. Let S denote the signal space. At each period t, the regulator observes a signal $s_t \in S$. Consequently, the sequence $(s_1, s_2, \dots, s_t)$ constitutes the regulator's scientific evidence at time t. In this framework, a signal represents any new body of information capable of modifying the regulator's scientific assessment and may consist of a peer-reviewed toxicological study, an epidemiological investigation reporting previously unknown health effects, new exposure measurements collected through workplace monitoring, environmental observations identifying unexpected persistence, or regulatory data generated through mandatory post-market surveillance. All informational content is heterogeneous and some studies provide more reliable evidence based upon larger samples or established methodologies,[59] while other may suffer from methodological limitations, conflicting results, or substantial measurement error. Formally, signals are generated according to an unknown data-generating process linking observable evidence to the underlying state of nature. This relationship is denoted by $\ell(s \mid \omega)$, where ℓ represents the likelihood of observing signals s conditional upon the true state ω. The likelihood function plays a familiar role in statistical inference and it measures how informative different observations are about the underlying scientific reality.[60] At this stage, however, no assumption is made regarding the

[59] Cf., *inter alia*, Andrew M McIntosh, Michael Sharpe and Stephen M Lawrie, 'Research Methods, Statistics and Evidence-Based Practice' in Eve C Johnstone, David Cunningham Owens, Stephen M Lawrie, Andrew M McIntosh and Michael Sharpe (eds), *Companion to Psychiatric Studies* (8th edn, Churchill Livingstone 2010) 157; Lu Liu, Benjamin F Jones, Brian Uzzi and Dashun Wang, 'Data, Measurement and Empirical Methods in the Science of Science' (2023) 7(7) *Nature Human Behaviour* 1046; Mark L Taper and Subhash R Lele (eds), *The Nature of Scientific Evidence: Statistical, Philosophical, and Empirical Considerations* (University of Chicago Press 2010).

[60] See, for further reference, *ex plurimis*, Yudi Pawitan, *In All Likelihood: Statistical Modelling and Inference Using Likelihood* (2nd edn, Oxford University Press 2026).

regulator's knowledge of the likelihood itself because, for emerging technologies and engineered nanomaterials, the mapping between observed evidence and underlying risk may itself be only partially characterized.

3.2.3 Risk, Statistical Uncertainty, and Model Uncertainty

The distinction between ordinary risk and Knightian uncertainty[61] might become clearer once the informational structure is specified. Assume first that the regulator knows the likelihood structure $\ell(s \mid \omega)$, which is uniquely specified and known. In such a case, uncertainty would concern only the realization of the state $\omega \in \Omega$. Scientific information, summarized by the filtration $(\mathcal{F}_t)_{t\geq 0}$, generates well-defined conditional probability measures $P(\cdot \mid \mathcal{F}_t)$ over $(\Omega, \mathcal{F})$. Beliefs would be therefore represented by a single posterior distribution. Given a regulatory action $\alpha \in A$, the induced social outcome is $x(a, \omega)$, and social welfare is given by $W(x(a, \omega))$. The regulator expected welfare at time t is, therefore, $\mathbb{E}[W \mid \mathcal{F}_t] = \int_\Omega W(x(a, \omega)) dP(\omega \mid \mathcal{F}_t)$. Whithin this framework, learning would consist of updating $P(\cdot \mid \mathcal{F}_t)$ through Bayesian learning[62] as new evidence – i.e., signal s – is observed through the likelihood $\ell(s \mid \omega)$. Conventional regulatory contexts can be approximated in this way and the inferential problem might

[61] Note that Knightian uncertainty refers to situations in which outcomes are known but probabilities cannot be assigned, in contrast to risk where both outcomes and probabilities are known, and it can be further distinguished into bounded uncertainty where probabilities lie within known intervals and pure uncertainty where no probabilistic bounds can be specified, forming a continuum from near-risk to complete indeterminacy. See, for further references, Cass R Sunstein, 'Knightian Uncertainty in the Regulatory Context' (2025) 9(3) *Behavioural Public Policy* 614.

[62] Given that probabilities are not fixed frequencies, but degrees of belief are conditional on available information, and therefore subject to revision, we might consider Bayesian learning as the process by which an agent updates probabilistic belief in response to new evidence in a logically consistent way. For a more detailed discussion of these aspects, see Robert A Jacobs and John K Kruschke, 'Bayesian Learning Theory Applied to Human Cognition' (2011) 2(1) *WIREs Cognitive Science* 8.

be related to parameter estimation within a fixed probabilistic model[63], relying on the existence of a unique probability measure over Ω.

By contrast, the structure of many nanotechnologies is fundamentally different and uncertainty, in many applications involving emerging technologies, might not be limited to unknown parameter values, extending to the validity of the probabilistic model itself and, in all these instances, the regulator might face uncertainty not only about outcomes but also about their probabilistic representation.[64] In this case, the mapping between states and observable signals is not pinned down by a single function $\ell(s \mid \omega)$, being a set of admissible likelihoods $\mathcal{L} = \{\ell_\alpha(s \mid \omega): \alpha \in A\}$, in which each ℓ_α induces, via Bayesian rule, a corresponding conditional probability measure $P_\alpha(\cdot \mid \mathcal{F}_t)$ over $(\Omega, \mathcal{F})$. The regulator, in this framework, would be characterized by a set of admissible posteriors $\mathcal{P}_t = \{P_\alpha(\cdot \mid \mathcal{F}_t) : \alpha \in A\}$ and, in this environment, expected social welfare would be no longer uniquely defined. Indeed, for each regulatory action $\alpha \in A$, the regulator would obtain a set of conditional expected welfare level

$$\mathcal{W}_t(a) = \{\int_\Omega W\big(x(a,\omega)\big) dP_\alpha(\omega \mid \mathcal{F}_t) : P_\alpha \in \mathcal{P}_t$$

And, consequently, the regulatory problem would become a mapping of a set of expected welfare values, which might be defined as

[63] For instance, in the regulation of pharmaceutical consequences, adverse-event probabilities can often be estimated with substantial confidence from extensive clinical data. Cf., in this regard, Fei He, Jiaxuan Liu, Xin Shen, Zijing Wang, Qiao Li and Guohui Li, 'Adverse Event Profile for Nanoparticle Albumin-Bound Paclitaxel Compared With Solvent-Based Taxanes in Solid-Organ Tumors: A Systematic Review and Meta-Analysis of Randomized Clinical Trials' (2022) 56(8) *Annals of Pharmacotherapy* 898.

[64] For instance, different toxicological mechanisms may imply distinct exposure–response functions and alternative models of nanoparticle behavior can generate divergent predictions regarding environmental persistence. Competing biological hypotheses may yield, in this perspective, materially different assessments of long-term risk, with available evidence often insufficient to discriminate between them.

$$\max_{a \in A} F(\mathcal{W}_t(a)).$$

3.3 The Expected Utility Benchmark and Its Domain of Validity

The informational structure developed in the previous Section was deliberately neutral with respect to the decision rule adopted by the regulator, specifying how information is generated and accumulated through the filtration $\mathcal{F}_t$ without imposing conditions on how such information is mapped into regulatory choices. In this Section we introduce subjective expected utility, providing a reference representation of choice under uncertainty, against which departures can be formally defined.

3.3.1 Preferences over Regulatory Acts

Given the environment described *supra*, a regulatory decision consists of the choice of an action $\alpha \in A$, which induces a state-contingent outcome $x(a, \omega) \in X$ and a state-contingent level of social welfare $W\big(x(a, \omega)\big)$. From an ex ante perspective, each regulatory action is identified within a function from states of nature to consequences such as the act associated with a can be defined as

$$f_\alpha : \Omega \to \mathbb{R}, \qquad f_\alpha(\omega) = W\big(x(a, \omega)\big)$$

The regulator's choice problem can thus be reformulated as the selection of an act from the set $\mathcal{F} = \{f_\alpha \mid a \in A\}$ and the central question would therefore be how these acts are ranked before the realization of the true state ω.

The ranking of such acts requires the specification of a preference relation over the set $\mathcal{F}$. Let $\succeq$ denote a weak preference relation defined on $\mathcal{F}$, where $f \succeq g$ indicates that

the regulator considers the regulatory act f at least as socially desirable as the regulatory act g, prior to the realization of uncertainty. The strict preference and indifference relations are defined in the usual way: $f \succ g$ if $f \succeq g$ and not $g \succeq f$, and $f \sim g$ if both $f \succeq g$ and $g \succeq f$. Since each act $f \in \mathcal{F}$ is a state-contingent mapping $f : \Omega \to \mathbb{R}$, the preference relation $\succeq$ is defined over uncertain prospects rather than deterministic outcomes. The regulator must therefore evaluate regulatory policies ex ante, before the realization of ω, on the basis of their entire distribution of consequences across possible states of nature.

To obtain a tractable representation of such interdependent evaluations, assume, in this benchmark subsection only, that the preference relation $\succeq$ satisfies the standard axioms of rational choice under uncertainty: completeness, transitivity, continuity, and the sure-thing principle.[65] These axioms impose internal consistency on choices across contingent events and ensure that preferences are separable with respect to common consequences in disjoint states. Under these conditions, preferences over acts admit the existence of a real-valued functional

$U : \mathcal{F} \to \mathbb{R}$ such that for any $f, g \in \mathcal{F}$, $f \succeq g \iff U(f) \geq U(g)$. The functional $U(\cdot)$ assigns to each regulatory act a scalar measure of its social desirability, aggregating its state-contingent consequences into a single ranking criterion. The structure of this

[65] See Jean-Yves Jaffray and Peter P Wakker, 'Decision Making with Belief Functions: Compatibility and Incompatibility with the Sure-Thing Principle' (1993) 7(3) *Journal of Risk and Uncertainty* 255. Note that Section 3.3.3 relaxes the sure-thing principle, replacing it with the weaker axioms of Gilboa and Schmeidler, which accommodate Ellsberg-type ambiguity aversion, to obtain the multiple-priors representation that underlies the remainder of the paper. See, in this regard, Itzhak Gilboa and David Schmeidler, 'Maxmin Expected Utility with Non-Unique Prior' (1989) 18(2) *Journal of Mathematical Economics* 141.

representation is further characterized by the Savage representation theorem,[66] with a unique subjective probability measure P over the measurable space $(\Omega, \mathcal{F})$ such that the functional U can be expressed in expected utility form.

To see this structure more clearly, note that since each regulatory act $f \in \mathcal{F}$ maps states of nature into welfare consequences, the evaluation of an act can be written as a function of its state-contingent values and this evaluation takes the form of a probability-weighted aggregation of such consequences across all states. Formally, there exists a probability measure P such that $U(f) = \int_{\Omega} f(\omega) dP(\omega)$. Given that each regulatory act is defined by $f_{\alpha}(\omega) = W\big(x(a, \omega)\big)$, the regulator's evaluation of a licensing decision can be written as $U(a) = \int_{\Omega} W(x(a, \omega)) dP(\omega)$.The regulatory choice problem is thus represented as

$$\max_{a \in A} \int_{\Omega} W(x(a, \omega)) dP(\omega).$$

establishing the benchmark case of decision-making under uncertainty, in which all relevant uncertainty is summarized by a unique subjective probability measure.

3.3.2 Limits of the Single-Prior Representation

The Savage representation is a classic and coherent benchmark for decision-making under uncertainty by reducing all informational limitations to a single subjective probability measure P. Within this framework, scientific learning is captured by Bayesian updating of beliefs, while the underlying probabilistic structure linking

[66] For further references, see Mohammed Abdellaoui and Peter P Wakker, 'Savage for Dummies and Experts' (2020) 186 *Journal of Economic Theory* 1.

evidence to outcomes is assumed to be stable and fully specified. Nevertheless, this informational structure might become increasingly restrictive in environments characterized by emerging technologies, such as engineered nanomaterials, where both the magnitude of potential harms and the mechanisms through which they arise remain only partially understood. In such contexts, uncertainty extends to the adequacy of the probabilistic model used to describe them.

In particular, the mapping between observable scientific evidence and underlying states of nature may not be uniquely determined. Different scientifically plausible models may generate distinct likelihood functions consistent with the same set of observed data, especially when evidence is sparse, heterogeneous, or subject to measurement error. In the specific case of engineered nanomaterials, this indeterminacy has an identifiable source: because toxicological and environmental behavior depends on shape, surface chemistry, aggregation state, and the medium of dispersion rather than on chemical identity alone (see Section 2, and the gold/silver nanoparticle contrast discussed therein), a single observed signal s – e.g., a toxicological study on one nanoform – is generally uninformative about the likelihood function governing a structurally similar but not identical nanoform. This is the formal counterpart of the "relational" rather than "intrinsic" property structure described in the Introduction: the object of regulation does not have a stable, context-independent mapping from state to signal, which is precisely the condition under which$\ell_\alpha(s \mid \omega)$. fails to be uniquely pinned down

Instead, there may exist a family of admissible likelihood structures, each compatible with the available scientific evidence and each inducing a different probabilistic

representation of future outcomes. In such a setting, Bayesian updating might no longer be sufficient and the assumption of a unique subjective probability measure P over $(\Omega, \mathcal{F})$ becomes too restrictive to capture the epistemic structure of regulatory decision-making. The regulator's beliefs can no longer be summarized by a single posterior distribution, since different admissible models may generate different posteriors even when conditioned on the same information set $\mathcal{F}_t$. This observation motivates a departure from the standard Savage framework.

