A Systems Engineering Approach to Regulating Autonomous Systems

by

David Paul Britton

Department of Mechanical Engineering and Materials Science
Duke University

Date:_____

Approved:

_____
Mary Cummings, Supervisor

_____
Patrick Codd

_____
Michael Waitzkin

Thesis submitted in partial fulfillment of
the requirements for the degree of
Master of Science in the Department of
Mechanical Engineering and Materials Science in the Graduate School
of Duke University

2017

ABSTRACT

A Systems Engineering Approach to Regulating Autonomous Systems

by

David Paul Britton

Department of Mechanical Engineering and Materials Science
Duke University

Date:_____

Approved:

_____
Mary Cummings, Supervisor

_____
Patrick Codd

_____
Michael Waitzkin

An abstract of a thesis submitted in partial
fulfillment of the requirements for the degree
of Master of Science in the Department of
Mechanical Engineering and Materials Science in the Graduate School of
Duke University

2017

Abstract

Autonomous systems are emerging across many industries. From unmanned aircraft to self-driving cars to closed-loop medical devices, these systems offer great benefits but also pose new risks. Regulators must grapple with how to manage these risks, challenged to keep pace with technological developments and exhibit appropriate precaution without stifling innovation. Seeking inspiration for a viable approach to the regulation of autonomous systems, this thesis draws from the practices of systems engineering, an interdisciplinary field of engineering aimed at managing the risks of complex projects. By comparing systems engineering practices to regulatory options, current regulations, and the inherent challenges of regulating emerging technologies, this thesis concludes that a systems engineering-based approach to regulating autonomous systems offers great potential for managing the risks of autonomous systems while also driving innovation.

Contents

List of Tables

List of Figures

Acknowledgements

Thanks first to my parents, without whom I would not be privileged enough to spend my time writing about such things as regulating our future robot overlords. Thanks to Professor Cummings for her direction and guidance on this thesis and for taking in a mutt looking for shelter in the engineering school. The help of the rest of my committee—Buz Waitzkin and Patrick Codd—is also much appreciated, as is the advice and comradery of my classmates and other members of the Duke community.

1. Introduction

Highly-automated technologies are emerging in almost all industries. At the highest levels of automation, *autonomous systems* take in information about the unstructured world around them, process that information to analyze the problems they face, and use that analysis to make and act upon decisions in the face of uncertainty, all without significant human intervention. Emerging autonomous systems include self-driving vehicles, which use sensors to view nearby obstacles, processing that information along with stored mapping data in order to safely navigate to a desired destination[1]; artificial intelligence-based financial trading systems, which track market conditions and individual stocks and make independent decisions on when to buy or sell[2]; and even new medical devices which monitor a patient's physiological condition and alter the rate of drug delivery or direct other medical intervention without caregiver input.[3]

Autonomous systems hold great promise—as a recent White House report put it, these technologies have "the potential to help address some of the biggest challenges that society faces."[4] Potential benefits of autonomous systems include increased access

---

[1] *E.g.*, Waymo (formerly the Google self-driving car project) at waymo.com.
[2] *E.g.*, Tech Trader AI at https://www.techtrader.ai/
[3] *E.g.*, Medtronic MiniMed closed-loop insulin pump, at medtronicdiabetes.com.
[4] Executive Office of the President, "Preparing for the Future of AI", at Introduction, 2016. (Obama Adminsitration) https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf

to education and justice, public health, mobility, and transportation safety.[5] But autonomous systems also come with potential negative consequences. Calls for increased government regulation of autonomous systems are growing more prominent.[6]

Technology regulation typically focuses on lowering risks: that is, on reducing the potential negative consequences associated with an industry, activity, or product. Because the economic assumption is that firms will maximize their own efficiency, one view of technology regulation is that it always limits the use of a technology that would otherwise be the best choice for an individual industry participant.[7] To the extent that regulation can affect innovation at all, then, the effect would only be to reduce incentives to invest in more economically efficient technologies.[8] But competing research demonstrates that regulation can actually create the need for innovation, driving technological progress towards societal goals.[9] Where regulation incentivizes more efficient, or otherwise 'better' technology, regulation can actually promote the benefits of

---

[5] *Id.*

[6] Examples can be found across the world. For one high-level example, *see* "Robots: Legal Affairs Committee Calls for EU-Wide Rules," European Parliament Press Release, Jan. 2017 http://www.europarl.europa.eu/news/en/news-room/20170110IPR57613/robots-legal-affairs-committee-calls-for-eu-wide-rules

[7] *See, e.g.,* Adam B. Jaffe et al., "Environmental Regulation and the Competitiveness of U.S. Manufacturing: What does the Evidence Tell Us?", *Journal of Economic Literature* (March 1995) http://www.ucl.ac.uk/cserge/Jeffe%20et%20al%201995.pdf.

[8] *Id.*

[9] Nicholas A. Ashford and Ralph P. Hall, "Regulation-Induced Innovation for Sustainable Development", *Administrative & Regulatory Law News*, Spring 2012, 21-3; Ashford and Hall, "The Importance of Regulation-Induced Innovation for Sustainable Development" *Sustainability*, 2011, http://www.mdpi.com/2071-1050/3/1/270/pdf.

new technology rather than just protect against its risks.[10] The overarching challenge of regulating emerging technologies is in designing regulations that do both: regulations for autonomous systems should both encourage fulfillment of the technology's potential and manage the associated risks.[11]

The first step of risk regulation is risk assessment [12]: that is, the starting place is to identify the risks of autonomous systems. Generalizing across industries or applications for autonomous systems, the risks fall into three broad categories. First, as technology takes over the completion of tasks once left for humans, mistakes in completing those tasks shift from human error to technological failure. Although reducing the risk of human error is often cited as a main benefit of autonomous systems,[13] that only helps if the autonomous system is more reliable than the person was. Problematically, these are new technologies in new roles, sensing and interpreting information in as-yet-unproven ways that may not be entirely understood even by their own programmers.[14] The unpredictability of autonomous systems may mean that failure modes are unforeseen, and therefore untested and unmanaged. When autonomous technology is running a safety-critical transportation, infrastructure, or

---

[10] *Id.*
[11] *See, e.g.,* Johnathan Wiener, "The regulation of technology, and the technology of regulation", *Technology in Society* (2004), available at: http://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=1960&context=faculty_scholarship.
[12] *E.g.,* Stephen Breyer, "Breaking the Vicious Circle".
[13] *See, e.g.,* Executive Office of the President, supra note 4.
[14]Ryan Calo, "Robotics and the Lessons of Cyberlaw", Cal. L.R., 2015 http://www.californialawreview.org/wp-content/uploads/2015/07/Calo_Robots-Cyberlaw.pdf

medical system, failure could be catastrophic. Thus, the first risk to address is that autonomous systems will fail in new and dangerous ways while completing tasks that technology could never before undertake without humans.

Second, the difficulty humans face when interacting with autonomy can lead to negative consequences. Human users of autonomous tools, previously accustomed to full knowledge and control over the information relevant to their job, lose awareness of the situation in which the autonomous system is acting as it internalizes more of the important information and decision-making.[15] The user, as well as other people nearby, may have trouble interpreting what the system is doing, and why the system is doing it.[16] Additionally, when relegated to a supervisory role monitoring automated systems, people get bored and tend to over-rely on automation. Together, these factors make it more difficult for a human to direct the activities of autonomous systems, live alongside them safely, or to detect mistakes and intervene successfully.[17] Breakdowns in the interactions between human and automation have already lead to disaster.[18]

---

[15] *See, e.g.,* Mica Endsley, "Automation and Situation Awareness," Automation and human performance: Theory and applications, R. Parasuraman & M. Mouloua, eds. (1996). http://www.aerohabitat.eu/uploads/media/Automation_and_Situation_Awareness_-_Endsley.pdf.

[16] See, e.g., Jan Bredereke and Axel Lankenau, "A Rigorous View of Mode Confusion" in Computer Safety, Reliability , and Security (2002), http://link.springer.com/chapter/10.1007%2F3-540-45732-1_4

[17] *E.g.,* Mary Cummings and Jason Ryan, "Who Is in Charge? The Promises and Pitfalls of Driverless Cars." Transportation Research Board 2014.

[18] *See, e.g.,* Madeline Elish, "Moral Crumple Zones", We Robot 2016. (examples include the Three Mile Island meltdown and commercial airliner crashes).

Third, increasing automation poses a variety of risks at a societal level. Many manufacturing jobs have already been replaced by robots, and some commentators suggest that developments in autonomous systems could extend the threat of replacement to almost all jobs, from commercial trucking to the practice of law.[19] The possibility that these advances may push us towards the "end of work" has been cited as a real risk.[20] Other societal risks relate to privacy implications of ever-present sensors on ubiquitous robots,[21] the potential that autonomous systems will learn and reinforce racial biases,[22] and ethical concerns with the use of autonomous systems to make life-or-death decisions or play key educational or caregiving roles. All three categories of future risks are hard or impossible to quantitatively assess in an accurate, predictive way, but they at least give a broad understanding of what to worry about.

The second step of risk regulation is risk management: having identified the risks, how should they be dealt with?[23] Four concepts from risk management will play a role in managing the risks of autonomous systems. The first has been labelled the

---

[19] *E.g.,* Jason Millar and Ian Kerr, "Delegation, Relinquishment and Responsibility: The Prospect of Expert Robots", *Robot Law* (eds. Calo, Froomkin, Keer) 2016.
[20] "A world without work" Atlantic 2015 http://www.theatlantic.com/magazine/archive/2015/07/world-without-work/395294/
[21] Matthew Reuben and William D. Smart, "Privacy in Human-Robot Interaction: Survey and Future Work", We Robot 2016 http://robots.law.miami.edu/2016/wp-content/uploads/2015/07/Rueben_Smart_PrivacyInHRI_WeRobot2016.pdf.
[22] *E.g.,* "Discrimination by algorithm: scientists devise test to detect AI bias," Guardian 2016, https://www.theguardian.com/technology/2016/dec/19/discrimination-by-algorithm-scientists-devise-test-to-detect-ai-bias.
[23] Breyer, "Breaking the Vicious Circle."

"pacing problem": regulatory schemes struggle to keep pace with technological change.[24] The inertia provided by administrative law—for example, through requirements for notice-and-comment rulemaking—keeps agencies playing catch up, while existing laws were often drafted with a static view of technology in mind.[25] Institutional expertise also lags as, for example, robots and artificial intelligence enter industries whose traditional regulators are unfamiliar with advanced computing and must acquire the technical knowledge needed to understand autonomous systems.[26] Dealing with the pacing problem requires regulations and institutions that can adapt to the challenges of evolving technologies.[27]

A second, related concept is the "precautionary principle."[28] Where the risks of a new technology are unproven or unpredictable, the precautionary principle suggests (or in alternative formulations, demands) that the technology be banned until scientific methods show that the risks are acceptably low.[29] Most persuasive when the risks include long-term, irreversible damage, implementing the precautionary principle requires by definition defending strict regulations without scientific proof that the

---

[24] Gary Marchant, "The Growing Gap Between Emerging Technologies and the Law", in *Growing gap between emerging technologies and legal-ethical oversight*, Springer 2011.
[25] *Id.*
[26] Ryan Calo, "The Case for a Federal Robotics Commission."
[27] *See, e.g.*, Andrea Renda, et al., "Selecting and Designing European ICT Innovation Policies," *EU Science Hub*, 2016. ("The nature of the ICT ecosystem determines a growing need for flexible, adaptive regulation.").
[28] Jonathan B. Wiener, "Precautionary Principle," chapter in *Principles of Environmental Law* (Ludwig Krämer and Emanuela Orlando, eds.) of the *Encyclopedia of Environmental Law* (Michael Faure, ed.) (IUCN and Edward Elgar, forthcoming 2017)
[29] *Id.*

regulated action or technology is harmful.[30] Because of the unknowns associated with autonomous systems, deciding whether or not to follow the precautionary principle will be a necessary part of designing regulatory schemes.

A third risk management concept to keep in mind is the idea of "risk-risk tradeoffs."[31] Any tool to manage risk introduces its own, different risks. If someone decides to bike to work rather than drive, for example, she reduces her contributions to the risks of climate change but increases her risk of serious injury in the case of a traffic accident. For autonomous systems, tradeoffs amongst the three categories of risks will emerge: for example, one way to try to manage the risks of autonomous system failures is to make a human responsible for monitoring the system and intervening to stop errors, but this increases the risk that flaws in that human-machine interaction will cause problems. Tradeoffs between autonomous systems' own risks and the risks tied to outside management mechanisms must also be considered. Risk-risk tradeoffs should be carefully considered in any regulatory design.[32]

The fourth risk management topic of interest is that technology is regulated through multiple modalities, a concept key in the field of cyberlaw.[33] The first modality is law, through which government tells people or companies what they cannot make, or

---

[30] *Id.*
[31] John Graham & Johnathan Wiener, "Confronting Risk Tradeoffs", *Risk v. Risk* (1995).
[32] *Id.*
[33] Laurence Lessig, "The Law of Horse: What Cyberlaw Might Teach", Harvard Law Review, Fall 1999. https://cyber.harvard.edu/works/lessig/finalhls.pdf.

sell, or do with technology, and imposes penalties for violating those rules.[34] Law is what most people think of when they think about regulation. But three other modalities also regulate how technology is created and used.[35] Social norms often constrain behavior—shaping, for example, what you wear or how you interact with different people—and can similarly constrain the use or development of technology.[36] Markets also regulate, affecting both the behavior of consumers (higher gas prices, for example, reduce driving), and the development efforts of companies trying to create products that meet demand.[37] Lastly, architecture regulates behavior: in the real world, this means the physical layout of space, rooted in city planning, roads and infrastructure design, or the arrangement of goods in a store, all of which shape how people go about their lives and the ways technology fits into it.[38] In cyberspace, architecture is more malleable, allowing computer code to shape use such that technology design regulates user behavior.[39] Regulation of autonomous systems will take all of these forms.

Regulatory design for autonomous systems that takes into account all of these considerations—the pacing problem, preferences for the precautionary principle, risk-risk tradeoffs, the four modalities for regulation, multiple categories of risk, and the

---

[34] *Id.* at 506-7
[35] *Id.*
[36] *Id.* at 507
[37] *Id.*
[38] *Id.*
[39] *Id.*

desire to promote technological opportunities while minimizing risks—is a big ask. In a search for inspiration, this thesis looks at how engineers manage these risks and considerations in practice. Managing an engineering program involves pushing the bounds of innovation—so that the new product is competitive and fulfills customer demands for new technology—while also managing a host of risks related to development, use, and system failures. Best practices for engineering complex systems, then, can provide a guide for regulatory intervention which hopes to facilitate innovation while minimizing risks.