3.3.3 Multiple Priors and Ambiguity

The failure of the single-prior representation stems from the fact that uncertainty – in the context of nanotechnologies, in particular – extends to the probabilistic structure linking scientific evidence to latent states of nature. Once the likelihood function is no longer uniquely determined, we might not always rely on a single Bayesian model to represent the evolution of beliefs and, in such an environment, uncertainty would be better represented as ambiguity over probabilistic models. Instead of a unique subjective probability measure P the regulator considers a set of admissible probability measures consistent with the available scientific evidence. Formally, for each information set $\mathcal{F}_t$., beliefs are represented by a correspondence $P_t : \mathcal{F}_t \rightrightarrows \Delta(\Omega)$, where $\Delta(\Omega)$, denotes the space of all probability measures over $(\Omega, \mathcal{F})$, and $P_t \subseteq \Delta(\Omega)$ is the set of priors that remain scientifically admissible given the information available at time t and each element $P \in \mathcal{P}_t$ corresponds to a distinct but observationally coherent model of how states of nature generate observable evidence.

Given this structure, the evaluation of regulatory acts can no longer be based on a single expected welfare criterion. Instead, each act $a \in A$ induces a set of expected welfare values, one for each admissible probability measure

$$\mathcal{W}_t(a) = \{\int_\Omega W\big(x(a,\omega)\big)dP(\omega) : P \in \mathcal{P}_t\}$$

The consequence is that the regulator is faced with a set-valued evaluation of each regulatory policy, reflecting the fact that different admissible models may yield different welfare assessments. To aggregate such a set of evaluations into a single decision criterion, a representation mapping the sets of expected welfare values into real numbers could take the form $V_t(a) = \mathcal{F}(\mathcal{W}_t(a))$, where $\mathcal{F}(\cdot)$ captures the regulator's attitude toward ambiguity and its different specifications correspond to different behavioral models, ranging from worst-case evaluation to smoother ambiguity-aversion functionals. The regulator evaluates each act according to its worst-case expected welfare across all admissible priors:

$$V_t(a) = \min_{P \in \mathcal{P}_t} \int_\Omega W\big(x(a,\omega)\big)dP(\omega)$$

Generalizing the Savage framework and allowing for multiplicity of probabilistic models while preserving the expected utility structure at the level of each individual prior. Uncertainty, in this framework, is decomposed into *ordinary risk*, captured by each $P \in \mathcal{P}_t$, and ambiguity, captured by the multiplicity of elements in $\mathcal{P}_t$.

3.3.4 Dynamic Structure of Beliefs and Learning

The set of admissible probability measures $\mathcal{P}_t$ is time-dependent and reflects the evolving state of scientific knowledge. As new evidence becomes available, certain

probabilistic models may be reinforced by empirical observations, while others may be ruled out as inconsistent with the data. Accordingly, the regulator's informational environment is characterized by a dynamic process of belief revision driven by the arrival of scientific signals. Let $(s_t)_{t\geq 1}$ denote the sequence of observable scientific signals generated over time and $\mathcal{F}_t$represent the associated information filtration. Each signal provides partial information about the underlying state of nature and about the adequacy of competing probabilistic models. The evolution of beliefs is therefore governed by both the content of the signals and the structural assumptions embedded in each admissible likelihood specification.

Formally, the correspondence $\mathcal{P}_t$ evolves according to a belief updating operator

$$\mathcal{P}_{t+1} = \Phi(\mathcal{P}_t, s_{t+1})$$

where Φ maps the current set of admissible priors and the newly observed signal into a revised set of priors, capturing the idea that learning in the presence of ambiguity is set-valued: instead of converging to a single posterior distribution, the regulator refines the set of plausible probabilistic models as information accumulates. This framework allows for persistent disagreement among admissible models even asymptotically. Scientific learning may eliminate some models that become empirically inconsistent with observed evidence, but it may fail to identify a unique probabilistic representation when the available data are insufficient to fully discriminate between competing structures. This dynamic structure of beliefs implies that the regulator's evaluation of regulatory acts is itself time-dependent. Since the set of admissible priors evolves over time, the ambiguity-adjusted assessment of social welfare associated with any regulatory action

must also be updated sequentially. This creates a feedback loop between scientific information, belief formation, and regulatory evaluation.

3.3.5 Dynamic Consistency and the Rectangularity Condition

The updating operator Φ introduced above is postulated at the level of generality needed for the qualitative results of Sections 3.6-3.10, but its unrestricted form conceals a substantive technical question that must be addressed before the recursive value function of Section 3.7 can be treated as well defined: whether evaluating a regulatory policy today, using the full multi-period worst-case criterion over $\mathcal{P}_0$, agrees with evaluating it recursively, one period at a time, by applying Φ sequentially and taking worst-case continuation values at each step. If the two evaluations diverge, the regulator's plan may be dynamically inconsistent: a policy that is optimal from the vantage point of $t = 0$ may cease to be optimal from the vantage point of $t = 1$, even though no new information has arrived beyond what was already anticipated, purely as an artifact of how $\mathcal{P}_t$ is updated.

This is a familiar difficulty in the decision-theoretic literature[67] but it has a known resolution, that can be obtained if and only if the sequence of admissible-prior sets satisfies a condition termed *rectangularity*: informally, the set $\mathcal{P}_t$ must be generated as a product, across the remaining periods, of one-step-ahead conditional sets, so that any

[67] Cf. Larry G Epstein and Martin Schneider, 'Recursive Multiple-Priors' (2003) 113(1) Journal of Economic Theory 1 – introducing rectangularity as the condition characterizing dynamically consistent updating of multiple-priors preferences – and Larry G Epstein and Martin Schneider, 'Ambiguity, Information Quality, and Asset Pricing' (2008) 63(1) Journal of Finance 197 – for an application illustrating the practical consequences of non-rectangular updating.

admissible conditional model for period $t+1$ can be freely combined with any admissible conditional model for period $t+2$, and so forth, without this combination itself being ruled out by $\mathcal{P}_t$. Where rectangularity fails – for instance, if certain long-run scientific hypotheses are logically tied to certain short-run ones, so that admissibility today constrains which models remain admissible several periods hence in a way that is not separable period by period – the worst-case criterion computed directly over $\mathcal{P}_0$ need not coincide with the one obtained by recursively applying Φ, and Proposition 1's threshold characterization, together with Propositions 2 and 3, should be understood as holding for the regulator's *recursive* problem, which is the natural object for a repeatedly reviewed license but is not, in general, an equivalent restatement of the regulator's full-commitment problem at $t=0$.

For the remainder of the paper we maintain rectangularity of $\mathcal{P}_t$ as a standing assumption, consistent with the overwhelming majority of the literature that builds on the recursive multiple-priors framework[68]. This is not an innocuous simplification but rectangularity is most plausible when the sources of scientific ambiguity affecting different future periods are informationally separable – e.g., when disagreement

[68] Cf., in this regard, Pascal Maenhout Hao Xing and Anne G. Balter, 'Model Ambiguity versus Model Misspecification in Dynamic Portfolio Choice' (2026) 81(3) Journal of Finance 1741; Philipp Karl Illeditsch, 'Ambiguous Information, Portfolio Inertia, and Excess Volatility' (2011) 66(6) Journal of Finance 2213; Larry G Epstein and Martin Schneider, 'Learning under Ambiguity' (2007) 74(4) Review of Economic Studies 1275; Kiyohiko G Nishimura and Hiroyuki Ozaki, 'Irreversible Investment and Knightian Uncertainty' (2007) 136(1) Journal of Economic Theory 668; Garud Iyengar, 'Robust Dynamic Programming' (2005) 30(2) Mathematics of Operations Research 257; Pascal J Maenhout, 'Robust Portfolio Rules and Asset Pricing' (2004) 17(4) Review of Financial Studies 951; Fabio Trojani and Paolo Vanini, 'Robustness and Ambiguity Aversion in General Equilibrium' (2004) 8(2) Review of Finance 279; Larry G Epstein and Martin Schneider, 'Recursive Multiple-Priors' (2003) 113(1) Journal of Economic Theory 1; Zengjing Chen and Larry G Epstein, 'Ambiguity, Risk, and Asset Returns in Continuous Time' (2002) 70(4) Econometrica 1403.

concerns distinct physicochemical mechanisms (such as aggregation behavior, membrane permeability, environmental persistence) that can be independently confirmed or falsified by distinct lines of evidence – and least plausible when a single, currently unresolved scientific hypothesis – for instance, a unified toxicological mechanism linking nanoscale surface chemistry to bioaccumulation – would, if confirmed, simultaneously pin down the likely resolution of what otherwise appear as several distinct sources of ambiguity across future periods. In the latter case, the recursive characterization developed *infra* remains a useful and tractable approximation, but the regulator's genuinely optimal policy could in principle require a degree of forward commitment that a purely recursive, period-by-period threshold rule cannot replicate – an issue we flag rather than resolve, and one that connects naturally to the switching-cost extension we will further develop in Section 3.9, since imperfect recursiveness of this kind is itself a further reason why regulatory policy should not be expected to react to every incremental revision of $\mathcal{P}_t$ as though no cost attached to doing so.

3.4 Regulatory Control and Endogenous Information Production

Until now we have characterized the evolution of regulators' beliefs as a process driven by the arrival of exogenous scientific signals, treating learning as an informational consequence of external research activity, over which the regulator has no direct control. In practice, however, regulatory frameworks do not limit themselves to observing

scientific information, because they actively shape its production.[69] In the case of nanotechnology, for instance, regulatory authorities impose monitoring requirements, mandate post-market studies, define testing protocols, and require disclosure of experimental data, directly affecting both the quantity and the quality of the information that becomes available over time. To capture this feature, we might endogenize the information structure and let $m_t \in M$ denote the level of regulatory monitoring or informational effort chosen by the regulator at time t. The variable m_t represents a composite regulatory instrument that includes, among other dimensions, the intensity of surveillance, the stringency of reporting obligations, the frequency of inspections, and the scope of required scientific testing. The key assumption is that the distribution of scientific signals is affected by regulatory action. Formally, this is given by

$$s_{t+1} \sim F(\cdot \mid \theta, m_t)$$

where F denotes the data-generating process conditional on the underlying state of nature θ and the chosen level of monitoring m_t. Higher levels of monitoring should improve the informational content of signals, either by reducing measurement error, increasing sample size, or expanding the observable dimensions of potential harm. This specification implies that information is no longer an exogenous constraint on regulatory decision-making, but an endogenous object of choice. The regulator faces a problem of action and information design, in which regulatory interventions determine social

[69] Cf., in this perspective, Jochen Gläser and Grit Laudel, 'Governing Science: How Science Policy Shapes Research Content' (2016) 57(1) *European Journal of Sociology / Archives Européennes de Sociologie* 117; Anna Wesselink, Karen S Buchanan, Yola Georgiadou and Esther Turnhout, 'Technical Knowledge, Discursive Spaces and Politics at the Science–Policy Interface' (2013) 30 *Environmental Science & Policy* 1; Lorna Schrefler, 'The Usage of Scientific Knowledge by Independent Regulatory Agencies' (2010) 23(2) *Governance* 309.

outcomes directly – through authorization decisions – but also indirectly – by shaping the future evolution of knowledge – in such a way that

$$\mathcal{P}_{t+1} = \Phi(\mathcal{P}_t, s_{t+1}(m_t))$$

and belief updating depends on the informational structure induced by regulatory effort. In this perspective, the choice of m_t influences not current administrative costs and the future environment in which subsequent regulatory decisions will be taken, introducing a dynamic feedback mechanism between regulation and learning. According to this realistic expectation, a stricter monitoring increases informational precision, which in turn affects the evolution of ambiguity and the perceived desirability of future regulatory action. Regulatory policy, in this light, may be understood as a mechanism that jointly determines outcomes and the evolution of the informational structure governing future decisions of the legislator.

3.5 Regulatory Actions and Dynamic Control Problem

Given the informational structure and the endogenous evolution of beliefs described above, the regulator faces a dynamic decision problem in which policy interventions jointly affect current welfare and the future evolution of the epistemic environment. Regulatory choice, in this framework, should be modelled as a sequence of intertemporal decisions affecting both outcomes and information. At each date t, the regulator selects a policy vector $\pi_t = (L_t, m_t, T_t)$ where $L_t \in \{0,1\}$ denotes the licensing decision, $m_t \in M$ represents the level of monitoring or informational effort, and T_t captures the set of contractual or administrative obligations imposed on the innovator, including reporting requirements, disclosure rules, and compliance constraints. The licensing decision L_t

determines whether the technology is allowed to operate in the market. Conditional on authorization, the policy instruments m_t and $T_t$determine the informational intensity of regulation and the institutional framework governing the behavior of the innovator during commercialization.

In this context, the implementation of regulatory policy might generate both immediate and dynamic effects. On the one hand, licensing and contractual restrictions directly affect social welfare through their impact on economic surplus and potential harm, while, on the other hand, monitoring and disclosure requirements influence the informational structure of the environment by affecting the distribution of future scientific signals, as previously described. To capture the cost of information production, assume that monitoring is associated with a convex cost function $k(m_t)$, with k' and k'' both strictly positive, reflecting the increasing marginal cost of obtaining higher-quality or more precise scientific information. Given this structure, the regulator's problem becomes inherently dynamic: current decisions determine not only current welfare but also the evolution of ambiguity and the informational constraints faced in future periods. In particular, the evolution of beliefs is given by

$$\mathcal{P}_{t+1} = \Phi(\mathcal{P}_t, s_{t+1}(m_t)), \qquad s_{t+1} \sim F(\cdot \mid \theta, m_t)$$

and the dynamic structure of regulation is characterized by a sequential decision-making problem under an evolving set of admissible probability measures.

3.6 Ambiguity-Adjusted Harm Index and State Reduction

To obtain a tractable representation of the regulatory problem, it might be useful to refine the relevant informational content of $\mathcal{P}_t$, for decision-making purposes, in light of an

ambiguity-adjusted index.[70] The key observation is that, although the regulator faces ambiguity over probability measures, policy evaluation ultimately depends on the induced set of expected social welfare levels and, for each admissible prior $P \in \mathcal{P}_t$we can further refine the associated expected social harm of a given regulatory policy $a \in A$ as

$$H_t(a; P) = \int_{\Omega} h\big(x(a, \omega)\big) dP(\omega)$$

where $h(\cdot)$ denotes the social harm component embedded in the welfare function decomposition. Given ambiguity over priors, the regulator evaluates each policy through the entire set of admissible expected harms and the ambiguity-induced harm hence is

$$\mathcal{H}_t(a) = \{H_t(a; P) : P \in \mathcal{P}_t\}$$

Regulatory decision-making requires a scalar evaluation criterion, so the relevant object is probably an aggregation of $\mathcal{H}_t$, reflecting the regulator's attitude toward ambiguity. If we consider, as a benchmark, the worst-case evaluation criterion, selecting the highest admissible expected harm associated with each regulatory action, we might define a *ambiguity-adjusted social harm index*:

$$A_t(a) = \sup_{P \in \mathcal{P}_t} \int_{\Omega} h\big(x(a, \omega)\big) dP(\omega)$$

which provides a scalar summary of the worst-case expected harm associated with a given policy, conditional on the information available at time t. It collapses the infinite-

[70] For this purpose, we impose an additional structural assumption on the welfare function: social welfare associated with a given regulatory action is additively separable into a deployment benefit and a harm component, $W(x(a,\omega)) = b(a) - h(x(a,\omega))$, where $b(a)$ captures the surplus generated by market access (which we take to be a constant B conditional on $Lt = 1$, and zero otherwise) and $h(\cdot)$ captures the harm component introduced above. This separability assumption, which mighe be standard but non-trivial in the present ambiguity-based framework, is what allows the infinite-dimensional welfare comparison across states to collapse into the scalar reduced-form objective used in Section 3.7.

dimensional uncertainty set $\mathcal{P}_t$ into a single state variable that is directly relevant for regulatory decision-making. For expositional simplicity, and to highlight the evolution of the informational environment independently of specific actions, we define the reduced-form state variable

$$A_t \equiv \min_{a \in A} A_t\,(a).$$

This reduction is defined for the specific case in which harm components $h\big(x(a,\omega)\big)$ is separable in the regulator's action, i.e., $h\big(x(a,\omega)\big) = h(\omega)$ for all $a \in A$ such that $L = 1$ and harm realization, conditional on deployment, is dependent on the state of nature. Under this assumption, A_t is a policy-independent statistic of the informational environment, alias the worst-case expected harm from deployment given the current admissible priors.

The evolution of this reduced state is induced jointly by scientific learning and regulatory monitoring. In particular, since both $\mathcal{P}_t$ and the signal process $s_{t+1}(m_t)$ evolve over time, the ambiguity-adjusted harm index follows a controlled stochastic process:

$$A_{t+1} = \Psi(A_t, s_{t+1}, m_t)$$

where $\Psi(\cdot)$ captures the combined effect of new evidence and regulatory information production on the perceived worst-case harm. This reduction will allow the regulatory problem to be reformulated in recursive form, with A_t serving as the fundamental state variable governing intertemporal policy choices.

3.7 The Dynamic Optimization Problem

Given the reduced-form representation of the regulatory environment, the evolution of the system can be fully characterized by the scalar state variable A_t which summarizes the ambiguity-adjusted assessment of social harm. The regulator, in particular, observes A_t at each date and chooses a policy vector $\pi_t = (L_t, m_t, T_t)$that jointly determines current payoffs and the future evolution of the state. We have already anticipated that the regulatory problem is inherently intertemporal and current decisions affect not only contemporaneous social welfare, but also the trajectory of informational refinement through their impact on the distribution of future signals and, consequently, on the evolution of A_t. At each date t, conditional on A_t, the regulator solves:

$$\max_{\{L_t, m_t, T_t\}_{t=0}^{\infty}} \mathbb{E}\left[\sum_{t=0}^{\infty} \delta(L_t \cdot (B - A_t) - k(m_t))\right]$$

where B denotes the gross social benefit of successful deployment of the technology; A_t is the ambiguity-adjusted expected social harm at time t; $k(m_t)$ is the convex cost of information production; $\delta \in (0,1)$ is a social discount factor. The objective function reflects the fact that authorization decisions generate immediate social surplus net of expected harm and, at the same time, monitoring expenditures reduce current welfare but affect future beliefs through their impact on the informational structure. Within this conceptual framework, optimization is subject to the dynamic constraint governing the evolution of the epistemic state:

$$A_{t+1} = \Psi(A_t, s_{t+1}, m_t), \qquad s_{t+1} \sim \mathrm{F}(\cdot \mid \theta, m_t)$$

which embeds both the stochastic nature of scientific evidence and the endogenous role of regulatory monitoring in shaping information flows. Furthermore, the licensing decision is constrained by feasibility requirements on market access:

$$L_t \in \{0,1\}, \qquad \forall t \geq 0$$

while monitoring and contractual instruments are chosen from admissible compact sets $m_t \in \mathrm{M}, T_t \in \mathcal{T}$. The problem thus takes the form of a stochastic dynamic programming problem with endogenous information production and ambiguity-driven state evolution, in which the regulator determines, at the same time, whether the technology is authorized for market participation, the intensity of information acquisition, and the institutional structure governing ongoing compliance. In this structure, information is not an exogenous constraint but becomes a choice variable, such that the evolution of uncertainty becomes itself endogenous to the regulatory mechanism. This leads naturally to a recursive representation of the value function:

$$V(A_t) = \max_{L_t, m_t, T_t} \{L_t \cdot (B - A_t) - k(m_t) + \delta \mathbb{E}[V(A_{t+1}) | A_t, m_t]\}$$

which reflects the stochastic evolution of signals conditional on monitoring effort.[71]

3.8 The Core Structure of the Model and Main Result

The dynamic programming problem derived in the previous section provides a complete characterization of the regulator's decision environment under ambiguity and endogenous information production. Nevertheless, the economic content of the model can be further clarified by identifying its fundamental structural implications for the

[71] As a side note, the dependence of A_{t+1} on m_t captures the feedback loop between regulatory intensity and epistemic refinement.

nature of regulatory authorization, in which regulatory decisions are driven by a single evolving state variable A_t, summarizing the ambiguity-adjusted expected social harm associated with the technology. All policy instruments – licensing, monitoring, and contractual enforcement – operate by influencing either the current evaluation of A_t or its future evolution. Given this structure, the regulator's optimal policy can be characterized in terms of a threshold rule over the endogenous state variable. In particular, there exists a critical value $\bar{A}$ such that regulatory authorization is optimal if and only if the ambiguity-adjusted harm does not exceed this threshold.

Proposition (Provisional Legality as Optimal Policy)

There exists a threshold $\bar{A} > 0$ such that the optimal licensing rule takes the form

$$L_t^* = \begin{cases} 1 \; if \quad A_t \leq \bar{A} \\ 0 \; otherwise \end{cases}$$

Moreover, the optimal monitoring policy m_t^* is state-dependent and satisfies

$$m_t^* = m^*(A_t),$$

with $m^*(\cdot)$ increasing in A_t and optimal contractual structure T_t^* is such that it internalizes the marginal informational value of regulation on the future evolution of ambiguity.

Proof sketch

The result follows from the recursive structure of the value function

$$V(A_t) = \max_{L_t, m_t, T_t} \{L_t \cdot (B - A_t) - k(m_t) + \delta \mathbb{E}[V(A_{t+1}) | A_t, m_t]\}$$

We first establish that $V(\cdot)$ is weakly decreasing. Fix $A' > A$. For any feasible policy (L_t, m_t, T_t) attainable at state A', the same policy is feasible at state A, and generates

weakly higher current payoff, since $L_t\ (B - A_t)$ is decreasing in A_t. By the transition equation $A_{t+1} = \Psi(A_t, s_{t+1}, m_t)$, and assuming Ψ is weakly increasing in its first argument – i.e., higher current ambiguity does not, ceteris paribus, produce a lower continuation ambiguity – the continuation value inherits the same monotonicity by a standard contraction-mapping argument.[72] Blackwell sufficient conditions apply given $\delta \in (0,1)$ and boundedness of per-period payoffs. Hence $V(A) \geq V(A')$, and the result follows by induction on the value-iteration operator. The licensing decision L_t enters linearly in current payoffs and does not affect the cost of information production directly. Therefore, conditional on A_t, the optimal choice of L_t reduces to a static comparison between the net expected benefit of deployment $B - A_t$ and the continuation value of delaying authorization. The optimality of state-dependent monitoring follows from the fact that m_t affects the distribution of future states through its impact on the signal structure. Higher values of A_t increase the marginal value of information, hence optimal monitoring is increasing in the ambiguity-adjusted harm index.[73] Eventually, contractual instruments T_t operate as commitment devices shaping the informational environment and ensure incentive compatibility of the innovator's behavior, reinforcing the dynamic consistency of the optimal policy.

[72] Note that the monotonicity of Ψ in A_t will be verified in the functional form adopted for the numerical simulation in Section 4.

[73] Formally, this follows from a single-crossing argument: the objective is supermodular in $(m_t,\ A_t)$ whenever $\frac{\partial^2[\delta \mathbb{E}[V(A_{t+1}) \mid A_t, m_t]]}{\partial m_t\, \partial A_t} > 0$ and the marginal reduction in continuation harm from an additional unit of monitoring is itself increasing in the current level of ambiguity, a condition satisfied by the proportional-learning specification $\lambda(m_t)A_t$ adopted in Section 4.2. Under supermodularity, Topkis's theorem guarantees that the $\arg max\, m_t^*(A_t)$ is monotone non-decreasing in A_t. See, for more detailed discussion of this characteristic, Rabah Amir, 'Supermodularity and Complementarity in Economics: An Elementary Survey' (2005) 71(3) *Southern Economic Journal* 636.

3.9 Regulatory Hysteresis under Switching Costs

Proposition 1 characterizes the optimal licensing rule as a single, memoryless threshold: authorization is granted whenever $A_t \leq \overline{A}$ and withdrawn the instant A_t crosses above it, with no cost attached to the act of switching itself. This "bang-bang" property is a direct consequence of the frictionless specification of the licensing decision in Section 3.5, in which L_t enters the flow payoff linearly and independently of L_{t-1}. This is, nevertheless, a stylized feature, because granting and revoking a license are costly administrative acts: granting requires a formal review process and generates reliance by the innovator and by downstream users of the technology, while revocation typically requires its own proceeding, is subject to notice and appeal rights, and imposes disruption costs on firms, employees, and any product users who must transition away from the technology. A licensing rule that treats these transitions as costless might, more in general, understate the value of inertia and overstate the frequency with which the socially optimal policy actually changes authorization status.

This subsection extends the previously mentioned dynamic program to incorporate an explicit switching cost $\kappa \geq 0$, charged whenever licensing status changes between consecutive periods, and shows that the optimal policy takes the form of a two-threshold hysteresis band rather than a single cutoff. Therefore, licensing status L_{t-1} becomes an additional state variable, and the flow payoff is modified to

$$L_t \cdot (B - A_t) - k(m_t) - \kappa|L_t - L_{t-1}|,$$

so that the recursive value function becomes

$$V(A_t, L_{t-1}) = \max_{L_t,\, m_t,\, T_t} \{L_t \cdot (B - A_t) - k(m_t) - \kappa|L_t - L_{t-1}| + \delta\mathbb{E}[V(A_{t+1}, L_t)|A_t, m_t]\}.$$

Proposition 2 (Regulatory Hysteresis)

There exist two thresholds $\overline{A}_g$ and $\overline{A}_r$, with $0 < \overline{A}_g \leq \overline{A} \leq \overline{A}_r$, such that the optimal licensing rule is history-dependent:

$$L_t^* = \begin{cases} 1 & if\ L_{t-1} = 0\ and\ A_t \leq \overline{A}_g \text{ (grant)} \\ 1 & if\ L_{t-1} = 1\ and\ A_t \leq \overline{A}_r \text{ (remain licensed)} \\ 0 & otherwise \end{cases}$$

For $A_t \in \left(\overline{A}_g, \overline{A}_r\right]$, the optimal action coincides with the previous period's licensing status: neither granting an unlicensed technology nor revoking a licensed one is optimal, and the regulator optimally does nothing. As $\kappa \to 0$, both thresholds converge to the frictionless cutoff of Proposition 1, $\overline{A}_g, \overline{A}_r \to \overline{A}$, so that Proposition 1 is recovered as the switching-cost-free limit of the present result.