The engineering management principles detailed in this thesis are grounded in an interdisciplinary field known as *systems engineering*, which "focuses on how to design and manage complex systems over their life cycles."[40] A complex system, in this use of the term, is a collection of technical elements which work together to form the fully-functioning end product.[41] A robot, for example, usually includes many discrete elements—e.g., sensors, motors, end effectors, computer processors—which all need to be integrated into a single system for the robot to function as intended. To guide

---

[40] Blanchard and Blyler, Systems Engineering Management 1.3.2 (quoting Wikipedia) https://books.google.com/books?id=AEaXCwAAQBAJ&pg=PT47&lpg=PT47&dq=Systems+engineering+is+an+interdisciplinary+field+of+engineering+and+engineering+management+that+focuses+on+how+to+design+and+manage+complex+systems+over+their+life+cycles.&source=bl&ots=0Q31mqFoTt&sig=-NpUiSc6xsm-nBNrZu6nN_HFS1Q&hl=en&sa=X&ved=0ahUKEwi0gaTT8sbRAhXBSyYKHYjbAaUQ6AEIWjAM#v=onepage&q=Systems%20engineering%20is%20an%20interdisciplinary%20field%20of%20engineering%20and%20engineering%20management%20that%20focuses%20on%20how%20to%20design%20and%20manage%20complex%20systems%20over%20their%20life%20cycles.&f=false

[41] http://www.incose.org/AboutSE/WhatIsSE

development of these multifaceted systems, systems engineering takes a top-down approach, eliciting stakeholder needs, defining product requirements, and carefully tracing those requirements through design, implementation, and testing phases to validate that the end product meets expectations.[42] The systems-level thinking involved helps engineers address even the broadest levels of risk, and is considered key to the successful development of complex products.[43]

Chapter 2 provides background information on system engineering and details the steps of a systems engineering plan. Chapter 3 takes that systems engineering plan and offers options for regulatory intervention at each step, detailing how regulators might choose to leverage systems engineering practices to managing different kinds of risk. Chapter 4 then asks to what degree existing regulatory schemes already use systems engineering to manage risk, focusing on the aerospace, medical device, and automotive industries as likely fields of application for autonomous systems. Chapter 5 draws conclusions and makes recommendations for future application of these lessons. Together, these chapters show how systems engineering—or at least some of its elements—offer opportunities for regulators to manage the risks of emerging technology while also guiding innovation towards policy goals.

---

[42] *E.g.,* Blanchard and Blyler, Systems Engineering Management 1.3.2

[43] *See, e.g.,* Kristen Baldwin, Message from the DASD(SE), http://www.acq.osd.mil/se/ ("As the complexity of our systems has increased, so has the need for effective systems engineering throughout the life cycle.").

2. Systems Engineering and the V-Model

For companies developing innovative products, risks include delays in development, losing out to competitor firms that have better technology that customers prefer, and the financial and reputational impact of product failures linked to injuries, deaths, or other harms. Although seen through a different lens, these risks align with the concerns of policymakers who wish to promote innovation while avoiding product failures and keeping down regulatory burdens. To explore in detail how engineering firms deal with risk, this chapter provides an introduction to the field of systems engineering.

2.1 Systems Engineering: Origin, Definition, and Purpose

Historians of technology speak of a transformation in the nature of invention across the first half of the twentieth century.[1] Inventions at the turn of the century like the Edison light bulb, the Rontgen x-ray, and the Wright brothers' flight, were attributed to heroic geniuses who could independently tackle all aspects of a technical problem.[2] At the same time, however, the growing volume of scientific knowledge began to shape the structure of knowledge into silos of specialized expertise. At American universities, disciplines like mechanical, civil, and electrical engineering were set up as separate,

[1] "History of Technology", Eric Schatzberg, lectures Spring 2013, Madison, Wisconsin; "Science in the 20th Century", Richard Staley, lectures Fall 2014, Madison, Wisconsin.
[2] *See, e.g.*, "American Inventors: A History of Genius", Time-Life 2016.

distinct departments alongside older fields like physics, math, and chemistry.[3] Many new discoveries were fueled by detailed expertise in relatively narrow fields. But, in order to solve multifaceted problems, this level of specialization also required coordination between experts from different disciplines.

In the private sector, corporate research entities like Bell Labs and the General Electric Research Labs were formed as interdisciplinary teams to push the bounds of both basic research and to generate commercializable products based on complex technologies.[4] Individual inventive genius, the story goes, was replaced by teams of coordinated experts.[5] This trend culminated during World War II with the Manhattan Project, which organized over a hundred-thousand people across multiple facilities and specialties with the single goal of producing the atomic bomb.[6] Overcoming organizational challenges inherent in interdisciplinary development of complex technologies became key to innovation, especially with projects rushed by the war effort.[7]

---

[3] Richard Staley, lectures Fall 2014, Madison, Wisconsin; University history websites also verify: http://me.engin.umich.edu/about/history; https://engineering.purdue.edu/Engr/AboutUs/History; http://web.mit.edu/cheme/about/history.html; *See also,* C.P. Snow "The Two Cultures", 1959.
[4] Staley and Schatzberg, *supra* note 1.
[5] Staley and Schatzberg, *supra* note 1.
[6] Many accounts exist: for one unique angle that sheds light on the scale of the operation, see Denise Kiernen, "The Girls of Atomic City," Touchstone 2013.
[7] Alexander Kossiakoff and William N. Sweet, "Systems Engineering: Principles and Practice", Wiley 2003, 6.

Growing out of this historical moment, the term "systems engineering" was first used during the war years.[8] Gaining traction within the Department of Defense in the following years, systems engineering slowly gained prominence as the unifying engineering discipline that functioned to guide the engineering of complex systems.[9] Meanwhile, Cold War arms races in aircraft, jet propulsion, ballistic missiles, and control systems drove the development of complex technology.[10] With the advent of solid-state electronics, and more recently, world-wide connectivity, technological systems have grown in complexity in terms of the number of components and functions that make up a single product.[11] Systems engineering grew up alongside these technological changes, and today plays an essential role in the development of many technological systems.[12]

Today, definitions of the term "systems engineering" abound in the literature,[13] but Wikipedia's suffices as a starting point: "Systems engineering is an interdisciplinary field of engineering and engineering management that focuses on how to design and manage complex systems over their life cycles."[14] To clarify, the word "system" refers to an "assemblage or combination of things or parts forming a complex or unitary

---

[8] 'History of Systems Engineering", International Council on Systems Engineering, http://www.incose.org/AboutSE/history-of-systems-engineering
[9] Kossiakoff & Sweet at 6.
[10] *Id.*
[11] *Id.*
[12] International Council on Systems Engineering, note 8.
[13] *See* Blanchard, "Systems Engineering Management" at Section 1.3.2, quoting multiple definitions.
[14] "Systems Engineering", https://en.wikipedia.org/wiki/Systems_engineering

whole,"[15] while references to "engineering" imply the definition is talking about the creation of man-made things rather than study of ecological phenomena.[16] In other words, then, systems engineering guides people and organizations as they create new technologies made up of many interacting components. Rather than concentrating on an individual component, systems engineers think about the system as a whole.[17]

Systems engineering adds value to product development in several ways, each stemming from a systems engineer's perspective on the full system. System-level thinking allows a systems engineer to consider all objectives for the project.[18] Considerations include the business case, stakeholder needs, user preferences, regulatory compliance, production issues, and technical limitations or possibilities related to an idea for a new product.[19] Thus, one role for the systems engineer is to aggregate the expectations or needs of stakeholders for the system, and communicate them to designers and engineers to ensure that decisions made throughout design and development take those goals into account.[20] Systems engineers thus contribute to system success by making sure the final product matches stakeholder needs.[21]

---

[15] "System", definition at http://www.dictionary.com/browse/system?s=t.
[16] Kossiakoff & Sweet at 3.
[17] Mary Cummings, "Systems Engineering", lectures Fall 2015, Durham, NC.
[18] Kossiakoff & Sweet at Chapter 1.
[19] *Id.*
[20] *Id.*
[21] *Id.*

A systems engineer also serves as an interface between different engineering disciplines.[22] Components of a system developed by separate teams of engineers must eventually be integrated, and thus requires coordination of how the components will work together.[23] Further, the traits of a final project emphasized by different engineering disciplines is sure to vary, with individual specialists wanting the overall system to be designed to optimize their area of expertise.[24] As an example, Figure XXX shows how different specialists might view the best design for a missile.[25] When technical ideals for competing systems traits cannot be simultaneously implemented, systems engineers play the role of balancing the concerns of specialist engineers.[26]

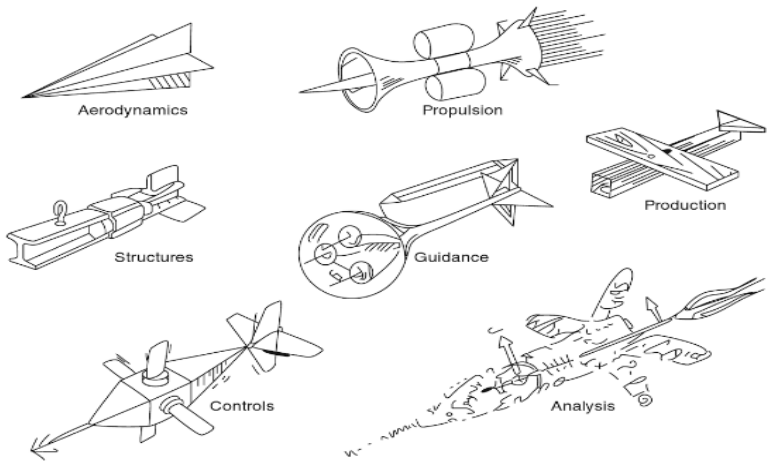

Figure 1: Ideal missile design as seen by different engineering teams.[27]

---

[22] *Id.*
[23] *Id.*
[24] *Id.*
[25] *Id.*
[26] *Id.*
[27] *Id.* at 18.

Together, and most importantly, the benefits of systems-level thinking make systems engineering the key risk management tool in the development of complex systems.[28] Systems engineers must make judgments about risk-risk tradeoffs, balancing competing interests between stakeholders, technical components, and engineering specialties.[29] Relevant risks include both business issues (e.g., budgets, customer needs, regulatory compliance, and competitive pressure) and technical problems (e.g., component failures and user errors).[30] Safety and reliability are of particular concern, with the technical risks balanced against competitive pressure to produce cutting-edge systems.[31] As one textbook puts it, "Selecting the most promising technical approaches, assessing the associated risks, rejecting those for which the risks outweigh the potential payoff, planning critical experiments, and deciding on potential fallbacks are all primary responsibilities of systems engineering."[32] Systems engineering plays a critical role in the development of safe systems.

Systems engineering is thus most useful when safety and reliability risks of emerging technology are most prevalent. According to one textbook:

> The characteristics of a system whose development, test, and application require the practice of systems engineering are that the system:

---

[28] *Id.* at 7-8, Chapter 8.
[29] *Id.* at Chapter 1.
[30] *Id.* at 7
[31] *Id.*
[32] *Id.* at 14

> (1) Is an engineered product and hence satisfies a specified need;
>
> (2) Consists of diverse components that have intricate relationships with one another and hence is multi-disciplinary and relatively complex;
>
> (3) Uses advanced technology in ways that are central to the performance of its primary functions and hence involves development risk and often relatively high cost.[33]

In other words, systems engineering helps most when a man-made product comprising multiple sub-parts relies on a new technology for which the possibilities and limitations are uncertain. Most autonomous systems are exactly this kind of product—they incorporate hardware elements ranging from well-known mechanisms (e.g, wheels and axels on self-driving cars) to innovative electronics (e.g., LIDAR systems and advanced sensors) to revolutionary software packages (e.g., self-learning algorithms for image recognition and navigation). In general, then, the risks associated with autonomous systems may best be managed through a systems engineering approach.

In order to explore how such a systems engineering approach works in practice, the next section details one popular framework known as the V-Model. Many versions of the V-Model exist, as do competing descriptions and formulations of systems engineering practices. But, for the purposes of eliciting a basic understanding of systems engineering's basic procedures for risk management, the story below is

---

[33] Kossiakoff & Sweet at 11

sufficiently generalizable and focuses on the concepts relevant to other chapters' discussions of law and regulation. The following does not intend to be a perfect description of how any given autonomous system is currently being developed, but rather provides a basic outline of risk management through proper systems engineering.

2.2 Frameworks for systems engineering: The V-Model[34]

The V-model is a visual representation of the main steps of systems engineering, sometimes referred to as the life-cycle building blocks.[35] On the horizontal axis is time, indicating that the activities in each block occur both in a certain order, but also overlap in time. This implies that some back and forth between steps may be necessary, and that some sub-projects may proceed to the next block while others stay behind. The vertical distribution indicates the level of system decomposition – that is, the blocks lower in the diagram deal with system components in more specific detail, while higher blocks look at the system from with a broader perspective. Blocks at the same level deal correspond to one another, as will be made clear in the descriptions that follow.

---

[34] The following discussion aggregates three practical sources for systems engineers, the MITRE Systems Engineering Guidebook, put out by a defense contractor to instruct their employees and subcontractors on best practices, a textbook on Systems Engineering by Kossiakoff and Sweet based on a practical course at John Hopkins University, and a course taught by Mary Cummings at Duke University titled "Introduction to Systems Engineering," Fall 2015. Where a statement in this subchapter is not individually footnoted, it comes from a generalized aggregation of these resources.

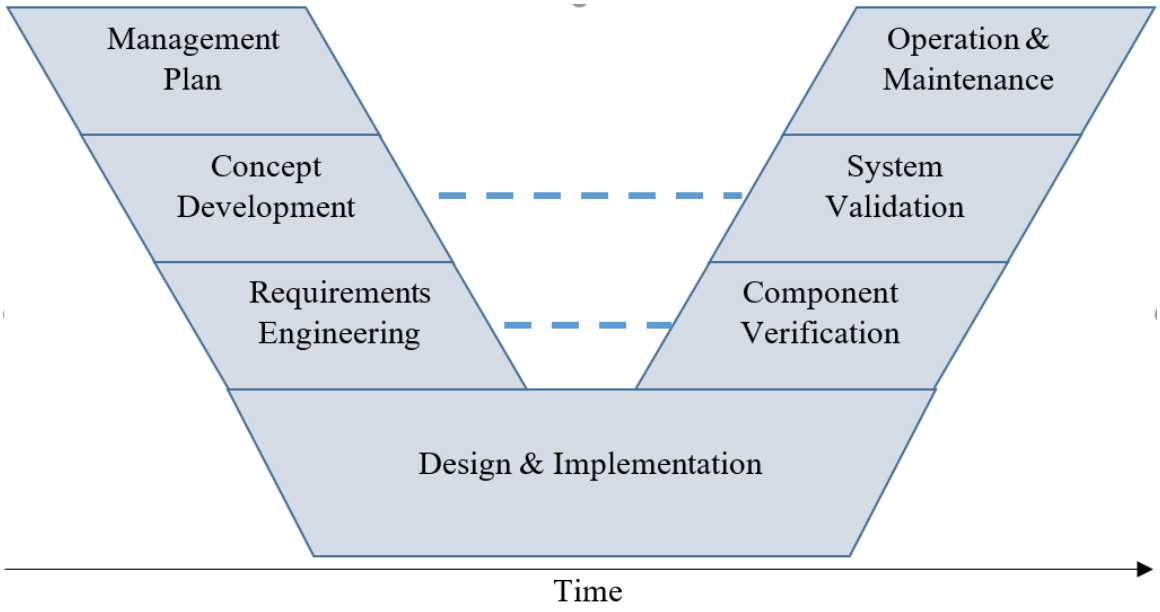

Figure 2: The V-Model[36]

2.2.1 Management Plan

In the management plan block, the system engineer takes an initial product idea or identified problem and plans how the project will proceed. Planning includes identifying initial budgets, project timelines, and necessary personnel. Other planning activities involve drafting a rigorous plan for stepping through the other blocks: these details will be spelled out below. The management plan provides the framework for disciplined adherence to all the other concepts laid out in this section. The plan must be

[36] Adapted from multiple sources by the author. *See generally*, Google Image results for "Systems Engineering V-Model", https://www.google.com/search?q=systems+engineering+v+model&espv=2&biw=1920&bih=950&source=lnms&tbm=isch&sa=X&ved=0ahUKEwi2rc6TncXSAhVLDsAKHWh6BU8Q_AUIBigB.

followed to rigorously to convey the value of tight connections between the life-cycle blocks that are key to the success of a system development project.

2.2.2 Concept Development

The Concept Development block requires systems engineers to develop a working description of what the system will do, what capabilities it will have, and how it will fit into its expected environment. This process takes two main phases. First, systems engineers must identify the end users targeted by the system and other stakeholders with an interest in the system's capabilities and interactions. Through interviews, observations, simulations, or other assessments, systems engineers can identify what the capabilities the system will need to achieve, as well as get a better understanding of the limitations or possibilities offered by the real-world situation in which the system will be used. Known as a needs assessment, this procedure compiles a list of user and stakeholder expectations that a successful system must meet.