Proof sketch

Define $V^0(A_t)$ and $V^1(A_t)$ as the value function restricted to $L_{t-1} = 0$ and $L_{t-1} = 1$ respectively. By the same monotonicity argument used in the proof of Proposition 1 – fixing $A' > A$ and noting that any policy feasible at A' remains feasible and weakly more attractive at A – both $V^0(\cdot)$ and $V^1(\cdot)$ are weakly decreasing in A_t, and $V^1(A_t) \geq V^0(A_t) - \kappa$ for all A_t, since a currently licensed regulator can always replicate the unlicensed regulator's continuation policy at the one-time cost of an unnecessary future switch. The existence of two distinct switching boundaries then follows from the standard value-matching and smooth-pasting argument used throughout the real-options literature on entry and exit under switching costs:[74] because switching sacrifices κ today

[74] The value-matching and smooth-pasting conditions used here to characterize the two switching boundaries are the standard technique in this literature for solving free-boundary problems of this kind.

in exchange for an uncertain continuation payoff that may reverse before the cost is recouped, it is optimal to grant a license only once A_t falls to a level strictly more favorable than the frictionless cutoff $\overline{A}$ and, symmetrically, to tolerate ambiguity strictly above $\overline{A}$ before revoking an existing license, since revoking prematurely risks having to re-grant, and pay κ again, once ambiguity subsequently falls. Eventually, this feature yields $\overline{A}_g \leq \overline{A} \leq \overline{A}_r$, with strict inequality whenever $\kappa > 0$. A full closed-form characterization of $\overline{A}_g$ and $\overline{A}_r$ hence requires specifying the functional form of $\Psi(\cdot)$ and solving the associated free-boundary problem numerically; this is left for future work building on the simulation framework of Section 4[75].

3.10 A Minimal Screening Extension: The Optimal License Menu under Private Information about Risk Type

We treated, until now, the innovator as a passive source of scientific signals, whose private information, if any, plays no strategic role in the regulator's problem. Nonetheless, the innovator may hold private information about the technology's underlying risk type prior to licensing, and the regulator's menu of licenses – differing in monitoring intensity and associated fees – might be designed to induce truthful revelation of that type. This subsection develops a minimal, tractable version of that

---

Cf., *ex plurimis*, Avinash K Dixit, 'Entry and Exit Decisions under Uncertainty' (1989) 97(3) Journal of Political Economy 620; Avinash K Dixit and Robert S Pindyck, *Investment under Uncertainty* (Princeton University Press 1994) ch 7; Øksendal, 'A simplified exposition of smooth pasting' (1999) 63(3–4) Economics Letters 173; Shackleton and Sødal (2005), 'Smooth Pasting as Rate of Return Equalization', Economics Letters 89(2), 200.

[75] The institutional implication of Proposition 2 is developed further in Section 5: administrative procedures that make revocation costly and deliberate – e.g., formal proceedings, advance notice, evidentiary hearings, rights of appeal – are not mere procedural safeguards but, instead, credible institutional mechanism through which the legal system can separate $\overline{A}_r$ from $\overline{A}_g$ and implement the welfare-maximizing hysteresis band rather than volatile threshold rule of our first Proposition.

screening problem, sufficient to characterize the qualitative structure of the optimal menu; a fully dynamic mechanism in which the innovator's type interacts with the continuously evolving ambiguity state A_t of previous Sections is left, as originally indicated, for future research.

Consider a static complement to the dynamic model in which the innovator privately knows its technology's risk type $\theta \in \{L, H\}$ at the moment of application, with H denoting the higher-risk type and $\Pr(\theta = L) = q \in (0,1)$ common knowledge. Licensing a type-θ technology under monitoring intensity m generates social welfare $W(m, \theta) = B - D(m, \theta) - k(m)$, where $D(m, \theta)$ is expected social harm, decreasing and convex in m, with $D(m, H) > D(m, L)$ for every m and $D_m(m, H) < D_m(m, L)$: monitoring is both more urgently needed and more effective at the margin for the riskier type, consistent with the increasing $m^*(\cdot)$ characterized in Proposition 1. Compliance with monitoring intensity m imposes a private cost $C(m, \theta)$ on the innovator, increasing and convex in m, satisfying the single-crossing condition $C_m(m, H) > C_m(m, L)$ for all $m > 0$: a riskier technology requires more extensive remediation, disclosure, and corrective investment per unit of additional monitoring, so that intensive monitoring is costlier at the margin for type H than for type L. An innovator of true type θ who is allocated the contract (m, T) – i.e., monitoring intensity and regulatory fee – earns private profit $\pi = R - C(m, \theta) - T$, where R is the (type-independent) private value of market access.

The regulator offers a menu $\{(m_L, T_L), (m_H, T_H)\}$ to maximize expected social welfare net of transfers,

$$q[W(m_L, L) + T_L] + (1 - q)[W(m_H, H) + T_H],$$

subject to the individual-rationality and incentive-compatibility constraints

$$IR_\theta: \quad R - C(m_\theta, \theta) - T_\theta \geq 0, \qquad \theta \in \{L, H\},$$

$$IC_\theta: \quad R - C(m_\theta, \theta) - T_\theta \geq R - C(m_{\theta'}, \theta) - T_{\theta'}, \qquad \theta' \neq \theta.$$

Proposition 3 (Optimal License Menu under Adverse Selection)

Under the single-crossing condition $C_m(m, H) > C_m(m, L)$, the socially optimal incentive-compatible menu satisfies:

(i) $m_L = m^*(L)$: the low-risk type receives the first-best, undistorted monitoring intensity. There is, therefore, "no distortion" for the type whose incentive constraint does not bind;

(ii) $m_H < m^*(H)$: the high-risk type receives monitoring intensity strictly below its own first-best level, creating a downward distortion introduced to relax the low-risk type's temptation to misrepresent itself as high-risk;

(iii) IR_H binds – i.e., the high-risk type earns zero surplus – and IC_L binds – i.e., the low-risk type is exactly indifferent between its own contract and the high-risk contract – while IR_L is slack – i.e., the low-risk type earns a strictly positive information rent equal to $C(m_H, H) - C(m_H, L) > 0$.

Proof sketch

The result might be framed as a two-type specialization of the standard regulation model under adverse selection.[76] Because $C_m(m, H) > C_m(m, L)$ for all m, a type-L innovator

[76] See, in this regard, David P Baron and Roger B Myerson, 'Regulating a Monopolist with Unknown Costs' (1982) 50(4) Econometrica 911 and Jean-Jacques Laffont and Jean Tirole, *A Theory of Incentives in Procurement and Regulation* (MIT Press 1993) ch 1-2, for the general adverse-selection regulation model of which the two-type menu characterized in Proposition 3 is a direct application.

can always deliver any monitoring intensity m more cheaply than a type-H innovator, so that L – not H – is the type tempted to misrepresent itself, since claiming to be H would let L pocket the difference between the fee T_H (calibrated to H's higher compliance cost) and its own, lower, true cost of delivering m_H. Consequently only IC_L and IR_H can bind at the optimum; IC_H and IR_L are then verified to be slack given (i)-(iii). Substituting the binding constraints into the regulator's objective and optimizing pointwise over m_L and m_H yields an unconstrained (first-best) optimum in m_L, since L's allocation does not enter the rent term paid to itself, while the choice of m_H trades off its direct welfare contribution $W(m_H, H)$ against the marginal increase in the rent $C(m_H, H) - C(m_H, L)$ that a higher m_H requires paying to L under IC_L; because this rent term is strictly increasing in m_H by single-crossing, the optimal m_H is set strictly below $m^*(H)$, establishing (ii).

Proposition 3 formalizes, in a static and minimal setting, the qualitative claim advanced in the Introduction: an incentive-compatible license menu optimally couples risk type with a systematic tradeoff between monitoring intensity and the innovator's net return from licensure, with low-risk technologies receiving efficient, comparatively light-touch treatment and earning an information rent for their truthfulness, while high-risk technologies are monitored less intensively than first-best efficiency would independently require – precisely because monitoring the high-risk type at its own efficient level would make misrepresentation too attractive for the low-risk type to resist. Embedding this static screening logic into the fully dynamic, continuously evolving ambiguity state A_t of Sections 3.6-3.9 – so that the innovator's report can itself be

revised, and reputational capital accumulated or lost, as monitoring evidence accrues over the life of the license – remains, as indicated previously, a substantial extension left for future research.

4. Numerical Simulation of Dynamic Regulatory Learning

The dynamic regulatory framework developed in the previous Chapter provided an analytical characterization of optimal licensing under ambiguity, modelling regulatory interventions as endogenous learning processes in which monitoring simultaneously affects current regulatory costs and the future evolution of scientific information. Within the proposed framework, it could be interesting to check the evolution of ambiguity and its interaction with regulatory effort and informational refinement. In particular, if ambiguity-adjusted social harm generates a nonlinear stochastic system, it might be useful to investigate dynamic properties through numerical simulation, using a computational experiment to investigate the internal implications of the theoretical model under controlled institutional environments. The simulation evaluates, in particular, how different institutional architectures influence the dynamic evolution of regulatory uncertainty when all other structural assumptions remain unchanged, isolating – using Monte Carlo simulations to estimate multiple stochastic trajectories – the contribution of monitoring policies to the process of scientific learning.[77]

[77] Unlike comparative institutional analyses based exclusively upon qualitative legal reasoning, numerical simulations allow the entire probability distribution of regulatory outcomes to be examined through time. In light of the fact that regulatory systems do not produce a single deterministic trajectory of scientific knowledge and information accumulates through inherently stochastic processes, including experimental discoveries, epidemiological evidence, environmental observations, technological innovation, measurement error, and unforeseen biological mechanisms, we intend to check the entire distribution of possible future regulatory states generated under alternative institutional arrangements. For additional

4.1 Approximations of the Dynamic Regulatory Process and Simulations

In Chapter 3 we have characterized regulatory decision-making processes on nanotechnologies as a stochastic dynamic control problem, in which the evolution of ambiguity adjusted social harm depends upon the regulator's current informational state, the arrival of new scientific evidence, and the endogenous monitoring effort selected through the regulatory mechanism.[78] Formally, we have represented the dynamics of the informational environment as $A_{t+1} = \Psi(A_t, s_{t+1}, m_t)$, where A_t denoted the ambiguity-adjusted social harm index, s_{t+1} the newly observed scientific evidence, and m_t the intensity of regulatory monitoring. The function $\Psi(\cdot)$ was introduced as a reduced representation of the complex mechanisms through which scientific information modifies the regulator's perception of future risk. The simulation replaces this abstract transition function by constructing an ambiguity-adjusted harm index evolving according to the stochastic difference equation

$$A_{t+1} = \max\{0, A_t - \lambda(m_t)A_t + \varepsilon_t$$

where the first component – i.e., $-\lambda(m_t)A_t$ – represents systematic informational learning in which, across time, the regulator acquires additional scientific evidence regarding the technology under examination and the parameter $\lambda(m_t)$ captures the speed

references in this perspective, see, *ex plurimis* , Martin Straume and Michael L Johnson, 'Monte Carlo Method for Determining Complete Confidence Probability Distributions of Estimated Model Parameters' in *Methods in Enzymology* vol 210 (Academic Press 1992) 117.

[78] The simulation that follows should be read as a numerical illustration of the qualitative dynamics implied by the theoretical model, rather than as an empirically calibrated forecast of regulatory outcomes for any specific class of nanomaterials. No attempt is made to calibrate $\lambda(mt)$ or $\sigma(mt)$ against toxicological, epidemiological, or environmental-persistence data; the exercise instead isolates, under stylized and admittedly simplified assumptions, the structural consequences of endogenizing monitoring intensity within a learning process, holding all else constant across the three institutional regimes.

at which regulatory learning takes place, while the second component of the transition equation – i.e., $\varepsilon_t \sim N(0, \sigma(m_t)^2)$ – captures the stochastic nature of scientific discovery and regulatory learning, in which the random disturbance accounts for the fact the production of scientific evidence is never perfectly predictable.[79] The operator of the transition equation – i.e., $\max\{0,\cdot\}$ – imposes a non-negativity constraint on the ambiguity-adjusted harm index because we might expect that complete scientific certainty would not create negative harm while, instead, reducing it toward zero.

In the simulations we assume that the European regulatory framework is founded upon the principle that scientific uncertainty should not justify regulatory inaction, while uncertainty constitutes a sufficient reason for intensifying regulatory oversight and promoting additional scientific knowledge on the subject.[80] The second regulatory regime that we implement in the simulations is the American system, which – as we have seen in Chapter 2 – places greater emphasis upon facilitating technological innovation and market entry, accepting that scientific knowledge evolve through the ordinary operation of markets and information is less heavily centralized, leaving space

[79] In particular, while it is true that new information may emerge from laboratory experiments, epidemiological investigations, adverse-event reporting, technological innovation, environmental observations, judicial proceedings, or entirely unforeseen mechanisms that alter the regulator's understanding of the technology under consideration, the timing, magnitude, and informational content of these discoveries cannot be known ex ante. To account for this phenomenon, we introduced unavoidable randomness into the evolution of the ambiguity-adjusted harm index. For further details on these randomici informational mechanisms please refer, *ex plurimis*, to Raghu Garud, Joel Gehman and Antonio Paco Giuliani, 'Serendipity Arrangements for Exapting Science-Based Innovations' (2018) 32(1) *Academy of Management Perspectives* 125.

[80] Cf., *inter alia*, Marjolein BA van Asselt and Ellen Vos, 'Wrestling with Uncertain Risks: EU Regulation of GMOs and the Uncertainty Paradox' in Asa Boholm (ed), *New Perspectives on Risk Communication: Uncertainty in a Complex Society* (1st edn, Routledge 2011); Jale Tosun, 'How the EU Handles Uncertain Risks: Understanding the Role of the Precautionary Principle' (2013) 20(10) *Journal of European Public Policy* 157; Ragnar Lofstedt and Frederic Bouder, 'Evidence-based Uncertainty Analysis: What Should We Now Do in Europe? A View Point' (2021) 24(5) *Journal of Risk Research* 521.

to soft law regulatory mechanisms subsequent monitoring allowing emerging evidence to modify regulatory assessments.[81] We simulate dynamic regulatory processes on nanotechnologies, rejecting the premises that normative intensity should remain fixed according to a single regulatory philosophy, treating monitoring as an endogenous policy instrument whose optimal intensity depends upon the regulator's current informational state. In this light, the central intuition underlying DL is that neither the precautionary approach nor the post-market regulatory implant are universally optimal. During the earliest stages of technological development, when ambiguity is substantial and scientific knowledge remains incomplete, the precautionary logic of intensive monitoring generates considerable informational benefits by accelerating learning. Once the regulator has acquired sufficient evidence and ambiguity has been substantially reduced, however, the marginal value of additional monitoring progressively declines. DL therefore integrates the principal strengths of both regulatory traditions within a single adaptive institutional framework in which precaution is employed when uncertainty is greatest and information possesses its highest social value. The simulation operationalizes this adaptive principle by allowing monitoring intensity to evolve as a function of the ambiguity adjusted social harm index.

To investigate these stochastic dynamics, the model was implemented numerically in Python, following the reduced-form state equation introduced in Section 4.2. to generate

[81] Cf., in this perspective, among others, Fred L Block and Matthew R Keller, *State of Innovation: The US Government's Role in Technology Development* (Routledge 2015); Brett M Frischmann, 'Innovation and Institutions: Rethinking the Economics of US Science and Technology Policy' (2000) 24 *Vermont Law Review* 347; Ashish Arora, Andrea Fosfuri and Alfonso Gambardella, *Markets for Technology: The Economics of Innovation and Corporate Strategy* (MIT Press 2004).

– via Monte Carlo simulations[82] – independent realizations of the same stochastic process. The simulations were performed using a sample of 50,000 independent trajectories for each regulatory regime and each simulated trajectory extends over 80 discrete regulatory period, representing successive stages of the regulatory learning process during which new scientific evidence becomes available and regulatory beliefs are updated. Rather than analyzing individual trajectories separately, for every regulatory period, the simulated observations are transformed into a continuous probability density function through Kernel Density Estimation,[83] as implemented in the *gaussian_kde* routine provided by the *SciPy* statistical library. Repeating the estimation across all regulatory periods produces a sequence of probability density functions – describing the complete temporal evolution of the regulatory learning process – which are finally assembled into a continuous three-dimensional probability surface using the visualization facilities provided by the *Matplotlib* library.

[82] Monte Carlo simulations allow the analysis to capture the full distribution of possible regulatory outcomes, providing statistically robust estimates of expected dynamics under stochastic uncertainty. Cf., for further reference, Jiaxin Zhang, 'Modern Monte Carlo Methods for Efficient Uncertainty Quantification and Propagation: A Survey' (2021) 13(5) *WIREs Computational Statistics* 1; Reuven Y Rubinstein and Dirk P Kroese, *Simulation and the Monte Carlo Method* (3rd edn, John Wiley & Sons 2016); HJ Pradlwarter and GI Schuëller, 'On Advanced Monte Carlo Simulation Procedures in Stochastic Structural Dynamics' (1997) 32(4) *International Journal of Non-Linear Mechanics* 735.

[83] Kernel Density Estimation is a non-parametric technique that estimates the underlying probability density directly from the simulated observations without requiring any prior assumption regarding its analytical form. See, for further details, Artur Gramacki, *Nonparametric Kernel Density Estimation and Its Computational Aspects* (Springer International Publishing 2018).

4.2 Simulation Results

The graphical results reported in Figure 1-3 portray the probability distribution generated by each regulatory regime over time, summarizing the evolution of fifty thousand independent realizations of the stochastic learning process.

4.2.1 The European Precautionary Model

The first probability surface – in Figure 1 – reports the evolution of ambiguity-adjusted social harm under the European precautionary regulatory framework. The simulated European trajectory confirms, numerically, the qualitative prediction that was underlined in Section 3.8, alias a general monitoring rule which is not responding to the currently accumulated evidence and producing a single, undifferentiated learning path across the entire population of regulated technologies. In particular, it seems that the European precautionary architecture treats early "good news" and early "bad news" symmetrically in terms of subsequent regulatory intensity, in a form of institutional rigidity that the dynamic licensing framework of Section 3 was designed to correct.

Figure 1: European Regulatory Framework

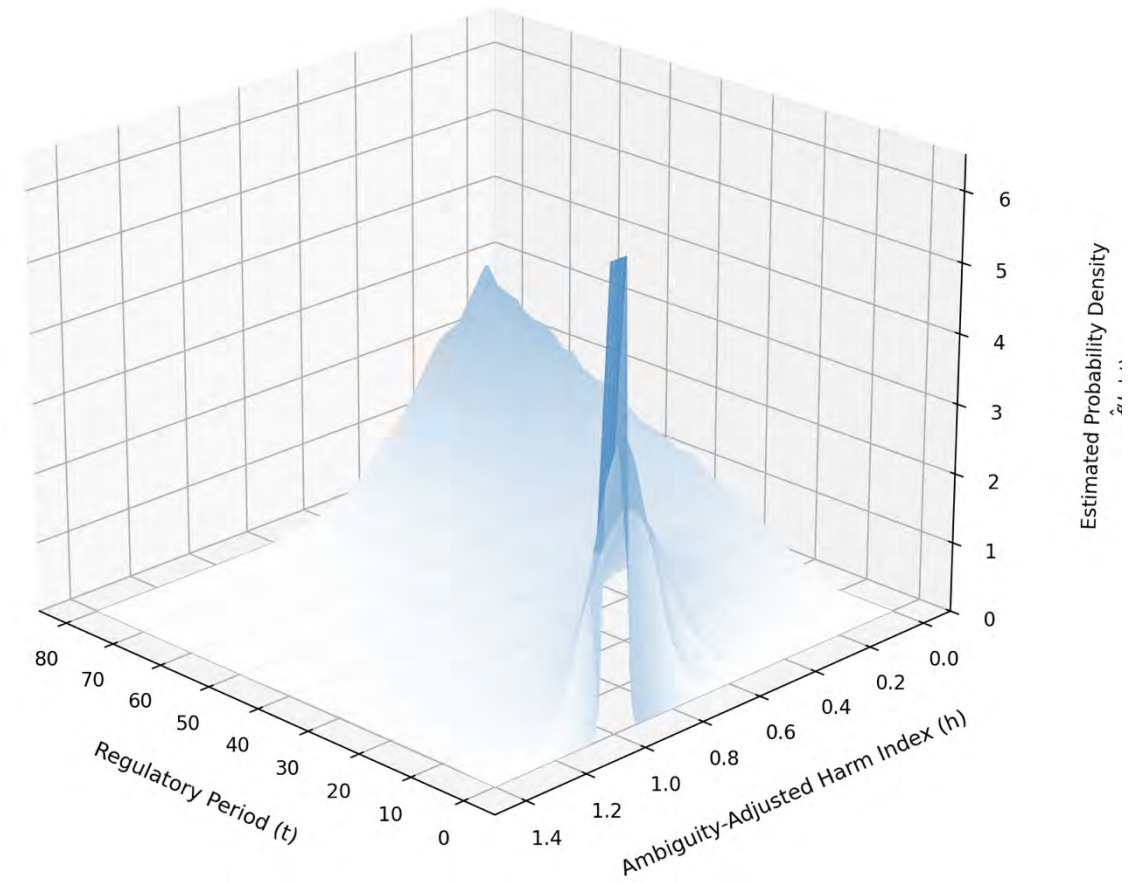

Note: The figure reports the evolution of the estimated probability density of the Ambiguity-Adjusted Harm Index under the European precautionary regulatory framework. The surface is obtained from 50,000 independent Monte Carlo simulations, each evolving over 80 regulatory periods from identical initial conditions. Probability densities are estimated at each period using Gaussian Kernel Density Estimation (KDE) and represented as a continuous three-dimensional surface. The horizontal axis denotes regulatory time, the vertical axis reports the estimated probability density, and the third dimension represents the Ambiguity-Adjusted Harm Index. Higher concentrations indicate a greater proportion of simulated trajectories converging toward the corresponding informational state.

The graph shown in Figure 1 displays a single, sharply peaked density spike at $t = 0$, located exactly at $h = 1$, in which the deterministic starting point, shared by all trajectories, collapses at the outset. Moving along the time axis, the surface broadens and its peak migrates steadily toward lower values of h, while the standard deviation rises and then gradually contracts again as the constant, non-adaptive monitoring rate slowly pulls the distribution downward. The simulated European trajectory confirms,

numerically, the central theoretical prediction for a constant, non-state-contingent monitoring rule, in which the population of trajectories never separates into distinguishable cohorts. After the deterministic spike at $t = 0$, the density mass broadens quickly – reflecting the accumulation of heterogeneous early evidence across otherwise identical nanomaterials – and then drifts as a single coherent ridge toward lower values of h. The rate of descent is visibly slow relative to the complete time horizon period: even by $t = 80$, the bulk of the density has not fully collapsed toward $h = 0$. This is the numerical signature of a regime in which λ does not respond to the information already accumulated. A technology that has generated reassuring early evidence is monitored exactly as intensively, and therefore learns exactly as slowly, as one that has generated worrying early evidence. The welfare cost of this rigidity is structural, not incidental: it is a direct consequence of treating m_t as a policy constant rather than a policy function of A_t, exactly as characterized in Proposition 1 (Section 3.8), where the optimal m_t^* is required to be increasing in A_t.

4.2.2 The American Post-Market Framework

The second probability surface shows the h index evolution of the ex post US regulatory framework. According to the results shown in Figure 2., density mass barely migrates from its initial position across almost the entire time horizon period, while its dispersion increases roughly monotonically instead of contracting. This is a strong numerical illustration of the paper's core critique of the U.S. framework (Section 2.2.3): reliance on post-market correction is structurally mismatched to a risk profile – chronic, delayed, causally diffuse harm – that requires front-loaded, systematic information production.

Figure 2: American Regulatory Framework

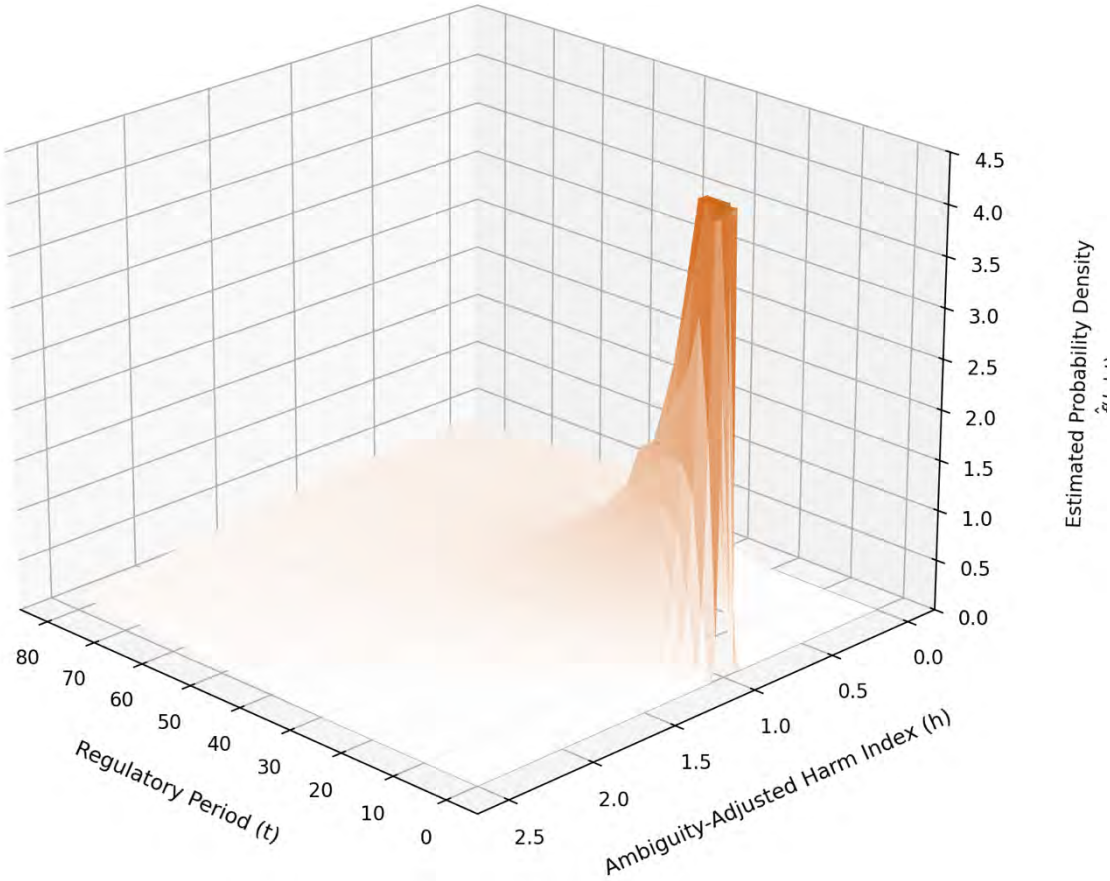

Note: The figure reports the evolution of the estimated probability density of the Ambiguity-Adjusted Harm Index under the American ex post regulatory framework. The surface is obtained from 50,000 independent Monte Carlo simulations, each evolving over 80 regulatory periods from identical initial conditions. Probability densities are estimated at each period using Gaussian Kernel Density Estimation (KDE) and represented as a continuous three-dimensional surface. The horizontal axis denotes regulatory time, the vertical axis reports the estimated probability density, and the third dimension represents the Ambiguity-Adjusted Harm Index. Higher concentrations indicate a greater proportion of simulated trajectories converging toward the corresponding informational state.