Second, system engineers develop a "concept of operations." The concept of operations is a description of "the proposed system in terms of the user needs it will fulfill, its relationship to existing systems or procedures, and the ways it will be used."[37] A concept of operations defines who will use a system, what it will do, where it will do those things, and, to the extent relevant to an end user or important stakeholder, how it

[37] MITRE SE Guidebook at 285, quoting an Office of Management and Budget document available at http://www.abelia.com/498pdf/OCD-DID.PDF.

will accomplish those tasks. This is a working document, meant to be evaluated by both stakeholders and engineers to assess whether it meets the identified needs while being realistically achievable. Development and evaluation of a concept of operations helps to identify risks from multiple perspectives and note their importance early in the product life-cycle, as well as elicit an understanding of stakeholders' risk tolerance.

The concept of operations thus provides a description of the system to be met by more detailed requirements definition and design, as discussed below. In other blocks of the V-Model, certain activities must be traced back to the concept of operations to ensure that development is always aimed towards fulfilling user and stakeholder needs.

2.2.3 Requirements Engineering

The Requirements Engineering block consists of establishing and managing "requirements." As described in one guidebook,

> A requirement is a singular documented need—what a particular product or service should be or how it should perform. It is a statement that identifies a necessary attribute, capability, characteristic, or quality of a system in order for it to have value and utility to a user.[38]

Good requirements are unambiguous, measurable, testable, and traceable statements that form the intermediary step between the operational needs embodied in the concept of operations and detailed system design. They define the criteria that must

[38] MITRE SE Guidebook at 301

be met by designs and implementation in order to fulfill the needs identified in the Concept Development block, and are directly tested on the right side of the V-Model.

Requirements engineering involves several activities. First, systems engineers must elicit and collect requirements from users and stakeholders in an extension of the needs assessment activities of the Concept Development block. Elicited requirements should be prioritized or rated on necessity—for example, by labelling whether the need refers to a must-have, a want-to-have, or a nice-to-have feature. Second, requirements must be drafted and analyzed. Good requirement drafting offers clear direction to engineers without specifying design, does not introduce conflicting or unfeasible requirements, and makes obvious the criteria for successfully meeting each requirement. Requirements drafting is iterative, with user input and further analysis driving adjustments as needed. Requirements can quantitively capture stakeholder tolerances for failures, and thus include an understanding of the risks carried by the finish product.[39]

Third, requirements engineering recognizes that evolving user needs or technical discoveries can lead to changes in requirements after they are initially defined. Requirements engineering adds value by defining clear objectives for design, implementation, and verification, but also must recognize the risk that these

---

[39] For example, "The system shall have a mean time between failures of greater than 500 hours." MTIRE Guidebook at 355.

requirements will be forced to change later as a result of changing customer or stakeholder demands. A late change in requirements will generally lead to delays and sunk costs, as time and money was already spent on work that may now be useless. System engineers must try to minimize the risk of such delays, or at least identify areas where requirements are less stable so that changes can be anticipated.

2.2.4 Design & Implementation

In the Design & Implementation building block, employees work to turn requirements into a finished product. Systems engineers carry out several tasks during these phases.

First, systems engineers help to determine the system's architecture, a term that refers to the partitioning of operational capabilities and requirements into different components that can be developed by separate teams. This is particularly important if component development will be contracted out to other firms. Later, systems engineers will play a role in integrating these components back into a unified system. Second, systems engineers work with designers and engineers to determine which requirements are readily achieved and which present more difficult problems. Challenging technical problems can then be focused on from the management level or prioritized in terms of testing or funding, helping to reduce associated risks. Third, systems engineers work to trace design decisions back to the relevant requirements, in an effort to assure that the

end product matches stakeholder needs. Traceability to requirements is the key systems engineering concept at play in the design and implementation phase.

Technical risks must be mitigated through careful design, development, and implementation. Systems engineers must maintain awareness of high-risk components and focus attention on assuring safety, reliability, or redundancy where needed. As design and development reaches the stage where early solutions or prototypes can be tested, systems engineers oversee experimentation and component tests. At the end of the design and implementation block, a completed system is formed. The story has crossed the bottom of the V-Model: later steps will focus on testing and evaluating the product to judge its readiness for deployment.

2.2.5 Component Verification

In the V-Model, the Component Verification block is at the same vertical level as the Requirements Engineering Block. This represents that component verification consists of ensuring that the implemented components meet each relevant requirement. In other words, systems engineers use this block to check whether each component does what it is supposed to do, as defined in Requirements Engineering. Some requirements may be easily or instantly verified, like a size or weight limit. Other kinds of requirements require more involved testing, but should have been drafted in the Requirements Engineering stage with verifiability in mind. Systems engineers are responsible establishing testing procedures for these requirements and certifying test

results, which may prove difficult for novel technologies like autonomous software.[40] Risks of system failure are managed at this stage by establishing that each component works as intended, and works as well as intended.

2.2.6 System Validation

In the V-Model, the System Validation block is at the same vertical level as the Concept Development block. This represents that system validation consists of ensuring that the full system—made up of integrated components—meets the user's needs as outlined in the concept of operations. Like component verification, system validation requires planning of testing procedures and execution of those tests. One hopes that user needs are concrete enough to be able to demonstrate satisfaction through testing.

Validation tests often involve real target users in the real operational environment, at least at the final stage of testing. Systems engineers should test all possible scenarios faced by the system, including environmental influences, varieties of user input, and all the ways that system components might fail. For deterministic systems—those that react in the same way to the same situation every time—exhaustive testing may plausibly address all scenarios, particularly where outside influences are limited. But for autonomous systems—which may adjust their actions based on machine learning techniques, and make probabilistic decisions in response to outside

---

[40] *E.g.*, Missy Cummings, "The Brave New World of Driverless Cars", TR News (in press), available at http://hal.pratt.duke.edu/publications.

events—proving that the system will act as needed in all cases is much harder. Sufficient testing protocols remain the target of academic research.[41]

User and stakeholder risk aversion—that is, how low the probability of system failure must be to meet the user's needs—can play a role in the stringency of the required tests. Proof of a zero risk technology may be out of reach, due to technical limitations in the product, risk-risk tradeoffs, or inadequate testing procedures. System validation does not necessarily require perfection—it merely ensures that the earlier blocks of the V-Model lead to a system that meets stakeholder expectations for error rates and associated adverse consequences. System validation evaluates how well the risks of a system have been managed, and assesses remaining issues.

2.2.7 Operation & Maintenance

The Operation & Maintenance block refers to deploying the system to real-world users. Systems engineering activities include continued monitoring to ensure that production models comply with requirements and user needs, eliciting feedback from users, and guiding updates or alterations. Unforeseen or underestimated risks which become prominent with real use can be identified and addressed. Lessons from a product may feed back into the beginning of the cycle for development of the next

---

[41] *E.g.*, *Id.*; Andrews, Abdelgawad, and Gario, "World Model Testing Autonomous Systems Using Petri Net", 2016 IEEE 17th International Symposium on High Assurance Systems Engineering, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7423135.

generation of the technology. The Operation & Maintenance block represents the continued need for a systems engineering perspective even after product reaches market.

2.3 Other Notes on the V-Model

The V-Model shows the basic steps that engineering firms use to manage risk. In practice, systems engineering is not as linear as the V-Model approach suggests, but the model does provide a useful basic understanding of key activities. Other systems engineering frameworks like the "spiral model" rearrange the same steps and core activities into a more flexible or iterative model. Software companies, and Silicon Valley start-ups more generally, are known to prefer a more 'agile' approach. In some industries related to autonomous technology, this creates tension between more traditional V-Model-like development by some players while collaborators and competitors growing out of the software world are accustomed to a less rigorous framework. In general, however, the systems engineering activities outlined above are key features of engineering risk management for any complex technology.

This chapter provide a basic understanding of systems engineering, providing the basis for understanding how a policy intervention fits into the product lifecycle or imagining ways to leverage systems engineering practices to achieve policy objectives. The following chapter takes this analysis a step farther, listing and discussing the

regulatory interventions possible at each V-Model block for the regulation of autonomous systems.

3. Regulatory Intervention Throughout the V-Model

Regulatory design for autonomous systems will require picking and choosing from a long list of policy options. Building on the systems engineering model described in the preceding chapter, this chapter provides a guide for regulatory intervention at each stage of the V-Model, noting how policy options affect systems engineering practices and vice versa. Each V-Model block is stepped through deliberatively to list the relevant policy options to that block and briefly describe the pros and cons of those options. Although the policy options described here are often framed as government interventions, other regulatory modes like private standards organizations, market forces, social norms, or architectural effects are either part of or independently capable of implementing the policy options. By understanding and leveraging the relationship between system engineering and regulation, regulators may be better suited to formulate adaptive regulations that appropriately address the risks of autonomous systems while encouraging innovation.

3.1 Management Plan

The management plan stage consists of drafting a plan for project management to be rigorously followed throughout the other steps of the chosen systems engineering model. The main regulatory approach relevant at the management plan stage is to require companies to follow a systems engineering framework. Such requirements may include specific preordination of how the plan should look: in defense acquisition

programs, for example, contractors must follow fairly specific instructions dictating their approach to systems engineering, including a requirement to prepare a systems engineering plan.[1] Less restrictively, regulators could work directly with a company to develop a mutually acceptable systems engineering plan from the beginning, as done by the FAA.[2] At the other end of the spectrum, regulators might require some showing of a systems approach at a later stage in the life cycle, without insisting on a specific model or framework. As described in Chapter 4, the FDA takes this approach through its design controls.[3]

Requiring a systems engineering plan to be formulated (and followed) is the most direct way for a regulator to leverage systems engineering for risk management. A regulator taking this approach puts faith in the process of systems engineering, as opposed to directing attention to technology-specific risks. Regulations can thus be more widely applicable, dependent on consistent organizational compliance rather than close government involvement with the details of a new technology. Regulators may feel less pressure to draft specific rules for unpredictable emerging technologies, leaving room for innovation.

On the other hand, requiring a systems engineering framework may eliminate new business models and discourage industry newcomers. Placing regulatory

---

[1] Department of Defense Instruction 5000.02 http://www.navysbir.com/docs/500002p.pdf, 81 and following.
[2] *Infra,* Subchapter 4.1.
[3] *Infra,* Subchapter 4.2.

requirements before the concept development stage may benefit incumbent firms, who develop institutional knowledge of regulatory requirements and relationships with regulators, and can therefore comply with procedural requirements on the first try. New players may not realize they were supposed to have followed a specific systems engineering approach until it is too late. This affects both the little guys—for example, graduate students who never hear about FDA design controls until years into their labs' medical device research—as well as bigger companies entering from other sectors, like Silicon Valley software firms pushing AI into regulated industries. Requiring a systems engineering management plan may challenge the more free-wheeling spirit of many creative start-ups.

3.2 Concept Development

Concept development activities relate mostly to engaging with stakeholders and understanding how a system will fit into its external environment. As this phase focuses on non-technological influences, policy makers without technical training may be most comfortable intervening at this level. As detailed below, policymakers can influence concept development activities in several ways, including (i) altering the problem to be solved, (ii) giving voice to underrepresented stakeholders, (iii) determining traits of the system's end users, and (iv) shaping the business case for a new technology.

3.2.1 When Law Defines the Problem

The concept development stage requires a systems engineer to think about how the new system will fit into the existing world, and thus the laws and policies that shape the existing environment must be considered. This is particularly important for autonomous systems, many of which are aimed at directly undertaking existing human activities. Take, for instance, self-driving cars. These cars must follow traffic rules, a well-established body of law that defines how a car must operate in its environment. Existing traffic norms and transportation infrastructure plays a role in what a user needs out of a self-driving car, and therefore play a key role in concept development. Law—and other government influences like infrastructure design—that shape the tasks an autonomous system must undertake need to be considered while developing a concept of operations.

Policymakers, then, may see room for altering these existing laws or programs to adapt to new technologies. Reforms that simplify the technical problems or grant exceptions may speed deployment of new systems, but also cause compatibility issues between new systems and incumbent technologies. Large scale infrastructure alterations, such as erecting barriers around roadways to keep out pedestrians, could make autonomous systems easier to deploy by lowering the uncertainty in their environments, but come at a high cost, may alienate incumbents, and must anticipate future technological developments.

3.2.2 Giving Voice to Stakeholders

Regulation can increase the voices of certain people in the process of eliciting stakeholder needs in an effort to understand the impact of a new technology on a wider range of people. In other contexts, this approach is prominent in support of environmental protection, where the distributed harms of environmental degradation make negatively-affected individuals less likely to be offered a seat at the table compared to more unified interests. Within the U.S. government, for example, many decisions must be accompanied by an environmental impact statement, intended to show that decision makers took environmental concerns into consideration.[4] Private standards and market forces have also raised environmental concerns to be part of a stakeholder analysis to some extent, through the reputational advantage conveyed by indications like LEED certification.[5] Such policies may therefore be a model for increasing the prominence of certain stakeholder views in concept development.

One option for regulation of emerging technologies, then, is to require or encourage consideration of the concerns of people whose needs may otherwise be ignored but whose input aligns with policy goals. For autonomous systems, giving voice to ethicists, labor unions, or interest groups may encourage companies to consider

---

[4] National Environmental Policy Act, 42 U.S.C. § 4321 et seq. (1969); see "Summary of the National Environmental Policy Act" at https://www.epa.gov/laws-regulations/summary-national-environmental-policy-act.

[5] "Leed Ceritification", U.S. Green Building Council, http://www.usgbc.org/articles/about-leed

these groups concerns more seriously and influence the development of more ethical, socially motivated systems. However, legal requirements to merely talk to these parties as part of a stakeholder needs assessment, without more, may end up with a fake exercise and a waste of everyone's time if input from the supported voices is ignored in the rest of the system development process.

3.2.3 Defining End-Users

The concept development phase attempts to define the system from the perspective of a user, making it crucial to determine who those users will be in terms of their training, psychological and physical capabilities, and other traits relevant to the system. Licensing of users can be a key policy tool in this phase, as it can tell a system engineer what minimum characteristics its users will have. Policy interventions include the possibility to both lessen existing licensing schemes and create new ones. With regard to self-driving cars, for examples, realizing the benefits of increased mobility for people who are not currently eligible for a driver's license due to age or health issues may require an elimination of the licensing requirements for automobile operators.[6] As one example of a new licensing scheme, the FAA has created an online registration

[6] "California: self-driving cars will not need licensed driver, given federal approval", *The Guardian*, Oct. 1, 2016. https://www.theguardian.com/us-news/2016/oct/01/california-self-driving-cars-licensed-drivers

system for commercial users of small drones to try to educate and track users of the new technology.[7]

User licensing schemes—and related training requirements or testing of potential licensees—allow systems engineers to determine if their users will have some minimum level of skill or knowledge in the relevant activity. Restricting licenses to highly-trained or highly-insured people can help to mitigate the risks of a new technology, but also limits access to the technology for consumers (and thus access to markets for companies). From doctors to accountants to welders, most industries have some form of licensing that delineates who can do specific tasks or use certain technologies. By automating some of these tasks, emerging technologies will challenge the current formulation of many such user licenses, as users roles inside system operation can dramatically shift.

3.2.4 Shaping the Business Case for a New Technology

Every industry has different forces that shape what products come out of it, and law and regulation plays a role in shaping the direction of innovation whether intentional or not.[8] In the health care sector, for example, evidence suggests that increased medical malpractice liability increases a market for new tools that reduce the

---

[7] *See, e.g.,* FAA, "Getting Started", https://www.faa.gov/uas/getting_started/

[8] *See generally*, Johnathan Wiener, "The regulation of technology, and the technology of regulation", Technology in Society (2004), available at: http://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=1960&context=faculty_scholarship

risk for error, but might discourage truly disruptive medical devices that take entirely novel approaches to patient care.[9] Liability law for users or other entities can thereby impact demand for new technologies. More directly, government spending in the United States accounts for slightly over half of all health care expenditures in the United States,[10] and therefore government spending decisions can influence what problems medical device developers are incentivized to solve. In other sectors, policies accommodating or eliminating business models like ride-sharing (e.g., Uber) might have significant effects on which technologies are developed.

Thus, instead of regulating technologies as such, policymakers may prefer to guide innovation by shaping what the customers want. Government purchasing decisions, consumer nudges,[11] or other industry-dependent options may be able to help certain technological traits be economically successful. This is both somewhat less invasive as a regulatory measure, and less certain in terms of how likely it is to achieve a policy goal. From the systems engineer's perspective, however, law and policy that shapes user needs can play a significant role in concept development.

---

[9] Alberto Galasso and Hong Luo, "Tort Reform and Innovation" Harvard Business School Working Paper 16-093, 2016. http://www.hbs.edu/faculty/Publication%20Files/16-093_14c952bf-4842-4ed7-b785-f4b8ae39875b.pdf

[10] "NHE Fact Sheet", https://www.cms.gov/research-statistics-data-and-systems/statistics-trends-and-reports/nationalhealthexpenddata/nhe-fact-sheet.html (adding together Medicaid, Medicare, and state and local health care expenditures).

[11] Richard Thaler and Cass Sunstein, *Nudge: Improving Decisions about Health, Wealth, and Happiness*, 2008.

3.3 Requirements Engineering

The Requirements Engineering phase focuses on defining the requirements a system must meet in later stages of development, and is therefore a key phase for ensuring the system's legal compliance. Because of the strong links between requirements and the rest of the V-Model, the legal options discussed with respect to other V-Model blocks make their way indirectly into the requirements drafting process. And where regulatory intervention provides rules that must be implemented technologically, some regulations can enter the requirements drafting process more directly.[12] Systems engineers recognize the need to translate legal texts into useful requirements in order to ensure legal compliance of the final product.[13]

Examples of requirements-level legal rules include NHTSA's Federal Motor Vehicle Safety Standards, which provide testable technical requirements for automobiles,[14] and the FAA's airworthiness standards, which do the same for aircraft.[15] These regulations not only attempt to be specific enough to map well onto a systems engineer's requirements document, but sometimes incorporate documentation from private standards setting organizations by reference to provide more detailed

---

[12] Paul N. Otto and Annie I. Anton, "Addressing Legal Requirements in Requirements Engineering", *15th IEEE International Requiremetns Engineering Conference* (2007).
[13] *Id.*
[14] *Infra,* Section 4.3.1.
[15] *Infra*, Subchapter 4.1.

expectations.[16] In extracting requirements from this type of legal text and incorporated references, and from any other private standards documents that a company intends to follow, systems engineers are challenged to collate all relevant rules and cross-referenced materials from potentially diverse locations into a single set of project-applicable statements.[17] Where court decisions, agency guidance documents, or jurisdictional overlap come into play, understanding when to apply which rules becomes even more challenging. [18]

Additionally, while some regulations impose bright line rules that map onto good requirements, others are designed as vague standards that often use the word "reasonable." As one example of a technologically-implemented rule subject to such a standard, health privacy laws require companies to "protect against any reasonably anticipated threats or hazards to the security or integrity" of certain health information without defining the term "reasonably anticipated."[19] Lawyers introduce this sort of ambiguity to make room for case-specific flexibility and unanticipated circumstances, but vague standards do not lend themselves to translation into unambiguous, verifiable

---

[16] *See, e.g.,* 49 C.F.R. § 571.104 (An example FMVSS: "Daylight opening means the maximum unobstructed opening through the glazing surface, as defined in paragraph 2.3.12 of section E, Ground Vehicle Practice, SAE Aerospace-Automotive Drawing Standards, September 1963.")
[17] Otto, note 12.
[18] *Id.*
[19] *Id.*; HIPAA at §164.306(a)(2).

requirements in the systems engineering process.[20] Systems engineers may therefore have difficulty turning these rules into compliant products.

In sum, regulators might seek to impose technology-specific rules of the sort that can be easily mapped into a requirements document by systems engineers. This is most useful where all systems in a given field are similar enough to follow the same rules—cars, for example, all serve the same purpose in the same general way, and so legal rules with technical specificity can be broadly applied.[21] Systems engineers are relatively likely to incorporate such rules into a requirements document accurately. Regulators might prefer vague standards with wider applicability to diverse circumstances and evolving technology, but systems engineers may have trouble turning the associated ambiguities into requirements that will generate the kind of compliance intended by regulators. On the other hand, vague rules may leave room for creative development and design choices. This tension should be kept in mind while drafting regulations for new technologies.

3.4 Design & Implementation

The Design and Implementation phase is where the regulatory approaches from all other phases come together to influence the creation of a compliant product.

---

[20] *See* id. , Woody Hartzog, "Juris Machina: Legal Aspects of Robotics", presentation at We Robot 2016.
[21] As opposed to fields where each system might address a different problem in a new way (e.g., medical devices)

Importing regulatory influences from other phases, Design and Implementation is less explicitly tied to legal intervention but nevertheless may be influenced by policymakers. Legal doctrines potentially of interest to policy makers at this stage are intellectual property law and employment diversity.

3.4.1 Intellectual Property Rights

At some point in the design process, business decisions about intellectual property rights need to be made. These decisions include both ensuring that the system does not infringe other firm's patents, as well as assessing how best to protect one's own intellectual capital. Patents, trade secrets, and copyright all provide avenues for companies to maintain control over their own research and development results, granting limited monopolies to help incentivize invention and commercialization.[22] Policymakers may therefore be inclined to alter incentives for innovation related to intellectual property rights in guide the direction or nature of innovation.

For autonomous systems, relevant policy options include encouraging (or discouraging) use of patent pools and compulsory licenses, open-source software, or data sharing.[23] In general, less restrictive intellectual property rights can benefit innovation that builds upon earlier work and facilitate interoperability between products developed by different companies, while more powerful intellectual property

---

[22] *E.g.,* David Lange et al., "Intellectual Property: Cases and Materials", Fourth Edition 2012.
[23] Ryan Calo, "Open Robotics", Maryland L. Rev. 2011.

rights grant larger incentives for investment in a protected product and can help a company justify large commercialization costs. For autonomous systems which learn from large databases of information, the proprietary or shared nature of important datasets can be even more important than the scope of patent rights.[24] The availability of data and existing software can help guide design and implementation of a system in one direction, while exclusionary patent rights might keep it from going other ways. Thus, policymakers could leverage intellectual property rights to affect design decisions in autonomous systems.

3.4.2 Employment diversity

Anecdotal stories, at least, reflect a notion that the demographic makeup of a product development team in the high-tech space can imbue their systems with sociological biases. One classic example is the automated sink, with some models failing to dispense soap to people with dark skin.[25] In artificial intelligence, examples already abound: Google's photo app labelled black people as gorillas, Nikon's camera software thought Asian people were blinking, and search engines are less likely to show ads for highly paid jobs to women.[26] This has led many to suggest that diversity within the

---

[24] *See* Brenda Simon & Ted Sichelman, "Data-Generating Patents", 2017, http://scholarlycommons.law.northwestern.edu/cgi/viewcontent.cgi?article=1270&context=nulr.
[25] "The Reason this 'Racist Soap Dispenser' Doesn't Work on Black Skin", TechMic 2015, https://mic.com/articles/124899/the-reason-this-racist-soap-dispenser-doesn-t-work-on-black-skin#.BnMEQvyRI.
[26]Kate Crawford, "AI's White Guy Problem", https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html

teams of engineers designing, building, and training autonomous systems is necessary to help reduce biases within said systems.[27] As one writer put it, "Otherwise, we risk constructing machine intelligence that mirrors a narrow and privileged vision of society, with its old, familiar biases and stereotypes."[28] The extent to which the programming team can reduce discrimination in the way a computer system that by definition learns from real world data already full of existing biases is not so certain. But one regulatory option for policymakers will be to try to increase diversity Design and Implementation teams through employment regulations or educational support.

3.5 Component Verification

Component Verification consists of testing parts of the system to make sure they meet the original requirements. Thus, the verification stage is when requirements elicited from legal rules begin to be tested, and so the issues discussed under the Requirements Engineering heading return. As made clear by the Volkswagen emissions debacle, poorly-designed tests that give fake results or encourage cheating can be disastrous for a regulatory program.[29] Thus, the goal is not just a showing of legal compliance, but faith in the accuracy of the verification process. Addressing this goal at

---

[27] http://spectrum.ieee.org/tech-talk/at-work/tech-careers/computer-vision-leader-feifei-li-on-why-ai-needs-diversity

[28] Kate Crawford

[29] *E.g.,* "Volkswage: The Scandal Explained" , BBC News, 2015, http://www.bbc.com/news/business-34324772; "Beyond Volkswagen, Europe's Diesels Flunked a Pollution Test", NY Times, 2016. https://www.nytimes.com/2016/02/08/business/international/no-matter-the-brand-europes-diesels-flunked-a-pollution-test.html.

a slightly abstract level, the main regulatory design question is about how involved regulators need to be in defining the necessary tests, carrying out those tests, or reviewing the results of the tests.

The options span a spectrum, as regulators might: (1) design, run, and analyze the results of experiments independently; (2) work with a company to design mutually acceptable tests, oversee testing procedures, and collaborate in reviewing the results; (3) review testing procedures and results only after tests are completed; (4) expect companies to self-certify that testing procedures and results were sufficient without substantive review of results; (5) or leave companies free of obligations to communicate with regulators about verification activities. A combination of these schemes, or blurred lines between them, are options as well.

At the first end of the spectrum, independence of testers from companies can work to increase faith in the neutrality of the testing program.[30] Slightly down the spectrum, collaboration between industry and regulators helps to keep the people with the most knowledge about a system involved in its testing and helps companies better understand regulator's expectations, so that risks can be better identified and managed.[31] Further on, reducing the role of regulators allows companies to take an approach that makes more sense for them, whether for technical or business reasons. Leaving

---

[30] *Infra* Subchapter 4.1, discussing the FAA Designee program.
[31] *Id.*

companies to comply with laws without supervision at the verification stage, helps to reduce administrative burdens and can be appropriate when market forces or the threat of large penalties incentivize compliance. Different industries and different risks lend themselves to different schemes along this spectrum.

3.6 System Validation

System Validation consists of assuring that the final system matches user needs. As this is also a test and evaluation stage, the discussion from the Component Verification section above remains relevant. But because system validation requires evaluating how the system functions in its intended environment, another important regulatory question is presented: When and how can an as-yet-unproven system be deployed into the real world in order to generate the data needed to validate the system? Exposing people to uncertain risks will often be a necessary part of the validation process. The most obvious regulatory example is FDA drug trials, carefully formulated over decades to try to balance ethical concerns, opportunity costs for patients, and the need to generate data related to the safety and effectiveness of a new drug. Over 90% drugs that start the first in-human trial phase fail to reach approval[32]—after real people have been harmed—making clear that trial participants are often

[32] "Clinical Development Success Rates 2006-2015", Biotechnology Innovation Organization, https://www.bio.org/sites/default/files/Clinical%20Development%20Success%20Rates%202006-2015%20-%20BIO,%20Biomedtracker,%20Amplion%202016.pdf.

exposed to higher levels of danger than agencies are willing to allow for the general public. As a class of technology, pharmaceuticals are not generally lumped in with autonomous systems or systems engineering, but the real-world-testing-risk for autonomous systems is of a similar nature and needs to be carefully managed.

In some industries with autonomous systems, regulators might want to require a new system to be approved for testing before exposure to its operational environment. FDA's Investigation Device Exemption process, which requires a showing of successful bench and animal experiments before a device is cleared for experimentation in humans, is one such example.[33] This may slow down innovation but seems to have few systemic consequences corrupting the validation process.

For other industries, regulators may choose to set aside geographical areas in which testing can take place in an attempt to at least localize the risks. The FAA, for example, has set up unmanned aircraft system test sites where researchers can experiment with drone technology within a designated airspace.[34] This is certainly useful for prototype or component testing. But if the segmented area does not contain all of the environmental factors a system will have to deal with in its full deployment, testing there cannot fully validate an autonomous system.

---

[33] *Infra,* Section 4.2.1.

[34] https://www.faa.gov/news/fact_sheets/news_story.cfm?newsId=15575

Another option is to require close supervision of autonomous systems during testing, with humans able to intervene to stop negative incidents. Some state laws for autonomous vehicles, for example, require the presence of a trained company representative in the cars during real-world testing, with the ability to override the system to prevent crashes.[35] Since one goal is to determine the statistical rates of crash incidents, the way these representatives meddle with what the car does can significantly affect the data. One Google-funded study out of Virginia Tech, for example, praised a great decline in the number of crashes per mile driven for the Google Car but did not appear to take into account the number of incidents prevented by the human users.[36] When noted, the number of human-prevented accidents may also be suspect, because of the difficulty in determining whether the human intervention was truly necessary to prevent a collision. This regulatory option thus makes analysis of test results difficult, but may be preferable compared to skipping a limited testing phase altogether or only testing in a sanitized testing environment.

The situations in which a system can be tested are thus the first regulatory decision: the second decision is determining when the test data is sufficient to prove that the system meets stakeholder or legal needs. For autonomous systems which interact

---

[35] For ongoing updates on state laws related to autonomous vehicle systems, monitor the National Conference of State Legislatures "Autonomous Vehicles Legislative Database", http://www.ncsl.org/research/transportation/autonomous-vehicles-legislative-database.aspx.
[36] Blanco et al., "Automated Vehicle Crash Rate Comparison Using Naturalistic Data", Virginia Tech Transportation Institute, Jan. 2016.

probabilistically with the physical world, the complexity of the systems and the uncertainty of the world around them makes it difficult or impossible to test every possible scenario in which a system interacts with its environment.[37] This problem has not yet been solved, and creation of testing protocols for autonomous systems is currently an active area of research.[38] One policy option is to fund testing research for autonomous systems in order to generate the tools regulators need.

At a more general level, stakeholder risk tolerance shapes how reliable a system must be shown to be, and therefore helps determine the endpoints for experimentation. Regulators—and businesses—will be forced to make difficult decisions about autonomous systems related to maximum error rates, especially for technologies intended to replace existing tasks ridden with human errors. Deciding how safe is safe enough is a technical challenge for experts running experimental tests, businessmen judging market tolerance for danger, and policymakers weighing the costs and benefits of stricter standards. Regulators looking closely at a given industry would do well to understand how those perspectives overlap in order to best craft policies around System Validation.

---

[37] *E.g.*, Andrews, Abdelgawad, and Gario, "World Model Testing Autonomous Systems Using Petri Net", 2016 IEEE 17th International Symposium on High Assurance Systems Engineering, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7423135

[38] *See* id.; "Adaptive Testing of Autonomous Systems," U.S. Naval Research Laboratory, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7423135; "GTRI to Develop Technology Roadmap for Test and Evaluation of Unmanned and Autonomous Systems", Georgia Tech Research Institute, https://gtri.gatech.edu/casestudy/autonomous-systems-roadmap. Many other programs are ongoing as well.

3.7 Operation & Maintenance

The previous stages offer opportunities for *ex ante* regulation; the Operation & Maintenance stage offers opportunities for *ex poste* legal intervention. In other words, opportunities abound for law and regulation to continue to influence autonomous systems after the systems start to reach end users. Relevant legal instruments in Operation & Maintenance include: (1) manufacturing controls; (2) postmarket surveillance and reporting requirements; (3) defect investigations and product recalls; and (4) civil liability.