The evidence reported in Figure 2 provides quantitative support for the qualitative critique developed in Section 2.2.3. Under the American regime's characteristically low and reactive monitoring intensity, the density surface shows negligible systematic drift away from the initial condition across nearly the entire simulated horizon: the bulk of probability mass remains clustered near $h = 1$ from $t \approx 0$ through $t \approx 80$. At the same

time, the dispersion of the distribution increases rather than contracts, indicating that uncoordinated, market-based information production injects noise into the regulator's assessment faster than it resolves ambiguity. The implication is that a purely reactive, low-intensity monitoring regime does not merely fail to accelerate learning about nanomaterial risk but it renders the distribution of possible future regulatory states strictly less predictable over time, reinforcing the argument that ex post correction might be structurally unsuited to harms that are chronic, delayed, and causally diffuse.

4.2.3 The Dynamic Licensing Framework

In Figure 3 we report a qualitatively distinct pattern of transient bimodality that is absent from both benchmark regimes. Because optimal monitoring intensity $m_t^* = m^*(A_t)$ is itself increasing in the current level of ambiguity (Proposition 1, Section 3.8), trajectories that experience a favorable evidentiary draw early in the regulatory process are subjected to progressively more effective monitoring and consequently separate, over the intermediate window of the simulation, into a visibly distinct "fast-converging" cohort, while trajectories with an unfavorable early draw remain clustered near the initial ambiguity level for longer. This produces two separated concentrations of probability mass connected by a saddle-shaped region of low density. As additional periods elapse, the slow cohort progressively catches up with the fast cohort and the two ridges merge into a single peak at low values of h, consistent with the eventual convergence guaranteed by the monotonicity of the value function established in Section 3.8. Unlike the European and American regimes, Dynamic Licensing is the only regulatory

environment generating cross-sectional heterogeneity in regulatory outcomes conditional on the realized path of accumulated evidence.

Figure 3: Dynamic Licensing Framework

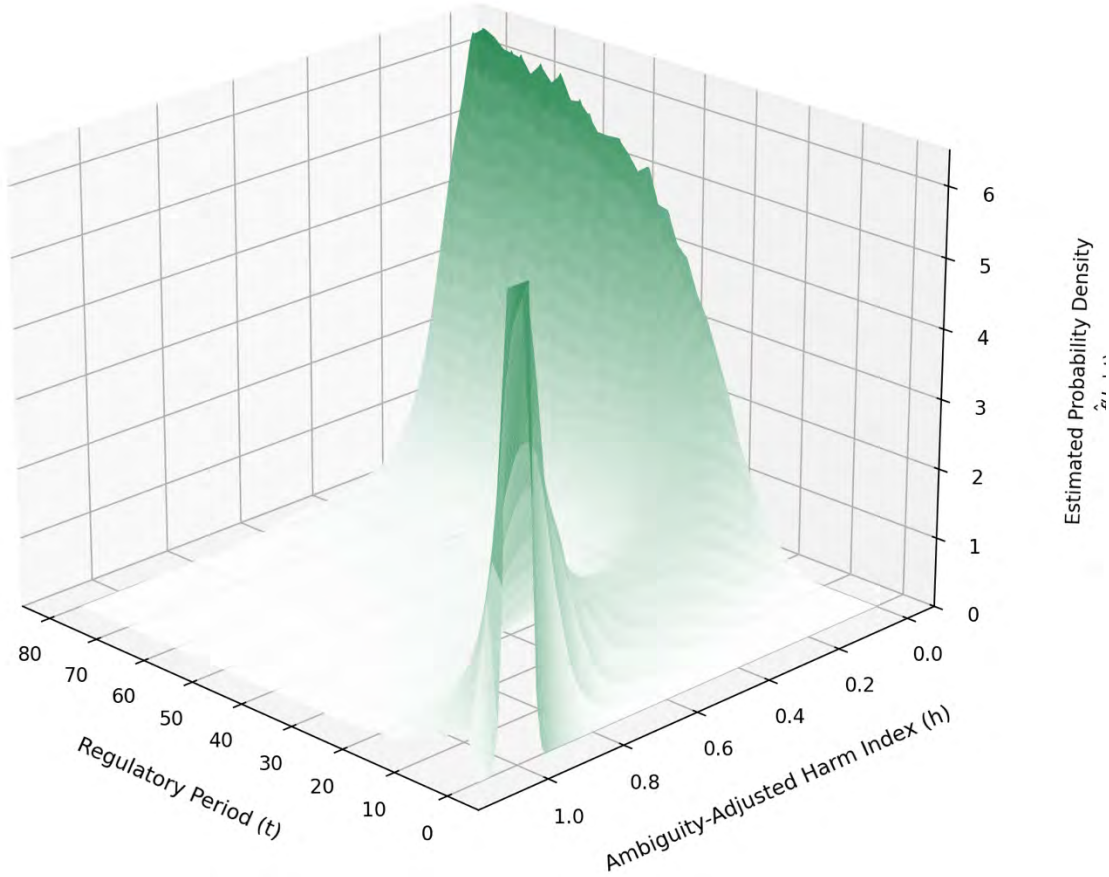

Note: The figure reports the evolution of the estimated probability density of the Ambiguity-Adjusted Harm Index under a Dynamic Licensing Framework regulatory framework. The surface is obtained from 50,000 independent Monte Carlo simulations, each evolving over 80 regulatory periods from identical initial conditions. Probability densities are estimated at each period using Gaussian Kernel Density Estimation (KDE) and represented as a continuous three-dimensional surface. The horizontal axis denotes regulatory time, the vertical axis reports the estimated probability density, and the third dimension represents the Ambiguity-Adjusted Harm Index. Higher concentrations indicate a greater proportion of simulated trajectories converging toward the corresponding informational state.

The pattern shown in Figure 3 is one of *transient bimodality.* Alongside the spike near $h = 1$ at low t – which is the result of the deterministic initial condition – a second distinct local peak emerges intermediate periods and consolidates around progressively lower values of h as t increases. From the graph is immediately evident that the surface

displays two separated concentrations of probability mass – a saddle-shaped valley visibly separating a "high-ambiguity" cohort of trajectories from a "low-ambiguity" cohort – before the lower h peak absorbs the remaining mass and the surface becomes unimodal again for the remainder of the horizon. The "bimodality" is the direct consequence of the proportional monitoring rule adopted for this regime, given that the *rate* of ambiguity reduction is proportional to the *current level* of ambiguity and trajectories that experience a favorable early sequence of informative shocks reduce their own h faster.[84] The population of trajectories is thus split into a fast-converging cohort and a slow-converging cohort, visible as two separated ridges on the surface, before the slow cohort eventually catches up and the two peaks merge.

Substantively, this implies that, unlike the European and American regimes, the Dynamic Licensing regime is the only framework generating meaningful cross-sectional heterogeneity in regulatory outcomes during the transition phase. Under Dynamic Licensing, otherwise similarly situated nanotechnologies can experience materially different regulatory trajectories depending on the realized path of accumulated evidence: some will see monitoring relaxed relatively early, others – subject to the identical policy rule but a less favorable draw of the stochastic learning process – will remain under intensive scrutiny for longer, according to the path-dependent responsiveness anticipated *supra*.

[84] Note that under the proportional rule, this effect further sustains a comparatively high level of monitoring relative to their – already lower – ambiguity, accelerating further their convergence. Trajectories that experience an unfavorable early sequence remain longer at higher h, and continue to draw the (proportionally higher) monitoring intensity appropriate to that higher ambiguity level, but converge more slowly in absolute time.

4.3 Expected Welfare Comparison

The density surfaces reported in Figures 1-3 illustrate how ambiguity evolves under each regime, but they do not by themselves quantify which regime a benevolent regulator would prefer. This subsection closes that gap by computing, directly from the simulated panels, the expected discounted social welfare associated with each regime under the Proposition 1 threshold licensing rule, $L_t = \mathbb{1}\{A_t \leq \bar{A}\}$, applied out-of-sample to the 50,000 simulated trajectories. For each regime, period-by-period social welfare is evaluated as $L_t\big(B - A_t - k(m_t)\big)$, discounted at rate δ and averaged across trajectories, with $k(m_t) = c\, m_t^2$ a convex monitoring cost applied only while the technology remains licensed. The licensing threshold is set at $\bar{A} = 0.8$, the gross benefit of deployment at $B = 1.5$, the discount factor at $\delta = 0.95$, and the monitoring-cost scale at $c = 5$; as with the baseline transition parameters of Section 4.1, these values are illustrative rather than calibrated, and are chosen only to keep the monitoring cost on the same order of magnitude as the harm index itself.

Under this parameterization, expected discounted welfare is approximately 13.13 under Dynamic Licensing, 11.71 under the European regime, and 3.83 under the American regime – a welfare gain for Dynamic Licensing of roughly $+12\%$ over the European benchmark and of roughly $+243\%$ over the American benchmark. The gain over the American regime is large and robust across a wide range of parameter choices, consistent with the qualitative pattern of Figure 2: a monitoring regime that fails to reduce ambiguity systematically also fails to accumulate the informational conditions for stable licensing, so that the license-retention rate under the American regime – the

share of the simulated population satisfying $A_t \leq \bar{A}$ – never exceeds roughly 46% of the population even after 40 periods, against 99% under Dynamic Licensing and 91% under the European regime at the same horizon.

The comparison with the European regime is more delicate, and worth reporting. Because the fast-learning cohort under Dynamic Licensing is monitored considerably more intensively than either constant-monitoring benchmark, the welfare advantage over the European regime is not parameter-free: repeating the computation across a range of values for the monitoring-cost scale c shows that Dynamic Licensing's welfare advantage over the European regime shrinks monotonically as c increases, crosses zero near $c \approx 20$, and reverses sign for still more convex monitoring-cost specifications. This finding is economically intuitive rather than troubling for the underlying theory: Proposition 1 guarantees that the *optimal* state-dependent policy $m^*(A_t)$ weakly dominates any fixed-monitoring benchmark, but the hazard-based cohort-switching mechanism used to approximate that policy numerically (Section 4.1, and the accompanying methodological footnote) is a stylized proxy for $m^*(A_t)$. The crossover result therefore identifies the boundary condition under which the theoretical dominance result translates into a realized welfare gain in practice: dynamic licensing outperforms static precaution when the marginal cost of additional monitoring intensity is not too steeply convex, and the advantage narrows – and can in principle disappear – as information production becomes very expensive at the margin. This is, if anything, a further argument for the incentive-compatible, cost-minimizing information production mechanisms discussed in Section 5.3, since the practical case for dynamic licensing

strengthens directly with the regulator's ability to keep the marginal cost of monitoring low.

4.4 Robustness of the Transient Bimodality

Because the two-cohort hazard-switching mechanisms used to generate Figure 3 are a numerical approximation of the continuous state-dependent monitoring rule (see the methodological footnote in Section 4.2.3), it is useful to check whether the qualitative phenomenon it produces – a transient separation of the simulated population into a fast- and a slow-converging cohort – survives changes in the timing and intensity of the underlying switching hazard, or whether it is an artifact of the single baseline parameterization reported in Figure 3.