3.7.1 Manufacturing controls

Regulating manufacturing helps to ensure that the products that reach end users are sufficiently identical to those which were approved by regulators during earlier stages in the V-Model. FDA and FAA both oversee manufacturing to some degree as part of their schemes to insure safety of the products that reach the market.[39] Possible modes of regulatory oversight include inspections of manufacturing facilities, evaluation and certification of finished products, or more passive measures like education of best practices guidelines. Certainly, regulatory approval processes in Component Verification or Systems Validation are rather pointless if no attempt is made to check if the actual, marketed products match what regulators approved. On the other hand,

[39] *Infra,* Section 4.2.2 and 4.1.1.

strict regulatory oversight of manufacturing imposes administrative costs and could limit business's flexibility in subcontracting or outsourcing certain manufacturing tasks.

3.7.2 Postmarket Surveillance and Reporting Requirements

Regulators may wish to keep evaluating systems after they reach end users. Goals here are to get more data to feed back into the System Validation activities, identify unforeseen failure modes, and generally track the effectiveness of premarket evaluations. One regulatory option is deploying regulators to directly observe or test systems in real-world deployment, as NHTSA does by purchasing cars from regular car dealerships and test them for compliance with safety standards.[40] Another option is to compel reporting from industry or users, generally focused on adverse incidents. FDA, for example, requires manufacturers and user facilities (i.e., hospitals, nursing homes) to file reports when bad things happen related to the use of their devices, including product quality issues, use errors, and therapeutic failures.[41] A more passive option would be to set up a process for voluntary incident reporting.

Postmarket surveillance has the benefit of providing data about the system in its actual use, which may not match exactly the anticipated use or testing conditions relied upon for premarket regulation. Also helpful is that the sample size available after a system reaches full-scale production and deployment will generally be much higher

---

[40] *Infra* Section 4.3.1

[41] *Infra* Subchapter 4.2; https://www.fda.gov/MedicalDevices/Safety/ReportaProblem/ucm2005291.htm

than was possible in the testing phases. Additionally, allowing some regulatory approval activities to wait until after full deployment allows innovations to reach consumers faster—an important regulatory goal in the health sector.

These advantages may be countered by the difficultly of aggregating full, accurate information across many users in the market. Outside of the control of an experimental protocol, reports about adverse incidents may be quite difficult to analyze in any rigorous, quantitative way, especially depending on the quality of information contained in those reports.[42] Regulators are likely to have less time and funding for postmarket evaluation, and may be disinclined to review data that might show them they were wrong to approve a system for market in the first place. Policymakers should be careful to design postmarket surveillance schemes so that they generate useful data that someone actually analyzes.

3.7.3 Defect Investigations and Corrective Remedies

Related to postmarket surveillance is the question of what to do once that surveillance discovers a potential issue. Regulators can be granted authority to conduct investigations of potentially unacceptable risks. This authority might include the right to require a company to submit records, allow inspections, or make employees appear for

[42] For example, take a look through some records in the FDA reporting database, and it quickly becomes clear that the data would not be easily analyzed, thanks to the free-form text entry it relies upon. https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfmaude/search.cfm

interviews. Regulators can use those powers to make particular findings about whether a system poses risks that must be mitigated.

Once such a risk is found, regulators then must decide how to deal with it. Recalling the products from the market is one option, and can place significant financial strain on a company if forced to buy the product back from its customers or replace the products with improved models. Repairing defective features in a product is also a possible remedy to the risk. Where the risk is embodied in hardware, the costs of repairmen and parts add up quickly. Software issues, on the other hand, may be solved through a relatively-cheap online update.[43] The incentives for companies to avoid software defects may therefore be smaller than those to avoid hardware problems, an important consideration as software becomes an integral part of the safety of autonomous systems. A final option is to impose fines, or restrict other company activity until a risk is mitigated.

3.7.4 Civil Liability

Lawsuits against manufacturers, designers, and users demanding payment of damages for harms caused by autonomous systems will certainly shape how these systems develop. Civil liability has the advantage of compensating victims for actual harm, and placing the economic costs of those harms directly on the responsible parties.

---

[43] Tesla took advantage of this- See *Infra* Subchapter 4.3.3

Civil liability can therefore deter the manufacture and sale of unsafe systems, or, at least ensure that companies take into account a projected cost of settling lawsuits brought by the people they harm. Because many low-income Americans do not have the economic means to get access to or proper legal representation the civil justice system,[44] civil liability may not capture the full cost of the harms inflicted by a system.

The availability of civil liability for product defects also puts state court judges or lay juries in the position of deciding whether an autonomous system was reasonably safe. In contrast, expert regulators can put products through careful evaluation and make cost/benefit decisions with potentially a lot more expertise and broader knowledge than the issues considered in a single courtroom. Contradictory court decisions across jurisdictions can also lead to difficult compliance problems for companies. One legal option, then, is federal regulatory preemption of state tort law. Medical devices which go through the full premarket approval process, for example, are immune from tort liability so long as the device in question complied with all applicable federal regulations.[45]

Legal interventions could also determine how liability is attributed in cases involving new technologies. In the Robot Law literature, much ink has been spilled

---

[44] "Access to Justice in the United States", The World Justice Project 2010, http://worldjusticeproject.org/sites/default/files/486481access_to_justice_in_the_united_states_virginia_lawyer_12-10.pdf

[45] *Riegel v. Medtronic*, 553 U.S. 312 (2008), interpreting the preemption provision of the Medical Device Amendments of 1976, 21 U.S.C. § 360k(a).

trying to sort out whether common law tort doctrine can properly address the new ways that autonomous systems might harm people as a matter of substantive law.[46] Some of this literature concludes that preordained forms of compensation for injured people will be necessary due to difficulties in assigning liability to anyone.[47] A related issue is how markets for insurance tie into liability related to autonomous systems: compelling insurance for certain players, as is common today for car ownership, may be a policy option in some industries.[48] Altogether, the way that law distributes the costs of harms caused by autonomous systems after they occur is sure to affect incentives for innovation up front.

3.8 Regulatory Intervention Conclusion

This chapter provided a list of regulatory options for autonomous systems, with reference to the V-Model. This provides a guide for policymakers who wish to keep the practices of systems engineering in mind when crafting regulatory oversight for new engineered technologies. The analysis does not—and cannot—lead to the conclusion that one specific combination of the options outlined in this chapter creates the correct

---

[46] Among many examples: Peter Asaro, "Robots and Responsibility from A Legal Perspective", IEEE, http://www.peterasaro.org/writing/asaro%20legal%20perspective.pdf; Matthew Scherer, "Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies", Harv. J. Law & Tech., 2016.; Ryan Calo, "Open Robotics", Maryland L. Rev. 2011.

[47] *E.g.,* Orly Ravid, "Don't sue me, I was just lawfully texting & drunk when my autonomous car crashed into you," 44 Sw. L. Rev. 175, 2014.

[48] "Self-Driving Cars and Insurance", Insurance Information Institute, http://www.iii.org/issue-update/self-driving-cars-and-insurance

regulatory scheme for a given technology. Rather, no one right answer can capture all industries, autonomous systems, or risks. In order to see how regulation and systems engineering already interact in U.S. federal regulation in a variety of ways, the following chapter explores three federal regulatory schemes sure to be relevant for some autonomous systems.

4. Existing Regulations and Systems Engineering

This chapter lays out the regulatory schemes for ensuring product safety in industries likely to house applications for autonomous systems. The discussion includes both long-established regulatory regimes and any recent policies which attempt to address new challenges of autonomous technology. More particularly, the chapter gives an overview of the federal government's regulation of aircraft safety under the Federal Aviation Administration (FAA), medical devices under the Food and Drug Administration (FDA), and automotive parts and vehicles under the National Highway Traffic Safety Administration (NHTSA). These examples reveal a range of options with respect to how deeply government regulators are involved in regulated firms' engineering management, from ongoing relationships based on control of the systems engineering framework to mere investigation of compliance after a product reaches the market.

4.1 Federal Aviation Administration and Systems Engineering

Airplanes, helicopters, and other aircraft are multifaceted technological systems which hold safety as a primary goal. The aerospace industry has long pushed the boundaries of human ingenuity, and has already incorporated significant automation into flight systems. Because aircraft are likely to continue this trend towards autonomous flight, review of FAA regulation of aircraft safety provides a window into

how one agency deals with safety-critical automation and related complex systems. As discussed in the following, the FAA is involved at all stages of a product's lifecycle.

4.1.1 Aircraft Certification

The FAA's "continuing mission is to provide the safest, most efficient aerospace system in the world."[1] To achieve one element of this mission, aircraft safety, the FAA starts its relationship an aerospace company before any particular innovation is conceptualized.[2] A document called the Partnership for Safety Plan defines the working relationship between a company and the FAA independent of any specific compliance project, in order "to build mutual trust, leadership, teamwork, and efficient business practices."[3] One goal of established communication is to reduce regulatory delays by identifying requirements for certification and special issues early in the product development process.[4]

A key aspect of FAA's ongoing relationships with companies is the Designees and Delegations system. Under statutory authority dating to at least 1958, FAA delegates responsibility to employees of regulated companies to perform certification tasks on behalf of the FAA.[5] These employees, known as Designees, provide efficient

---

[1] https://www.faa.gov/about/mission/
[2] "The FAA and Industry Guide to Product Certification", 2nd Edition, Sept. 2004, p. A1-30. In practice, this typically means FAA maintains relationships with established aerospace firms outside of particular projects.
[3] *Id.*
[4] *Id.* at 1.
[5] The FAA Act of 1958, Section 314; Regulations at 14 C.F.R. § 183.

collaboration between companies and FAA functions.[6] FAA officials are freed up to focus on novel or higher-risk safety issues.[7] While the legitimacy of the designee system is often questioned by the public when something goes wrong,[8] oversight safeguards attempt to ensure the integrity of the Designee system and create a working environment for Designees free of undue pressure from their employers.[9] Designees—roughly 3000 in number[10]—are generally established before any particular product is conceptualized,[11] a clear sign of FAA's reliance on sustained relationships with industry.

Once an aerospace company gets an idea for a new a new aircraft, aircraft engine, or aircraft propeller, FAA's aircraft certification process begins. The most onerous regulatory certification is "type certification", and refers to the approval process for the design of an entire aircraft, including subcomponents like jet engines and avionics.[12] Type certification consists of five phases: conceptual design, requirements definition, compliance planning, implementation, and post certification.[13] In other

---

[6] "The FAA and Industry Guide to Product Certification", 2nd Edition, Sept. 2004, p. A1-31.
[7] *Id.*
[8] *See, e.g.,* "Insight: Will Dreamliner drama affect industry self-inspection?" *Chicago Tribune*, 3/2/2013 http://articles.chicagotribune.com/2013-03-02/business/sns-rt-us-boeing-787-oversightbre92104w-20130302_1_faa-inspector-faa-administrator-michael-huerta-dreamliner
[9] "The FAA and Industry Guide to Product Certification" at A1-31–2.
[10] "Insight: Will Dreamliner drama affect industry self-inspection?" *Chicago Tribune*, 3/2/2013
[11] "The FAA and Industry Guide to Product Certification"at A1-31.
[12] Geoffrey M. Hand, "COMMENTS: Should Juries Decide Aircraft Design? Cleveland v. Piper Aircraft Corp. and Federal Preemption of State Tort Law", 29 U.S.F. L. Rev. 741, 754-756
[13] "The FAA and Industry Guide to Product Certification" at 5.

words, type certification utilizes a systems engineering framework very much like the V-Model laid out in Chapter 2.

The FAA's "Conceptual Design" phase consists of "familiarization meetings" between FAA and the company to discuss the new product idea and its concept of operations. The goal is to identify critical areas and difficult certification issues at the outset and to develop "a mutual understanding of potential new projects."[14] A "Project Specific Certification Plan" begins to be formulated during this phase, which includes a project timeline, checklists for moving on to the next phases, means of compliance, testing plans, and other project management information.[15]

The "Requirements Definition" stage consists of clarifying the product definition and identifying specific regulations and methods of compliance.[16] Specific regulations known as airworthiness standards set product requirements with significant detail.[17] The airworthiness standards address all safety-related aspects of the typical aircraft, from weight limits[18] and takeoff speeds[19] to fabrication methods,[20] passenger information

---

[14] "The FAA and Industry Guide to Product Certification", 2nd Edition, Sept. 2004, p. 7.
[15] *Id.* at A2-38.
[16] *Id.* at 8.
[17] Codified at 14 C.F.R. subchapter C; One relatively simple example: 14 C.F.R. § 23.783(f)(1) ("Each passenger entry door must qualify as a floor level emergency exit. This exit must have a rectangular opening of not less than 24 inches wide by 48 inches high, with corner radii not greater than one-third the width of the exit.").
[18] 14 C.F.R. § 23.25
[19] 14 C.F.R. § 23.51
[20] 14 C.F.R. § 23.605

signs,[21] and fuel pumps.[22] If a new or unusual feature of the product is not adequately covered by existing airworthiness regulations, a "special condition" may be needed.[23] A special condition is a new regulation particular to that aircraft, typically developed through notice-and-comment rulemaking, which fills the gap in existing airworthiness standards.[24] Communication between a company and the FAA at the Requirements Definition stage allows for earlier processing of a perceived need for a special condition.

The "Compliance Planning" phase consists of completing the Project Specific Compliance Plan, including finalizing the airworthiness standards to be met, test plans to show compliance with those standards, project schedule, and particular Designee responsibilities.[25] Other activities include analysis of product failure modes and consultation with stakeholders.[26] Collaboration between companies and FAA in the Compliance Planning phase ensures that both sides will be on the same page with respect to how the product will be required to demonstrate its safety.[27]

In the "Implementation" phase, the company works closely with the FAA to "ensure that all agreed product specific certification requirements are met."[28] Mapping

---

[21] 14 C.F.R. § 23.791
[22] 14 C.F.R. § 23.991
[23] 14 C.F.R. § 11.19.
[24] 14 C.F.R. § 11.38 ("Even though the Administrative Procedure Act does not require notice and comment for rules of particular applicability, FAA does publish proposed special conditions for comment [with some exceptions].")
[25] "The FAA and Industry Guide to Product Certification", 2nd Edition, Sept. 2004, p.10.
[26] *Id.*
[27] *Id.*
[28] *Id.* at 10.

onto the Component Verification and System Validation stages of this thesis's version of the V-Model, activities here include completion of all test plans, flight tests, inspections, and other compliance documentation. Many of these tasks are completed under the authority delegated to the company's Designee to carry out compliance activities on behalf of the FAA.[29] While a company can run product development tests without FAA involvement, a test only counts towards certification if the FAA was informed and arrangements were agreed upon prior to testing.[30] FAA must confirm that such tests meet the mutually-developed tests plans from the Compliance Planning phase such that credit for successfully passing those tests can be given.[31] Conformity rules and inspections trace the custody of tested products and check design characteristics in an attempt to ensure that the tested product is identical to the design being evaluated.[32] When all compliance activities are complete, the type certification approval is issued.[33]

The final phase is Post Certification. The focus here is on continued airworthiness through maintenance and operational awareness. To be fully approved to fly, an aircraft also requires a "production certificate", which recognizes the

---

[29] *Id.* at A2-42, A6-94.
[30] *Id.*
[31] *Id.*
[32] *E.g., id.* at A2-43; 14 C.F.R. § 21.53 ("(a) Each applicant must provide, in a form and manner acceptable to the FAA, a statement that each aircraft engine or propeller presented for type certification conforms to its type design. (b) Each applicant must submit a statement of conformity to the FAA for each aircraft or part thereof presented to the FAA for tests. This statement of conformity must include a statement that the applicant has complied with §21.33(a)").
[33] "The FAA and Industry Guide to Product Certification", 2nd Edition, Sept. 2004, p.10.

conformance of the manufacturing process with the approved design, and an "airworthiness certificate" which signifies that an individual aircraft matches the approved design.[34] The inspections necessary for these certifications are also often carried out by company's Designees and overseen by the FAA.[35] By the time a new aircraft is fully approved for commercial operation, as much as 8 years may have passed since the beginning of the type certification process.[36]

After certification procedures, the FAA requires continued management, reporting, and self-disclosure of airworthiness issues which become apparent after the product is on the market.[37] When a safety risk is discovered, FAA can implement an "Airworthiness Directive" through notice-and-comment rulemaking.[38] An airworthiness directive makes it unlawful to operate an aircraft unless one carries out the inspections, complies with the limitations, or acts to resolve unsafe conditions as specified in the new regulation.[39] One recent airworthiness directive, for example, "requires an inspection for, and replacement of, all non-conforming aft engine mount

---

[34] Geoffrey M. Hand, "COMMENTS: Should Juries Decide Aircraft Design? Cleveland v. Piper Aircraft Corp. and Federal Preemption of State Tort Law", 29 U.S.F. L. Rev. 741, 754-756
[35] *See, e.g.*, "That's the Ticket! How an FAA Designee helps ensure that Boeing delivers on its promises." *Boeing Frontiers*, Oct. 2007, http://www.boeing.com/news/frontiers/archive/2007/october/i_ca01.pdf .
[36]Boeing 787 certification, http://787updates.newairplane.com/Certification-Process
[37] *Id.* at A2-45.
[38] 14 C.F.R. § 39.
[39] *Id.*

retainers" on a certain model of Airbus airplanes.[40] Airworthiness directives are quite common, with 56 issued in the last 60 days when this passage was drafted.[41]

Through these phases, the FAA stays involved in the life of an aerospace technology from conception to retirement, with agency personnel (or at least Designees) involved in every block of the V-Model. FAA's regulatory design leverages the process of systems engineering to manage risks, while also setting technology-specific rules to govern new technology at the level of requirements engineering and observing tests to prove compliance with those rules. FAA thus serves as an example of full embrace of systems engineering by regulators of complex engineered systems.