Repeating the Dynamic Licensing simulation under two alternative parameterizations – a "faster switching" specification, in which cohort separation is concentrated earlier and more sharply - i.e., hazard peak at $t = 10$ rather than $t = 16$, over a narrower window – and a "slower switching" specification, in which separation is more gradual and delayed – i.e., hazard peak at $t = 25$, over a wider window – confirms that a clearly bimodal cross-sectional density persists in both cases, with the timing and amplitude of maximum peak separation shifting in the expected direction: earlier and somewhat less pronounced under faster switching (maximum separation of 0.32 in the Ambiguity-Adjusted Harm Index, attained at $t = 10$), and comparable in magnitude under slower switching (maximum separation of 0.44, attained at $t = 28$, against 0.53 at $t = 28$ under the baseline parameterization). What is robust across all three parameterizations, in other words, is the qualitative fact of transient cross-sectional heterogeneity itself –

showing the existence of a temporary window in which the simulated population separates into statistically distinguishable fast- and slow-learning cohorts before eventually re-converging – rather than the specific period at which that separation is maximal, which is naturally sensitive to how quickly the underlying regulatory learning process is assumed to operate. This evidence supports treating the transient bimodality of Figure 3 – i.e., of dynamic licensing structure – as a structural feature of state-dependent monitoring under the proportional-learning specification of Section 4.1, rather the result of any numerical artifact of one particular choice of hazard parameters.

5. From Theory to Institutional Design: Implementing Dynamic Licensing

5.1 From Optimal Policy to Dynamic Architectures

The formal analysis developed in Chapter 3 established that, under conditions of scientific ambiguity characteristic of emerging nanotechnologies, a licensing regime in which authorization status, monitoring intensity, and regulatory obligations evolve continuously as functions of an ambiguity-adjusted harm index – i.e., A_t – dominates both the traditional precautionary model based upon fixed *ex ante* authorization and the conventional *post-market* model based upon essentially static approval followed by ex post enforcement. The result is derived as a theorem of constrained welfare maximization that identifies the properties of an optimal regulatory mechanism under conditions of evolving scientific knowledge and demonstrates that, over the relevant

parameter space, a dynamically adaptive licensing regime might achieve a higher level of expected social welfare than either of the two institutional paradigms currently prevailing in the European Union and the United States. The institutional discussion that follows is framed around the nanotechnology context introduced in Section 2, both because that context motivated the analysis and because it offers a concrete setting against which abstract delegation principles can be tested. The underlying institutional logic, however, is not specific to nanomaterials, and Section 6 discusses its extension to other domains characterized by comparable epistemic conditions.

Yet the conclusion reached in previous Chapters remains, in an important sense, institutionally incomplete and two central questions arise: Which legal institutions are capable of performing the functions that, within the formal model, are represented by continuously updated regulatory variables? Conversely, which institutional arrangements would necessarily destroy the incentive structure upon which the model's optimality result depends? Viewed from this perspective, the implementation of a dynamic form of licensing becomes not a question of selecting particular legal instruments but of allocating institutional functions[85].

[85] It should be emphasized that the architecture proposed here is not without institutional precedent, even if no existing instrument implements it in the fully continuous, state-contingent form characterized by Proposition 1. Both the European Union and the United States already operate licensing mechanisms that condition initial market access on an explicit commitment to post-authorization evidence generation, coupled with a standing power to revise, restrict, or withdraw authorization as that evidence accumulates. In the European Union, the conditional marketing authorization mechanism under Article 14-a of Regulation (EC) No 726/2004 permits pharmaceutical products addressing unmet medical needs to reach the market on the basis of less complete data than would ordinarily be required, subject to specific obligations to complete post-authorization studies within a defined timeframe and to annual reassessment by the European Medicines Agency; the PRIME (PRIority MEdicines) scheme operationalizes an analogous logic further upstream, providing enhanced regulatory support and more intensive dialogue precisely for products whose risk profile remains most uncertain at the point of initial evaluation. In the United States, the FDA's accelerated approval pathway (21 C.F.R. §§ 314.510, 601.41) and Breakthrough Therapy designation (established by the Food and Drug Administration Safety and Innovation Act of

5.2 Institutional Competence and the Architecture of Dynamic Delegation

The translation of this form dynamic delegation into positive law depends upon the regulator's ability to revise authorization conditions continuously as scientific knowledge evolves. The question, in this perspective, is identifying the constitutional, administrative, and economic constraints and determine which components of the regulatory mechanism are sufficiently political to require legislative determination and which are sufficiently technical to justify administrative implementation.[86]

For the American context, this question has become a live constraint of positive constitutional law, given that federal agencies assert jurisdiction over nanomaterials through general-purpose statutes – TSCA's "unreasonable risk" standard, the FDA's general safety mandates – that were drafted without nanotechnology, or continuous

---

2012, 21 U.S.C. § 356(a)(4)) follow a structurally similar logic, conditioning early authorization on surrogate endpoints upon a binding commitment to confirmatory post-market trials, with withdrawal a standing legal possibility – and, on occasion, an exercised one – where confirmatory evidence fails to materialize or proves unfavorable. Outside pharmaceutical regulation, an increasing number of jurisdictions have adopted regulatory sandboxes that make authorization explicitly provisional and evidence-contingent for other technologies characterized by comparable ambiguity, from the UK Financial Conduct Authority's fintech sandbox (launched in 2016) to the regulatory sandbox provisions for artificial intelligence systems introduced by Articles 57-63 of Regulation (EU) 2024/1689 (the EU AI Act). Nonetheless, none of these instruments formalizes monitoring intensity as a continuous function of an ambiguity-adjusted harm index in the manner of $m^*(A_t)$, and none derives its review thresholds from an explicit welfare optimization of the kind developed in Section 3. In this regard, the dynamic licensing framework proposed here is a formalization and generalization of an institutional logic that already exists in fragmentary, sector-specific form, showing that the constitutional and administrative legal questions addressed in Section 5.2 are questions this fragmentary practice has already had to confront, although without the support of the unified theoretical framework presented here.

[86] The problem may be stated more generally. Legislatures possess the democratic legitimacy necessary to establish normative priorities, determine the objectives of public policy, and define the legal boundaries within which public administration may operate. Administrative agencies, by contrast, possess neither equivalent democratic legitimacy nor equivalent political discretion. Their institutional advantage lies in specialized expertise, the capacity to process large quantities of technical information, respond to evolving empirical evidence, and update regulatory decisions with a frequency that legislative institutions cannot realistically sustain. For a comprehensive discussion of these issues cf., among others, Harold H Bruff, 'Legislative Formality, Administrative Rationality' (1984) 63 *Texas Law Review* 207; Mathew D McCubbins, Roger G Noll and Barry R Weingast, 'Structure and Process, Politics and Policy: Administrative Arrangements and the Political Control of Agencies' (1989) 75 *Virginia Law Review* 431; Marc Allen Eisner, *Regulatory Politics in Transition* (2nd edn, Johns Hopkins University Press 2000).

ambiguity-contingent monitoring, in contemplation. Under the major questions doctrine articulated by the Supreme Court in *West Virginia v Environmental Protection Agency*, an agency asserting the power to resolve a question of "vast economic and political significance" must point to clear congressional authorization for that specific power, and courts should hesitate to infer such authorization from otherwise ambiguous or general statutory language.[87] A dynamic licensing regime of the kind characterized by Proposition 1 – in which an agency continuously recalibrates monitoring intensity and authorization status for an entire, economically significant technological sector on the basis of an internally computed ambiguity index, without express statutory reference to any such index or mechanism – is a plausible candidate for exactly this kind of judicial skepticism, particularly once the technology in question reaches the economic scale contemplated throughout this paper. The decision in *Loper Bright Enterprises v Raimondo*, overruling the deference framework of *Chevron USA Inc v Natural Resources Defense Council Inc*, reinforces the same constraint from a different angle: courts, not agencies, now bear ultimate responsibility for determining the best reading of an ambiguous statute, which forecloses the possibility that an agency could adopt the dynamic licensing framework developed in this paper as one reasonable interpretation, among others, of an existing general grant of regulatory authority.[88] The problem might

[87] Cf., in this perspective, West Virginia v Environmental Protection Agency, 597 US 697 (2022). See also, for the doctrine's earlier articulation and subsequent consolidation, Utility Air Regulatory Group v Environmental Protection Agency, 573 US 302 (2014); FDA v Brown & Williamson Tobacco Corp, 529 US 120 (2000); and, for academic discussion, Cass R Sunstein, 'There Are Two "Major Questions" Doctrines' (2021) 73 Administrative Law Review 475, and Daniel T Deacon and Leah M Litman, 'The New Major Questions Doctrine' (2023) 109(4) Virginia Law Review 1009.

[88] Cf., in these regards, Loper Bright Enterprises v Raimondo, 603 US 369 (2024), overruling Chevron USA Inc v Natural Resources Defense Council Inc, 467 US 837 (1984), and, for a discussion of the implications for technology-specific regulatory delegation of the kind considered in this paper, see Kristin

eventually be that each of the current real-world instruments – e.g., conditional marketing authorization, accelerated approval, statutory regulatory sandboxes[89] – rests on express legislative authorization for the specific evidence-contingent mechanism the agency then administers, rather than on an agency's own interpretation of a general safety mandate.[90]

The optimization model, in particular, implicitly assumes the existence of an institutional actor capable of continuously observing the evolution of the information filtration, recalculating the ambiguity-adjusted harm index, modifying monitoring obligations, and revising authorization status whenever new evidence justifies regulatory intervention. Within the model these operations appear as straightforward transformations of state variables. Within a constitutional order they might become exercises of public authority, with additional challenges in terms of autonomy and control,[91] given that the model distinguishes between structural parameters and continuously evolving variables. Nevertheless, while structural parameters define the

E Hickman and Amy J Wildermuth, 'Harmonizing Delegation and Deference After Loper Bright' (2025) 100(6) New York University Law Review 1924; Michelle Benedetto Neitz, 'Ready or Not: How Congressional Dysfunction and Loper Bright Enterprises v Raimondo Will Shift US Regulation of Emerging Technologies to the Federal Bench' (2025) 78(1) SMU Law Review 119; Blake Emerson, 'Administrative Answers to "Major Questions": On the Democratic Legitimacy of Agency Statutory Interpretation' (2018) 102(5) Minnesota Law Review 2019.

[89] See note 85 for further reference.

[90] The lesson for implementing dynamic licensing in the American context, therefore, is that the threshold $\overline{A}$, the two-threshold hysteresis band of Proposition 2, and the basic structure of the license menu characterized by Proposition 3 are the kind of major, economically significant policy choices that – under current doctrine – Congress would need to authorize expressly, with the responsible agency retaining discretion only over the genuinely technical, recurring recalibration of monitoring intensity within that legislatively specified structure. This is, in substance, precisely the allocation of institutional competence for which Section 5.2 argues below on independent, welfare-theoretic grounds; the major questions doctrine supplies an additional, doctrinally grounded reason for drawing the line in the same place.

[91] See, in this perspective, Tom Christensen and Per Lægreid, 'Regulatory Agencies – The Challenges of Balancing Agency Autonomy and Political Control' (2007) 20(3) *Governance* 499.

architecture of the mechanism itself most of the objectives of the regulatory system, the interests receiving legal protection, and the procedural guarantees attaching to regulatory decisions, remain stable throughout the operation of the mechanism. These elements are not updated in response to new scientific evidence. What changes dynamically is the authorization status of nanotechnologies, which is updated according to how the expected welfare consequences of continued commercialization change. None of these adjustments alters the normative objectives pursued by the legal system. They merely improve the regulator's ability to pursue those objectives under progressively richer informational conditions.

The interesting aspect and particular urgency within the regulatory governance of nanotechnology is that the informational environment differs fundamentally from that characterizing more mature technological sectors, given that – in this sector – scientific uncertainty is endogenous to the regulatory process itself. In the field of nanotechnologies, new evidence may confirm previous assessments, contradict established assumptions, reveal previously unknown exposure pathways, or fundamentally alter prevailing toxicological models,[92] but, under such conditions, any attempt to specify ex ante every regulatory consequence within primary legislation necessarily assumes a stability of scientific knowledge that the underlying model

[92] Cf., *ex plurimis*, Kirsten Gerloff, Brigitte Landesmann, Andrew Worth, Sharon Munn, Taina Palosaari and Maurice Whelan, 'The Adverse Outcome Pathway Approach in Nanotoxicology' (2017) 1 *Computational Toxicology* 3; Richard D Handy and Benjamin J Shaw, 'Toxic Effects of Nanoparticles and Nanomaterials: Implications for Public Health, Risk Assessment and the Public Perception of Nanotechnology' (2007) 9(2) *Health, Risk & Society* 125; James E Trosko and Brad L Upham, 'A Paradigm Shift Is Required for the Risk Assessment of Potential Human Health after Exposure to Low Level Chemical Exposures: A Response to the *Toxicity Testing in the 21st Century* Report' (2010) 29(4) *International Journal of Toxicology* 344.

explicitly rejects. The consequence of this instability is an institutional paradox. The more uncertain the scientific environment becomes, the less capable legislation is of determining detailed regulatory outcomes in advance. Yet the same uncertainty simultaneously increases the importance of maintaining legally constrained decision-making, since regulated actors require sufficiently stable expectations to justify investment, innovation, and compliance.