4.1.2 FAA Authority Beyond Product Regulation

Beyond evaluation of a product as such, the FAA also has authority to regulate pilots, passengers, maintenance crews, and cargo, as well as control over access to and coordination within the national airspace.[42] With respect to emerging technologies, this broad authority gives the FAA both more policy levers to pull and more issues to address. For example, the FAA responded—after years of delay—to consumer drones through policies targeting their users and locations of use, rather than attempting to

---

[40] Airworthiness Directive, Airbus Airplanes, Feb. 27, 2017. Available at http://rgl.faa.gov/Regulatory_and_Guidance_Library/rgad.nsf/0/9bf6e77316eaa83c862580d4006798ee/$FILE/2017-04-10.pdf

[41] https://www.faa.gov/regulations_policies/airworthiness_directives/ (checked Feb. 8, 2017).

[42] *See, e.g.*, https://www.faa.gov/about/safety_efficiency/

require design certification as outlined above for each new drone.[43] The FAA created registration websites for consumer drones, a process for licensing drone pilots, and clarified rules about flying near airports or sporting events, over crowds, and within the line of sight of the pilot.[44] A slightly different licensing scheme applies to commercial users of small drones as opposed to recreational users.[45] A small unmanned aircraft system no longer needs an airworthiness or type certification, so long as it is flown in accordance with the new rules for use.[46] Federal drone regulation has therefore taken a somewhat indirect approach to the new technology relative to safety evaluation of the products themselves.

In slowly granting small drone users access to the skies, the FAA exhibited its discomfort with ceding its oversight authority over any aspect of the national airspace. Instead, the FAA is more accustomed to involvement in all phases of a product's conception, development, testing, and use. FAA's traditional approach to aircraft design certification demonstrates a regulatory approach where the agency is present throughout of all steps of a systems engineering framework.

---

[43] *See, e..g*. https://www.faa.gov/uas/faqs/ Consumer drones as used here refers to unmanned aircraft systems weighing less than 55 pounds. Regulations at 14 C.F.R. part 107.
[44] *Id.*
[45] *Id.*
[46] *Id.*

4.2 Medical Device Regulation

Health and medical applications are another potential outlet for autonomous systems. Applications of autonomy in health care include surgical robots that complete surgical tasks without direct physician control, drug delivery systems that vary dosage rates in response to physiological measurements, and artificial intelligence-based diagnostic programs that take-in patient information and output treatment regimens. When such a technology falls within the statutory definition of a "medical device"—that is, it is an instrument, machine, or contrivance "intended for use in the diagnosis . . . cure, mitigation, treatment, or prevention of disease"—it also falls into the regulatory purview of the Food and Drug Administration.[47] Because the FDA has not given any indication that it will alter its existing frameworks when considering autonomous systems, this section describes the FDA's current regulatory scheme.

The following discussion reveals that the FDA does not become directly involved in the product lifecycle until the testing phase, although it expects medical device companies to have done some amount of systems engineering planning before it allows clinical trials to begin.

[47] 21 U.S.C. § 321(h); the Agency has discretion to choose not to regulate categories of products that fall in this definition. The Obama FDA generated some controversy by expanding the scope of what the FDA intended to regulated, but one might expect that trend to reverse under the new administration.

4.2.1 Design Controls and Investigation Device Exemptions

The FDA requires all medical device companies to follow a set of good manufacturing practices, a somewhat misleading label for a category of rules that also includes controls on how a device is developed.[48] Relevant to our discussion are the FDA's "design controls" which essentially mandate that a medical device company follow some version of a systems engineering framework.[49] Although the regulations do not specify a particular systems engineering framework that must be followed, an FDA guidance document lays out in detail suggested principles for design planning, input, output, review, verification, validation, and transfer, which are quite similar to the V-model activities laid out in Chapter 2.[50] Although the FDA does not directly monitor developer's engineering management early in the product life cycle like the FAA does, FDA's stance is that "Design controls increase the likelihood that the design transferred to production will translate into a device that is appropriate for its intended use."[51]

Communication between developers and the FDA begins when a medical device developer wants to start testing a new device on humans: in other words, just before the System Validation stage of the V-Model. Human trials cannot be conducted without

---

[48] 21 CFR § 820
[49] 21 CFR § 820.30 , guidance document at http://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm070627.htm
[50] *Id.*
[51] *Id.*

prior clearance from the FDA under an Investigational Device Exemption (IDE).[52] The process of acquiring an IDE requires meetings with the FDA and submission of significant information about the device. This information includes: device description, drawings, components, materials, principles of operation, and specifications; analysis of potential failure modes; proposed uses, patient populations, instructions, warnings, and training requirements; proposed plans for clinical evaluation criteria and testing endpoints; and summaries of bench or animal test data or prior clinical experience.[53] A lot of this information is related to concept development, requirements engineering, and other systems engineering activities much earlier in the product life cycle.

One result of the IDE procedure is that the applicant works with the FDA to determine the criteria, end-points, and objectives of clinical trials, when deemed necessary.[54] Because a device might attack a disease or condition in a novel way—that is, each device solves a slightly different problem—no standard metrics of success can necessarily apply to all new medical devices. Compulsory meetings before clinical trials thus serve to help determine what the goals of the testing will be.[55]

---

[52] 21 C.F.R. § 812.1

[53] Early Collaboration Meetings Under the FDA Modernization Act (FDAMA); Final Guidance for Industry and for CDRH Staff
http://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm073604.htm

[54] *Id.*

[55] Early Collaboration Meetings Under the FDA Modernization Act (FDAMA); Final Guidance for Industry and for CDRH Staff
http://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm073604.htm

4.2.2 FDA Validation Pathways

FDA then evaluates devices and clears them for marketing in the System Validation phase. Tasked to both protect the public health and to advance it through innovation,[56] the FDA's medical device evaluation programs attempt to ensure the safety and effectiveness of the systems in treating patients without placing overwhelming regulatory obstacles in the way of device developers.[57] Trying to strike this balance, the FDA currently has two main regulatory pathways that must be followed before placing a new medical device ono the market: premarket approval (PMA) and 510(k) clearance.

PMA is the more stringent of the two, and is applied to devices which are "represented to be for a use in supporting or sustaining human life" or which present a "potential unreasonable risk of illness or injury."[58] PMA requires the FDA to determine that sufficient, valid scientific evidence assures that the device is safe and effective for its

---

[56] *See, e.g.*, James Flaherty, Jr., *Defending Substantial Equivalence: An Argument for the Continuing Validity of the 510(k) Premarket Notification Process*, 63 FOOD DRUG L.J. 901. *FDA Mission Statement*, http://www.fda.gov/downloads/aboutfda/reportsmanualsforms/reports/budgetreports/ucm298331.pdf . (last accessed Feb. 9, 2017).

[57] As stated by William Maisel, FDA's acting director of the Office of Device Evaluation, at a public workshop on Robotic Assisted Surgical Devices (July 27, 2015), "[T] he first prong of our vision is that patients in the U.S. have access to high quality, safe and effective medical devices of public health importance first in the world. …if we set our evidentiary bars to high, then a lot of really great ideas will never make it. And so, we have to appropriately balance the availability of these technologies, getting these technologies to market and also make sure that they remain safe and effective. …[W] e also need to think about what is the cost of the development of the technology… [I]f Studies, the cost of developing a technology is too high, then many of those technologies will never make it to patients. And so, striking the right balance important." (transcript available at http://www.fda.gov/MedicalDevices/NewsEvents/WorkshopsConferences/ucm435255.htm).

[58] 21 U.S.C. § 360c(a)(1)(C)

intended use.[59] Thus, a PMA applicant generally must provide results from clinical investigations involving human subjects showing safety and effectiveness data, adverse reactions and complications, patient complaints, device failures, and other relevant scientific information.[60] The application is often reviewed by an advisory committee made up of outside experts.[61]

The FDA estimated in 2005 that reviewing one PMA application costs the agency an average of $870,000.[62] One survey of medical device companies found that it took an average of 54 months to reach approval from first communication with the FDA about an innovation.[63] The same survey found that the average total costs for a medical device company from the time of product conception to approval was $94 million, although these cost cannot all be attributed to compliance activities.[64] Fifty-two new devices received PMA approval in 2015.[65]

---

[59] 21 C.F.R. 814
[60] 21 C.F.R. 814.20(6)(ii)
[61] CRS report, page 12-13
[62] Gov't Accountability Off., *SHORTCOMINGS IN FDA's PREMARKET REVIEW, POSTMARKET SURVEILLANCE, AND INSPECTIONS OF DEVICE MANUFACTURING ESTABLISHMENTS*, TESTIMONY BEFORE THE SUBCOMMITTEE ON HEALTH, COMMITTEE ON ENERGY AND COMMERCE, HOUSE OF REPRESENTATIVES, 5(2009) http://www.gao.gov/new.items/d09370t.pdf .
[63] Josh Makower, Aabed Meer, & Lyn Denend, FDA IMPACT ON U.S. MEDICAL TECHNOLOGY INNOVATION: A SURVEY OF OVER 200 MEDICAL TECHNOLOGY COMPANIES (Nov. 2010), 23 (available at http://advamed.org/res.download/30).
[64] *Id.* at 28.
[65] DEVICES APPROVED IN 2015, http://www.fda.gov/MedicalDevices/ProductsandMedicalProcedures/DeviceApprovalsandClearances/PMAApprovals/ucm439065.htm.

The 510(k) pathway is more popular, with 3,006 clearances in 2015.[66] 510(k) applies to moderately risky devices, and clears a device for marketing if it is "substantially equivalent" to a "predicate" device already on the market. A predicate device is a device that was available on the market before 1976, or any device cleared since then via 510(k). The FDA will clear a device as substantially equivalent to an earlier device if:

> (1) The device has the same intended use as the predicate device; and
> (2) The device:
> (i) Has the same technological characteristics as the predicate device; or
> (ii) (A) Has different technological characteristics, such as a significant change in the materials, design, energy source, or other features of the device from those of the predicate device;
> (B) The data submitted establishes that the device is substantially equivalent to the predicate device and contains information, including clinical data if deemed necessary by the Commissioner, that demonstrates that the device is as safe and as effective as a legally marketed device; and
> (C) Does not raise different questions of safety and effectiveness than the predicate device.[67]

A 510(k) applicant must therefore submit information about the device's design, characteristics, and relationship to a predicate device, and any data backing up those claims.[68] In contrast to PMA, human-subject clinical trials for safety and effectiveness

---

[66] DEVICES CLEARED IN 2015, http://www.fda.gov/MedicalDevices/ProductsandMedicalProcedures/DeviceApprovalsandClearances/510kClearances/ucm432160.htm.
[67] 21 CFR 807.100(b)
[68] Congressional Research Service, FDA REGULATION OF MEDICAL DEVICES, 9 (2012)

are typically not required.[69] However, the FDA can respond to a 510(k) application by requesting additional information it deems relevant,[70] which may lead to frustration over the unpredictability of the clearance process.[71]

A 510(k) application is significantly cheaper for the FDA to review, at an estimated average cost of $18,200 per application.[72] A company's total costs from product concept to clearance is around $31 million on average.[73] Although FDA hopes to reach a final decision on each application within three months, U.S. companies reported an average time of 10 months from first submission of an application to clearance.[74] This faster timeline and the lower evidentiary requirements make 510(k) appealing to device companies over PMA.

Autonomous systems, however, are unlikely to be allowed through the 510(k) pathway, at least assuming a good faith application of the regulatory language requiring substantial equivalence to a predicate device.[75] No available predicate device is autonomous, meaning that a new autonomous device would have new technical features. An autonomous device, now actively completing a task that the alleged

---

[69] *Id.*
[70] 21 CFR § 807.100(a)(3)
[71] Makower, *supra* note XXX at 26 (A Stanford-based survey of 200 U.S. medical device companies found that over half ranked FDA regulatory performance as "mostly unpredictable" or "very unpredictable", as compared to less than 5% of respondents ranking European Union device regulation in the same category.).
[72] Gov't Accountability Off., *supra* note XXX at 5.
[73] Makower, *supra* note XXX at 7. (Again, these are total costs, not compliance costs).
[74] *Id* at 26.
[75] 21 C.F.R. § 807.100(b); No mechanism exists for a third-party to challenge a 510(k) clearance.

predicate device was used by a human to complete, almost certainly raises new questions of safety or effectiveness under 21 C.F.R. § 807.100(b)(2)(ii)(C), quoted above. In simpler terms, the new risks associated with autonomous systems imply that the FDA will demand more thorough review. Companies with autonomous medical devices must face the prospects of the longer, more expensive PMA process.

The FDA also has the ability to monitor devices after they are put on the market. Many PMA approvals require postmarket surveillance studies to gather further safety and efficacy data. Postmarket study may also be required by 510(k) clearance.[76] FDA regulations also mandate reporting of device-related adverse events by device manufacturers and health care facilities, and allow reporting of such events by patients.[77] Finally, FDA may issue recall orders for marketed devices which are found to pose health hazards.[78]

Overall, FDA regulation—and more particularly, the pre-market approval process likely to be required for most autonomous medical devices—focuses on the verification and validation stages of the product life cycle. The FDA also reaches into the management plan phase through its design controls, and the interconnected nature of the Systems Engineering process means that systems engineers must consider what the FDA will be looking for when formulating a concept of operations or drafting

---

[76] 21 U.S.C. § 360c(a)(1)(B).
[77] Congressional Research Service, *supra* note 46 at 15–16.
[78] 21 CFR § 810

requirements. Regulatory oversight of these other phases, however, is indirect. FDA thus falls somewhere in between full control of private actors' systems engineering frameworks and mere review of test results at the end of clinical trials. In terms of federal agency presence in a company's engineering management phases, the FDA's medical device regulation therefore is a middle ground for safety-critical technology regulation.

4.3 Automotive Regulation

A popular place to talk about emerging autonomous systems is in reference to the automotive industry. Self-driving cars and platooning trucks are among the range of technological dreams envisioned for public roadways. This section looks at the existing federal framework for the regulation of automotive technologies and summarizes NHTSA's recent statements related to autonomous vehicles.

Federal regulation of the automotive industry is centered in the National Highway Traffic Safety Administration, whose mission is to "save lives, prevent injuries, and reduce economic costs due to road traffic crashes through education, research, safety standards, and enforcement activity."[79] NHTSA has authority over motor vehicles and motor vehicle equipment, a term it interprets broadly to include all components, accessories, and software which impacts the safety of a vehicle.[80] With

[79] NHTSA Strategic Plan 2016-2020, https://www.nhtsa.gov/about-nhtsa
[80] Safety Act ;

respect to technical features of cars, trucks, motorcycles, and other motor vehicles on public roadways, NHTSA attempts to assure safety through two mechanisms: minimum safety standards and recall authority.[81]

4.3.1 Federal Motor Vehicle Safety Standards and Compliance Testing

NHTSA administers the Federal Motor Vehicle Safety Standards (FMVSS), which provide minimum safety requirements to be followed by vehicle manufacturers. [82] Established through notice-and-comment rulemaking, the FMVSS consist of 73 separate standards grouped generally into three categories: crash avoidance, crashworthiness, and post-crash survivability.[83] These minimum safety standards address most safety-related aspects of a vehicle, including headlights, brake lights, and turn signals; windshield defrosting and washing; brake systems; tires; mirrors; electronic stability control; door locks; seat belts; motorcycle helmets; bus emergency exits; flammability; and fuel system integrity.[84] The FMVSS can be very specific, dictating sub-component requirements as well as the objective tests needed to show compliance.[85] The FMVSS are

---

[81] NHTSA Enforcement Bulletin 2016–02: Safety-Related Defects and Automated Safety Technologies, Federal Register Vol. 85, No. 185 (9/23/2016), 65705-9., 65707.
[82] 49 C.F.R. § 571
[83]Review of Federal Motor Vehicle Safety Standards(FMVSS) for Automated Vehicles https://ntl.bts.gov/lib/57000/57000/57076/Review_FMVSS_AV_Scan.pdf
[84] 49 C.F.R. § 571
[85] *See, e.g.*, 49 C.F.R. § 571.104 S4.1.2 ("Wiped area. When tested wet in accordance with SAE Recommended Practice J903a (1966) (incorporated by reference, see § 571.5) each passenger car windshield wiping system shall wipe the percentage of Areas A, B, and C of the windshield (established in accordance with S4.1.2.1) that (1) is specified in column 2 of the applicable table following subparagraph S4.1.2.1 and (2) is within the area bounded by a perimeter line on the glazing surface 25 millimeters from the edge of the daylight opening.")

therefore very important in the Requirements Engineering phase of the V-Model, and then in all design, implementation, and verification tasks that must be traced back to those requirements.

The FMVSS apply to all relevant motor vehicles manufactured on or after the effective date of the standard,[86] but NHTSA does not independently test each vehicle for compliance with all FMVSS before it reaches the market. Instead, manufacturers of motor vehicles must self-certify that their vehicles comply with all relevant FMVSS. [87] Although the FMVSS sometimes include language related to verification testing, NHTSA does not require any particular certification process: instead, "the manufacturer takes whatever actions it deems appropriate" to provide certification before first sales.[88] The manufacturer is also expected to monitor compliance of production vehicles. [89]

After vehicles are available on the market, the Office of Vehicle Safety Compliance (OVSC) buys cars for testing from real-world new-car dealerships.[90] This procedure hopes to ensure that "the test specimens selected are a true representation of the product which could be purchased by the consumer."[91] The purchased vehicles are

---

[86] 49 C.F.R. § 571.7(a)
[87] Office of Vehicle Safety Compliance, Compliance Testing Program https://one.nhtsa.gov/cars/testing/comply/Mission/1_ovsc_1.html .
[88] Office of Vehicle Safety Compliance, Compliance Testing Program https://one.nhtsa.gov/cars/testing/comply/Mission/1_ovsc_1.html .
[89] *Id.*
[90] *Id.*
[91] *Id.*

then put through testing that targets an average of 30 of the 44 testable FMVSS.[92] Due to budget limitations—OVSC has a total annual budget of around $20 million[93]—the majority of vehicle makes and models are never tested by the government, although OVSC does prioritize testing targets to investigate the highest risk standards and vehicle models.[94] According to NHTSA, "[i]nstances of non-compliance, especially non-compliances having substantial safety implications, are rare."[95]

Importantly, the FMVSS do not bar a manufacturer from including other technological features in a vehicle. Autonomous technology is one area of technology not addressed by existing FMVSS that manufacturers are placing in new vehicles. And although some FMVSS refer to human drivers and may therefore need some revision,[96] NHTSA openly states the ramifications of its anti-precautionary approach:

> Therefore, if a vehicle is compliant within the existing FMVSS regulatory framework and maintains a conventional vehicle design, there is currently no specific federal legal barrier to an HAV being offered for sale.[97]

---

[92] *Id.*
[93] https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/fy-2013_budget_highlights.pdf
[94] Office of Vehicle Safety Compliance, Compliance Testing Program https://one.nhtsa.gov/cars/testing/comply/Mission/1_ovsc_1.html .
[95] NHTSA HAV Policy, p. 72 https://www.transportation.gov/sites/dot.gov/files/docs/AV%20policy%20guidance%20PDF.pdf
[96] Review of FMVSS of Autonomous Vehicles, 2016, https://ntl.bts.gov/lib/57000/57000/57076/Review_FMVSS_AV_Scan.pdf
[97] NHTSA HAV Policy, p. 11

However, NHTSA retains authority to evaluate technologies not addressed by the FMVSS. The following section addresses the agency's powers to investigate incidents to identify defects and order recalls, even without an on-point FMVSS.

4.3.2 Enforcement Authority: Defects and Recalls

NHTSA's statutory grant includes the investigative authority to determine that a "[1] vehicle or equipment contains a defect related to motor vehicle safety or [2] does not comply with an applicable motor vehicle safety standard," and order recalls in response.[98] The second half of this quotation refers to the FMVSS discussed above, but the first half grants authority to NHTSA to identify a vehicle as defective independent of any explicit safety standard.[99] This gives NHTSA a more flexible option for addressing emerging technologies than reliance on formal FMVSS.[100]

NHTSA considers something to be a "defect" if it "poses an unreasonable risk to motor vehicle safety."[101] A defect determination can be based on one of two findings. First, a defect exists when the engineering or root cause of a failure is known.[102] NHTSA recognizes and will act on such a defect regardless of whether there have been any real-

---

[98] 49 U.S.C. 30118(a)
[99] U.S. v. Chrysler Corp., 158 F.3d 1350-1 (D.C. Cir. 1998).
[100] NHTSA Enforcement Bulletin 2016–02: Safety-Related Defects and Automated Safety Technologies, Federal Register Vol. 85, No. 185 (9/23/2016), 65705-9., 65707.
[101] *Id.*
[102] *Id.* at 65708.

world errors, so long as there is "a likelihood that it will cause or be associated with a 'non-negligible' number of crashes, injuries, or deaths in the future."[103]

Second, even if the root cause of a failure is not known, a vehicle or equipment is defective "if it is subject to a significant number of failures in normal operation, including failures either occurring during specified use or resulting from owner abuse . . . that is reasonably foreseeable."[104] In other words, NHTSA can identify a defect based on failure rates without needing to understand exactly how the underlying mechanism or software works or fails.[105]

Once a defect posing ongoing threats to motor vehicle safety is identified, NHTSA notifies the manufacturer of the defect and the company's obligations to remedy the defect.[106] The manufacturer can then choose to remedy the defect through repair, replacement, or refund.[107] In other words, NHTSA's defect notifications can be seen as orders to recall a vehicle or part. NHTSA has the authority to carry out civil enforcement actions and impose civil penalties if manufacturers do not comply with orders to remedy defects.[108]

---

[103] *Id.* at 65708
[104] *Id.*
[105] *Id.*
[106] 49 U.S.C. § 30118(a)
[107] 49 U.S.C. § 30120
[108] NHTSA Enforcement Bulletin 2016–02: Safety-Related Defects and Automated Safety Technologies, Federal Register Vol. 85, No. 185 (9/23/2016), 65705-9., 65707.

In September 2016, NHTSA issued a guidance document asserting that its investigation and recall authority extends to emerging automated safety technologies. NHTSA interprets the definition of "motor vehicle equipment" to encompass software components including cloud computing, features added through over-the-air updates, and any other innovative system related to motor vehicle control or safety.[109] The guidance stresses that autonomous vehicle manufacturers have a "continuing obligation to proactively identify safety concerns and mitigate the risks of harm," while "strongly encourag[ing]" them to "resolve safety concerns before their products are available for use on U.S. roadways." The guidance is essentially a statement that NHTSA intends to use its existing investigation and recall authority to police any automated driving systems that reach the market, even though the agency is currently powerless to impose pre-market controls.

Meanwhile, NHTSA's practical ability to investigate defects may be quite limited. The office responsible for investigations and defect findings only employs 20 investigators, with a budget of less than $23 million dollars to fund investigations of related to defects in the 250 million cars on American roadways.[110] NHTSA's fiscal year 2017 budget request asks for an increase of $25 million for the safety defects investigation program, backed in part by the fact that "the advancement of in-vehicle

---

[109] NHTSA Enforcement Bulletin 2016–02: Safety-Related Defects and Automated Safety Technologies, Federal Register Vol. 85, No. 185 (9/23/2016), 65705-9

[110] "FY 2017 NHTSA Budget Estimate" at 28 and 63, https://www.nhtsa.gov/about-nhtsa/2017-budget.

electronics and automation will increase complexity of safety issues warranting attention and possible investigation."[111] The budgetary and expertise limits of NHTSA's investigation team are likely to limit the agency's ability to properly address the coming influx of autonomous systems through defect investigations.

4.3.3 Investigation of an Autonomous Systems

In fact, NHTSA has already investigated at least one highly-automated vehicle for a defect, issuing a report in January 2017 finding no defect but giving clues about how NHTSA will approach future investigations of autonomous systems. Following the first crash which occurred in Tesla's "Autopilot" mode, NHTSA opened an investigation into whether the autonomous system had a defect.[112] A report on the investigation was published in January 2017, and may give some clues as to how NHTSA will evaluate autonomous systems going forward.[113] Data from the car showed that Tesla's automated emergency braking, traffic-aware cruise control, and autosteering technologies were engaged at the time of the crash, which involved a Tesla Model S running into a semi-trailer that pulled across its path.

Beginning with analysis of the technical features to look for an identifiable root cause of failure, NHTSA first evaluated Tesla's emergency braking system, comparing it

---

[111] *Id.*
[112] "Tesla Autopilot Investigation Report", Jan. 2017, https://static.nhtsa.gov/odi/inv/2016/INCLA-PE16007-7876.PDF.
[113] *Id.*

to other manufacturers' automated braking technology and finding it in-line with the state-of-the-art.[114] Further, NHTSA concluded that "braking for crossing path collisions, such as that present in the Florida fatal crash, are outside the expected performance capabilities of the system."[115]

NHTSA then moved on to consider Tesla's traffic-aware cruise control and autosteer features. With regard to these partially autonomous features, NHTSA focused on their human-machine interaction components.[116] Tesla's Autopilot falls under what NHTSA calls a Level 2 automated system, meaning that the vehicle completes steering and speed control tasks, but the human driver must monitor the driving environment and take control when necessary.[117] At Level 2, the NHTSA investigation report indicates the agency expects a manufacturer to at least consider ways to:

> 1) provide the operator with information about system limitations;
> 2) include a method for monitoring driver engagement with the driving task and assisting the driver with maintaining attention to the environment;
> 3) minimize the potential for mode confusion to occur, through intuitive feedback from vehicle dynamics and/or warnings to the driver; and

---

[114] *Id.* at 3–4
[115] *Id.* at 4.
[116] *Id.* at 4–8
[117] *Id.* at 5.

> 4) consider restricting availability or performance when used on roads that are not in the intended use operating environments.[118]

The report goes on to find that Tesla sufficiently addressed these concerns. First, it notes that system limitations were communicated in the owner's manual, in release notes alongside software updates, a user agreement required before engaging Autosteer, and a dialog box which appears when Autosteer is activated.[119] Second, Tesla monitors driver engagement by detecting whether the driver's hands are the steering wheel—if the hands are not detected, the system gives visual and audible alerts and may slow the car until the driver responds.[120] Interestingly, the report notes that Tesla's updates since the crash improved these features by locking a driver out of autopilot mode after failing to respond to the alerts.[121] Third, the report approvingly highlights the warnings and timing of how the Tesla system hands control back and forth between driver and automation.[122] Fourth, the report notes that the Tesla system can be engaged anywhere,

---

[118] *Id.* at 5–6. More broadly, it appears that NHTSA will import human-machine interface requirements from earlier reports it has published, citing here Human Factors Evaluation of Level 2 and Level 3 Automated Driving Concepts – Concepts of Operation. (2014). DOT HS 812 044. National Highway Traffic Safety Administration. Washington, DC.

[119] *Id.* at 6.

[120] *Id.* at 7.

[121] *Id.*; This is interesting because it shows how the economic penalty associated with recalls can be avoided if a software update remedies the defect.

[122] *Id.*

but appears to be convinced that the system limitation warnings are enough to mitigate this issue.[123]

In assessing the potential human-machine interaction design, the report notes that a manufacturer must consider unreasonable use due to owner abuse, including distraction.[124] More specifically, NHTSA wanted assurance that Tesla considered the foreseeable situation where a driver's gaze was off the road for more than seven seconds.[125] Based on a report that NHTSA compelled Tesla to produce on its own design and testing,[126] NHTSA concluded that:

> The potential for driver misuse was evaluated as part of Tesla's design process and solutions were tested, validated, and incorporated into the wide release of the product. It appears that Tesla's evaluation of driver misuse and its resulting actions addressed the unreasonable risk to safety that may be presented by such misuse.[127]

Thus, trusting Tesla's verification and validation activities, NHTSA found no identifiable engineering flaw that would justify a finding of a defect.

The NHTSA report also considered that a defect determination can be based on failure rates. Because Tesla collects so much data from its cars, it was able to show NHTSA data that Tesla crashes had dropped by almost 40 percent per mile traveled

---

[123] *Id.* at 8.
[124] *Id.* at 10.
[125] *Id.* at 9–10.
[126] 49 C.F.R. § 510.7
[127] Tesla Autopilot Investigation Report at 10.

after the software update which installed the Autopilot feature. This drastic, quantified safety improvement certainly made a strong case in Tesla's favor. No defect was found, and NHTSA closed the investigation.[128]

Although this investigative report like this one does not create binding legal precedent, it does indicate the approach NHTSA is likely to take towards its only existing means of regulating autonomous vehicles: through its investigation and recall authority. NHTSA asked questions about design and testing phases, suggesting some interest in leveraging those systems engineering phases in its regulatory approach while showing awareness of the issues related to human-robot interaction. Overall, however, NHTSA's power to address new technologies is located in the Operation & Maintenance block of the V-Model, making real-world crash statistics a primary influence on regulatory enforcement decisions.

4.3.4 NHTSA's Federal Automated Vehicles Policy

Also in September 2016, NHTSA released a much-anticipated document entitled "Federal Automated Vehicles Policy."[129] As an agency guidance document, the policy does not create binding legal obligations: instead it merely represents the agency's beliefs on certain issues at the time of publishing, reflecting preferences and future

[128] *Id.* at 12.

[129] https://www.transportation.gov/sites/dot.gov/files/docs/AV%20policy%20guidance%20PDF.pdf

courses of action likely to change under the new presidential administration.[130] The document tries to do several things, as described below, but does none of them convincingly.