This concept of "*dynamic delegation*" might differ from conventional understandings of administrative discretion. Classical administrative law frequently conceives delegation as a transfer of decision-making authority subject to procedural constraints in terms of discretion and deference.[93] In this context, dynamic delegation requires defining a specific architecture of an adaptive regulatory mechanisms in which objectives, evaluative criteria, procedural safeguards, reviewability, and permissible methods of scientific assessment – already been specified through legislation – and agencies retain substantial scientific discretion, remaining institutionally incapable of redefining the underlying normative commitments of the regulatory system. The incentives compatibility of such a regime depends not only upon the regulator's ability to revise decisions but also upon the regulated firm's expectation that revisions will occur according to transparent and legally predictable criteria. If administrative authorities possessed unlimited discretion to modify authorization conditions in response to political pressure or changing regulatory preferences, rational innovators would discount

[93] See, in this light, Nuno Garoupa and Jud Mathews, 'Strategic Delegation, Discretion, and Deference: Explaining the Comparative Law of Administrative Review' (2014) 62(1) *The American Journal of Comparative Law* 1.

the expected benefits of voluntary information disclosure, with compliance investments resembling sunk costs vulnerable to arbitrary revision rather than contributions to a stable cooperative regulatory relationship.[94]

5.3 Incentive Compatibility and the Legal Production of Regulatory Information

The institutional architecture developed thus far presupposes that the regulator is able to continuously acquire information regarding the evolving safety profile of the regulated nanomaterial. This observation introduces a dimension of the implementation problem that cannot be resolved through administrative procedures alone. Scientific information does not exist independently of the institutional arrangements through which it is generated[95] but much of the information required for adaptive regulation originates outside the public administration.[96] In this perspective, the regulatory framework might itself determine whether socially valuable information will be generated, disclosed, verified, and incorporated into subsequent decision-making. Traditional administrative law frequently conceives regulatory obligations as mechanisms for ensuring compliance with pre-existing legal standards, while dynamic regulatory constraints require to transform private observation into public regulatory knowledge, thereby permitting

[94] For a a more complete overview in this direction, see Shlomi Sher, Johannes Müller-Trede and Craig R M McKenzie, 'Choices without Preferences: Principles of Rational Arbitrariness' (2025) 132(6) *Psychological Review* 1375.

[95] Cf. Fritz Sager and Anat Gofen, 'The Polity of Implementation: Organizational and Institutional Arrangements in Policy Implementation' (2022) 35(2) *Governance* 347; Yoon Kim and Jeffrey M Stanton, 'Institutional and Individual Influences on Scientists' Data Sharing Practices' (2012) 3(1) *Journal of Computational Science Education* 47; Lorna Schrefler, 'The Usage of Scientific Knowledge by Independent Regulatory Agencies' (2010) 23(2) *Governance* 309; M Granger Morgan and Jon M Peha, 'Analysis, Governance, and the Need for Better Institutional Arrangements' in M Granger Morgan and Jon M Peha (eds), *Science and Technology Advice for Congress* (Routledge 2003) 3.

[96] To make some example, evidence concerning the long-term behavior of novel nanomaterials is produced through toxicological studies, post-market surveillance, epidemiological observation, environmental monitoring, adverse-event reporting, and private research conducted by regulated firms themselves.

future licensing decisions to approach the welfare optimum identified *supra*. The distinction has significant institutional consequences. Scientific knowledge possesses the characteristics of a public good, because, once generated and verified, it benefits not merely the innovator but also competing firms, downstream users, consumers, and future applicants developing technologically related products.[97] From this point of view, the regulator might confront itself with strategic information asymmetry[98] and the central implication of the formal model is that legal systems should attempt to organize incentives toward the continuous production of reliable scientific evidence.[99]

Note that this conceptual shift does not imply reducing regulatory oversight, presupposing, on the contrary, robust public authority. Private information becomes socially valuable only because disclosure occurs within a legally structured environment characterized by independent scientific verification, mandatory reporting requirements, audit powers, procedural accountability, and credible sanctions for strategic manipulation.[100] From this standpoint, given that dynamic licensing evolves throughout

[97] See, on this point, Joseph E Stiglitz, 'Knowledge as a Global Public Good' in Inge Kaul, Isabelle Grunberg and Marc A Stern (eds), *Global Public Goods: International Cooperation in the 21st Century* (Oxford University Press 1999) 308 and Stephen Morris and Hyun Song Shin, 'Social Value of Public Information' (2002) 92(5) *American Economic Review* 1521.

[98] This problem has been highlighted, *ex plurimis*, in David P Baron and David Besanko, 'Regulation, Asymmetric Information, and Auditing' (1984) 15(4) *RAND Journal of Economics* 447.

[99] Public investment in universities and academic-driven research and innovation has been shown to have a positive impact on the reduction of barriers to entrepreneurship and information spillovers, as evidenced, among others, by Stefano Baruffaldi, Markus Simeth and David Wehrheim, 'Asymmetric Information and R&D Disclosure: Evidence from Scientific Publications' (2024) 70(2) *Management Science* 1052 and Daniel Feser and Till Proeger, 'Asymmetric Information as a Barrier to Knowledge Spillovers in Expert Markets' (2017) 13(1) *International Entrepreneurship and Management Journal* 211.

[100] In this perspective, to avoid regulatory capture, it is essential to maintain the highest level of scientific integrity, particularly given the sensitivity of the topic of nanotechnologies. See, on this topic, in particular, Genna Reed, Yogi Hendlin, Anita Desikan, Taryn MacKinney, Emily Berman and Gretchen T Goldman, 'The Disinformation Playbook: How Industry Manipulates the Science-Policy Process–and How to Restore Scientific Integrity' (2021) 42(4) *Journal of Public Health Policy* 622.

the life of the license itself, rather than terminating upon initial market approval, the regulated firm repeatedly confronts with new opportunity for information produced during commercialization to influence future informational outcomes, with the result that where administrative authorities disregard favorable evidence, alter evaluative standards retrospectively, or exercise scientific discretion unpredictably, incentives to generate additional information rapidly deteriorate. The practical implication is that the regulatory relationship should be structured as a sequence of repeated interactions in which disclosure decision shapes future credibility and the licensing regime acquires the characteristics of an *ongoing institutional relationship* in which information generation, regulatory adaptation, and commercial activity evolve simultaneously.

6. Conclusions

The analysis developed throughout this study originated from an apparently sector-specific question: whether existing regulatory frameworks governing engineered nanomaterials remain capable of providing an efficient and legally coherent response to technologies characterized by persistent scientific uncertainty. The comparative examination of European Union and United States law undertaken in Chapter 2 demonstrated that, despite important differences in institutional design, both jurisdictions continue to rely upon regulatory architectures whose fundamental logic remains static. Market authorization is still conceived primarily as the legal consequence of a sufficiently complete ex ante assessment, while post-market intervention operates largely as an exceptional corrective mechanism. Such an architecture reflects the implicit assumption that the principal function of regulation is to transform uncertainty

into certainty before commercialization occurs. Nevertheless, this assumption no longer corresponds to the conditions under which emerging technologies are developed and governed. In the case of engineered nanomaterials, scientific knowledge requires an ongoing interaction between laboratory research, commercial deployment, environmental exposure, epidemiological observation, and regulatory oversight itself. The relevant uncertainty, in this sector, concerns the continual coevolution of scientific understanding between the industry and the regulatory framework,[101] with the necessity to accommodate a more dynamic informational environment in an institutional structure capable of understanding – before regulating – emerging technologies.

From this diagnosis follows the principal theoretical contribution of the present work. Rather than treating the shortcomings of current regulatory regimes as isolated defects capable of correction through more demanding pre-market testing, expanded precautionary obligations, or intensified post-market surveillance considered independently, this paper has argued that the underlying problem should be framed as an institutional regulatory matter. Existing norms organize legal decision-making around a static conception of scientific knowledge, in the attempt to govern dynamic technological processes. The resulting tension manifests itself simultaneously as excessive regulatory rigidity, insufficient responsiveness to newly emerging evidence, inefficient allocation of monitoring resources, and weakened incentives for the production of socially valuable information.

[101] See, in these terms, Charles Sabel, Gary Herrigel and Peer Hull Kristensen, 'Regulation under Uncertainty: The Coevolution of Industry and Regulation' (2018) 12(3) *Regulation & Governance* 371.

To address this structural problem, we have developed an integrated theoretical framework combining economic analysis, decision theory under ambiguity, mechanism design, and economic law, demonstrating formally that, under conditions of persistent ambiguity, a regulatory regime continuously adjusting authorization status, monitoring intensity, and regulatory obligations in response to an ambiguity-adjusted harm index dominates both conventional precautionary regulation and purely reactive post-market supervision, in the sense that the optimal state-dependent policy characterized in Proposition 1 can never do worse, and generically does strictly better, than either fixed-monitoring benchmark. Two extensions refine this baseline result. First, once the realistic administrative costs of granting and revoking authorization are incorporated into the dynamic program, the optimal licensing rule is shown to take the form of a two-threshold hysteresis band rather than a single frictionless cutoff (Proposition 2), formalizing the intuition that provisional legality should be sticky, not volatile. Second, a minimal static extension of the model to settings in which the innovator holds private information about the technology's risk type characterizes the socially optimal incentive-compatible license menu, showing that truthful revelation requires monitoring the riskier type below its own first-best level and leaving an information rent to the safer type (Proposition 3). Unlike traditional binary authorization systems, a dynamic delegation system permits regulatory intervention to evolve as scientific information accumulates, thereby reducing both over-regulation and under-regulation while improving, at the same time, overall social welfare.

The simulations presented in Chapter 4 provided quantitative support for these analytical results, while also revealing a genuine limit on their unconditional application. Although

necessarily dependent upon simplifying assumptions, the computational analysis confirmed that dynamic licensing generates a large and robust welfare improvement over the American, post-market benchmark across the parameter values considered. The comparison with the European, precautionary benchmark proved more delicate: because the numerical approximation of the optimal monitoring rule used in the simulation carries a real informational cost, the welfare advantage of dynamic licensing over static precaution narrows as the marginal cost of monitoring becomes more steeply convex and can – in principle – reverse for sufficiently costly information production. This finding does not contradict – however – the theoretical dominance result of Proposition 1 – which concerns the true optimal policy $m^*(A_t)$, not its numerical approximation – but it does mean that the practical case for dynamic licensing over static precaution is strongest, and should be argued as such, precisely where incentive-compatible, cost-minimizing information production mechanisms are most effective at keeping the marginal cost of monitoring low. More importantly, the simulations illustrated the dynamic mechanisms through which adaptive regulation gradually improves decision quality by reallocating monitoring effort toward periods and products characterized by the greatest informational value, and a supplementary robustness check confirmed that the transient cross-sectional heterogeneity this generates is not an artifact of one particular parameterization of the underlying learning process. In doing so, they reinforced one the central claims of this paper: effective regulation under uncertainty depends both upon the quantity of scientific evidence available and the institutional capacity to incorporate that evidence continuously into regulatory decision-making –

and, no less importantly, upon keeping the cost of producing that evidence low enough that flexibility remains worth its price.

The analysis would nevertheless have remained incomplete had it ended with the demonstration of an economically superior regulatory mechanism. Optimal policy does not automatically become legally implementable policy. For this reason, Chapter 5 addressed the transformation of the abstract optimization model into an administratively operational legal framework. The institutional architecture developed there sought to demonstrate that the welfare properties identified by the formal model survive only if implementation preserves two complementary forms of action. On the one hand, it might be required by legislatures to determine the normative objectives and structural parameters of the regulatory system, while leaving scientifically specialized administrative authorities to retain responsibility for the continuous recalibration of regulatory decisions as knowledge evolves. On the other hand, under conditions of persistent scientific ambiguity, the stability of substantive regulatory outcomes becomes increasingly difficult to reconcile with rational regulation whenever scientific knowledge evolves continuously after commercialization. In this perspective, it would be necessary to incentivize scientific knowledge as regulatory information.

The broader significance of the proposed framework extends beyond nanotechnology itself. Although engineered nanomaterials provided the empirical context motivating this research, the institutional logic developed throughout this paper addresses a more general category of regulatory problems increasingly characteristic of contemporary technological innovation. Artificial intelligence, advanced biotechnology, synthetic biology, quantum technologies, autonomous systems, and other emerging technological

domains share an important structural characteristic: scientific understanding evolves throughout commercialization rather than preceding it. In each of these contexts, regulatory institutions face the challenge of governing technologies whose risk profiles remain partially indeterminate at the moment legal authorization must nevertheless be granted. The institutional principles developed in this paper – e.g., adaptive authorization, continuous evidence generation, incentive-compatible information disclosure, and procedurally structured regulatory revision – may therefore possess wider applicability than the specific regulatory context from which they originated.

This broader perspective also suggests a more general theoretical conclusion concerning the relationship between law and scientific uncertainty. Throughout much of modern regulatory frameworks, uncertainty has frequently been conceptualized as an obstacle to effective regulation, a temporary condition to be reduced before sufficiently reliable legal decisions become possible. The analysis here developed shows that persistent uncertainty should not necessarily be regarded as a failure of scientific knowledge or regulatory design. In rapidly evolving technological sectors it constitutes a normal structural feature of the regulatory environment. The objective of legal institutions should therefore not be to postpone regulatory intervention until uncertainty disappears – an objective that may never be attainable – but to develop institutional arrangements capable of learning while regulating. Economic optimization and legal implementation have often been treated as analytically separate stages: economists identify efficient regulatory outcomes, while lawyers determine whether existing institutions can implement them. The approach developed here suggests that the institutional architecture through which regulatory decisions are produced forms part of the

optimization problem itself. Incentives, regulatory competences, procedural guarantees, and mechanisms for generating reliable information are not external legal constraints imposed upon an economic model but indispensable components of the conditions under which welfare-maximizing regulatory outcomes become attainable.