First, the policy encourages manufacturers of automated vehicle systems to voluntarily submit information relating to the development of their technologies. The policy document lays out what categories of information should be included,[131] and makes recommendations for what manufacturers should do about those categories. For example, the section on cybersecurity states that "Manufacturers and other entities should follow a robust product development process based on a systems-engineering approach to minimize risks to safety, including those due to cybersecurity threats and vulnerabilities."[132] The document's recommendations are all at this or a higher level of abstraction and vagueness, and are not accompanied by any enforcement mechanisms. At best, these suggestions may be helpful in the sense that they give industry a heads-up on what NHTSA might look for in a defect investigation, and may be interesting for occasionally suggesting carmakers follow a systems engineering framework. The Trump administration—with Tesla and Uber executives on presidential advisory

---

[130] "Trump administration re-evaluating self-driving car guidance.", Reuters, Feb. 26, 2017, http://www.reuters.com/article/us-usa-trump-selfdriving-idUSKBN1650WA .
[131] *Id.* at 15 "The Safety Assessment would cover the following areas: • Data Recording and Sharing • Privacy • System Safety • Vehicle Cybersecurity • Human Machine Interface • Crashworthiness • Consumer Education and Training • Registration and Certification • Post-Crash Behavior • Federal, State and Local Laws • Ethical Considerations • Operational Design Domain • Object and Event Detection and Response • Fall Back (Minimal Risk Condition) • Validation Methods"
[132] *Id.* at 21.

panels—is unlikely to even try to make companies turn over this information in reality.[133]

Second, the document makes repeated explanations about how to propose a new FMVSS for agency consideration and eventual notice-and-comment rulemaking.[134] NHTSA appears to be hoping that an outside party will undertake the regulation drafting and cost-benefit analysis needed to develop a formal standard for an autonomous safety system.[135] One takeaway is that the agency appears to not have a plan to do that work itself, at least in the near term.

Lastly, the document lays out a model state policy for automated vehicles. Because differing state rules may make it difficult for autonomous car developers to make systems that can cross state lines legally, NHTSA hopes for some amount of standardization while trying to preserve the states' independent roles in licensing, liability, and policing. For purposes of the discussion of systems engineering that this thesis focuses on, the most interesting provisions of this model policy—as well as state actions to date—relate to allowances for testing autonomous systems in real environments. NHTSA recommends that states implement an application process for

---

[133] "Musk, Kalanick Join Trump Strategic Policy Forum", http://www.businessinsider.com/tesla-ceo-elon-musk-uber-travis-kalanick-join-donald-trump-strategic-policy-forum-economic-team-2016-12; "Trump administration re-evaluating self-driving car guidance.", Reuters, Feb. 26, 2017, http://www.reuters.com/article/us-usa-trump-selfdriving-idUSKBN1650WA .

[134] *Id.* at 7, 48, 87. The appendix section titled "NHTSA's Current Regulatory Tools" consists solely of a guidance for well-supported petitions for new rulemaking.

[135] *See, e.g., id.* at Appendix A.

testing of highly-automated vehicles in their jurisdictions. From the federal perspective, more state-allowed testing means more testing data from more localities to support NHTSA's attempts to determine the safety of a new system in diverse environments. The model policy, however, is written at a high level: instead of providing text that could be copy-and-pasted into a state law, NHTSA offered guidelines that could encompass many possible state-level implementations. Overall, the policy guidance document offers little practical advice.

4.3.5 Conclusions about NHTSA

In conclusion, NHTSA's general regulatory framework for automotive safety is focused on limited post-market testing and investigations after incidents. Where a specific regulation does not set a relevant safety standard—as is the case for new autonomous technologies—NHTSA allows the technology unless an investigation reveals an unreasonable safety risk, often with a focus on statistical outcomes. Although federal standards set some technical design requirements which must be achieved by the end product, NHTSA shows almost no involvement in a manufacturer's systems engineering process. NHTSA therefore demonstrates safety regulation taking the least-invasive approach with respect to regulated companies' engineering management.

4.4 Comparative Takeaways

FAA, FDA, and NHTSA therefore differ on the stage of systems engineering at which companies need to start working directly with their relevant agency. Figure 3

shows this point-of-first-contact overlaid on the V-Model. Earlier intervention allows FAA to maintain continuous relationships with aerospace companies, a small group relative to other industries, and give early advice about new product ideas. Abstaining from direct involvement before the real-world operations stage lets NHTSA make it easier for carmakers to sell new variations of each car model every year without significant regulatory burdens. FDA's middle ground approach may reflect a concern with the life-or-death nature of medical device risks but also a sense that a lack of funding or poor bench testing results will eliminate many medical devices inventions before the risk to human materializes.

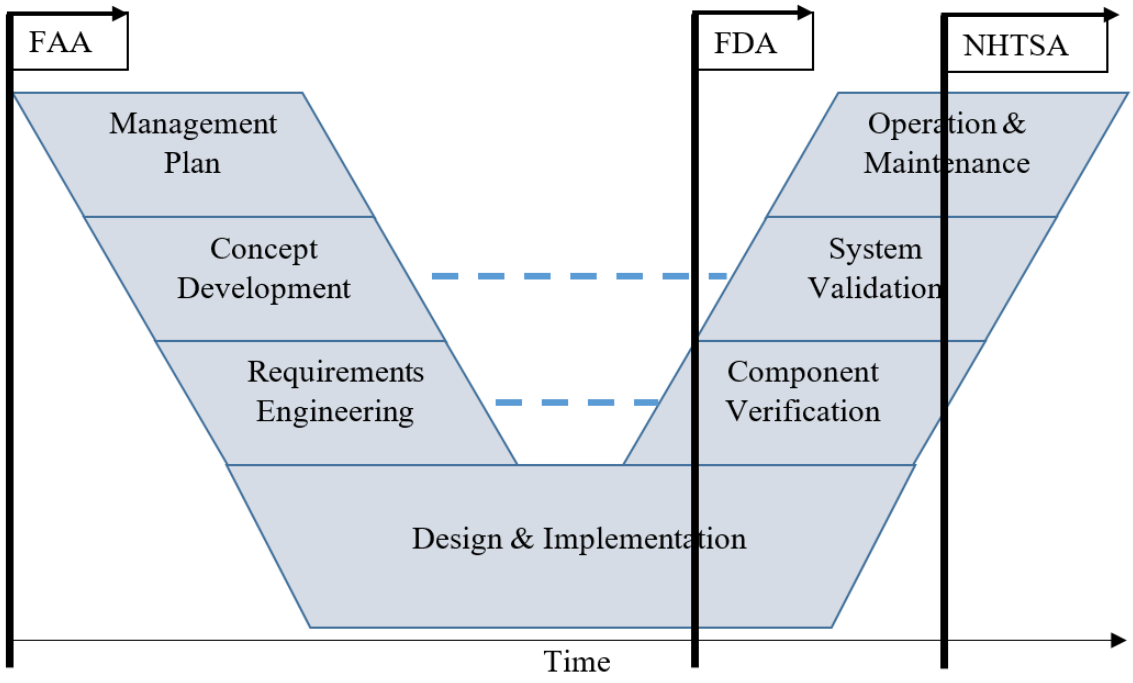

Figure 3: Point of First Company Contact with Agency in Safety Regulation

However, the point-of-first-agency-contact comparison does not tell the whole story. Rather, both FDA and NHTSA reach backwards across the V-Model to impact steps that occur before the agency becomes directly involved with a particular product. Table 1 shows how each agency's major regulatory tools fit into the V-Model blocks. Of course, the interconnectedness of V-Model blocks means that regulatory interventions in one V-Model block also influences neighboring blocks. For example, FDA's design controls require systems management planning at the beginning of the V-Model, and are therefore listed in the Management Plan row of the table, but the plan created there necessarily guides how a company goes through the remaining phases. The blank boxes for all three agencies in the Design & Implementation means that companies are fortunately left flexibility to innovate, but of course the systems designs are bounded by the need to produce a system compliant with rules flowing into the V-Model through other systems engineering blocks.

Table 1: Major Agency Regulatory Tools by Systems Engineering Block.

SE Block	FDA	FAA	NHTSA
Management Plan	Design Controls	Mutually-developed Partnership for Safety Plan & Project Specific Compliance Plan.	
Concept Development		Familiarization meetings.	
Requirements Engineering		Airworthiness standards, development of 'special conditions'.	FMVSS
Design & Implementation			
Component Verification	Test results part of IDE, 510(k), and PMA applications.	Mutually-designed test plans, Agency observation of tests (often using Designees).	FMVSS-defined tests; Self-certification.
System Validation	IDE approval; mutually-designed test plans; federal review of results to grant premarket approval.	Mutually-designed test plans; Agency observation of tests (often using Designees); conformity inspections; issuance of type certificate.	
Operations & Maintenance	Manufacturing controls; postmarket study; adverse event reporting; recall authority.	Production certificate; airworthiness certificate; mandatory issue reporting; airworthiness directives; user licensing and use restrictions.	Post-market compliance testing; defect investigations; recall/remedy authority.

Comparing the point-of-first-agency-contact visualized in Figure 3 with the regulatory interventions of Table 1, both the FDA and NHTSA show a gap in time between when a company must begin complying with agency rules and when that agency actively engages with the development of a given system. The longer this gap gets, the higher the losses will be for a company that is told that it failed to comply with

agency rules earlier in the V-Model. While earlier contact and collaboration between agencies and companies may sound like an increase in regulatory burdens, reducing regulatory uncertainty and the costs of noncompliance through early contact would likely result in overall saved regulatory compliance costs while benefiting the efficiency of an industries' innovative process.

In sum, these comparison show that two of the three agencies use systems engineering principles to adaptively regulate new technologies. FAA facilitates innovation by mutually determining regulatory expectations early in the product life-cycle; FDA requires a system engineering plan while relying heavily on extensive system validation tailored to the device in question. Although significant differences in regulatory structure exist, these two agencies explicitly leverage systems engineering to help manage the risks of emerging technologies and adapt certification standards to meet the needs of the latest innovation. Both have already cleared products with highly-automated features without having to alter their basic regulatory frameworks. Objective, quantitative proof of how well these regulations work is outside the scope of this paper, but what is clear is that the inclusion of systems engineering concepts helps the regulations adapt to changing technology.

NHTSA does not rely on system engineering in the same way, with its current framework rejecting a precautionary approach to new technological features on automobiles. This difference may be part of why it has felt so much pressure in recent

years to make new announcements about what it will do about self-driving cars. NHTSA certainly has a procedure in place for dealing with new technologies—by waiting to see how well they work and for the industry to converge on a standard—but that approach may not be sufficient to handle the complexity of autonomous systems. Lack of clarity, vague policy documents, and understaffed and underfunded case-by-case investigations currently leave holes that might be best solved by a fuller embrace of systems engineering in the regulatory scheme.

5. Conclusions and Recommendations

This final chapter returns to the principles of risk regulation discussed in the introduction to see what this paper's discussion of systems engineering illuminates about how regulators should approach emerging autonomous technologies. With the lessons of previous chapters in mind, analysis of the pacing problem, the precautionary principle, risk-risk tradeoffs, and multi-modal regulation reveal that a systems engineering approach starts to solve many of the problems associated with the regulation of emerging technologies

The first principle was the pacing problem: the idea that regulators and regulations cannot keep up with the changing pace of new technology.[1] Our three examples of regulatory schemes present two general approaches to the pacing problem. On one side, NHTSA's approach to the pacing problem is to acquiesce to it. By allowing new technologies as a default and waiting to respond to problems through post-market investigations, NHTSA accepts that it will always lag behind industry.[2] NHTSA waits to respond with rulemaking until the new technological feature becomes well-understood in the industry, reducing its need to act in the face of uncertainty about where technology is headed.[3] On the other side, FAA and FDA use systems engineering as a

---

[1] *Infra,* Chapter 1.
[2] *See Infra*, Subchapter 4.3.
[3] *Id.*

tool to identify, mitigate, and test the risks associated with each new innovation before it is allowed on the market.[4] Regulators taking this approach must develop the technical knowledge necessary to understand and evaluate each new product, and innovation efforts may be stifled by the delays this causes, but the underlying regulatory frameworks need not be altered to address emerging technologies. Working collaboratively with regulated companies to develop requirements and test plans can help combat the difficulties inherent in keeping regulators educated on the latest advancements. Reliance on systems engineering principles in premarket regulation allows for adaptation of regulatory oversight and standards to match the needs of new technologies on an *ex ante* basis.

The contrast in these approaches brings us to the second concept: the precautionary principle, the idea that a technology should be strictly limited until safety is proven.[5] FAA and FDA are both precautionary in nature: a new product is not allowed to be put to real-world use without agency permission.[6] These precautionary approaches use systems engineering principles and system validation testing to generate the data needed to show that a new system is safe. Requiring agency review of regulated companies' systems engineering practices before a product reaches the market therefore captures the concerns of the precautionary principle, but also actively

---

[4] *See, infra* Subchapters 4.1, 4.2, 4.4.
[5] *Infra* Chapter 1.
[6] *Infra* Subchapters 4.1 and 4.2.

motivates data generation to prove acceptably low risks of an innovative technology. NHTSA makes the opposite presumption, namely that a new technological feature on a car is safe enough to put on the market unless affirmatively shown to pose an unreasonable risk.[7] NHTSA's presumption is understandable when one considers that until now new car features—like back-up cameras, tire pressure monitors, or blind spot warnings—were quite unlikely to make a car more dangerous. With autonomous systems entering the game and posing new kinds of risk, the same can no longer be said with such confidence. The complexities and unknowns of autonomous systems—the associated risks of system failure, human-system interaction breakdowns, and social disruptions[8]—seem to justify a more precautionary approach.

The third concept is risk-risk tradeoffs, the idea that each risk management tool carries its own risks.[9] For any regulatory agency dealing with product safety, the major risk-risk tradeoff to consider is between the injuries caused by allowing a product on the market too early versus the injuries that could have been eliminated by a new product had it been let on the market sooner. Awareness of this risk-risk tradeoff makes clear that both extreme regulatory skepticism and complete non-regulation have their drawbacks. In practice, systems engineers are always managing this tension between budgeted timeline and product quality. Our examples from the FDA and FAA show

---

[7] *Infra* Subchapter 4.3.
[8] *Iinfra* Chapter 1.
[9] *Id.*

how regulators have embraced this idea, with systems-engineering-based regulations tailored to the needs of specific industries.[10] Thus, systems engineering can provide a foundation for regulators attempts to strike a balance between the risks of over- and under-regulation.

Lastly, the introduction noted that regulation includes more than just law, as social norms, markets, and architecture shape human behavior and technological change.[11] Take the three alternatives to law in turn. First, social norms shape adoption of a new technology, particularly important for autonomous systems that put humans in new relationships with machines. The risks of autonomous systems that involve safety-critical human-robot interaction will be entangled with changing social norms related to people's social conduct relative to autonomy: one sees early hints of these issues in NHTSA's investigation of Tesla's autopilot, where the agency recognized that carmakers should assume that users will be inattentive when using automated tools.[12] Although social norms for a new technology may be difficult to predict *ex ante*, the concept development phase and system validation testing phases of a systems engineering framework can help engineers—and regulators—try to grasp the social norms and user intuitions that will impact the safety of a new autonomous system.

---

[10] *Infra,* Subchapter 4.4.
[11] *Infra,* Chapter 1.
[12] *Infra*, Section 4.3.3.

Second, markets regulate behavior, perhaps an obvious point in a discussion largely centered on the traits of products that profit-driven companies will try to create. Companies building autonomous systems want to make money, and thus need their systems to reflect levels of risk tolerable to their customers. Systems engineering practices can help companies make sure that they build autonomous systems that people will pay for, by engaging important stakeholders early in the concept development and requirements engineering stages, and tracing that input through design and implementation decisions. Monitoring the way markets shape risk management can guide regulators decisions on when to intervene.

Third is architecture, the idea that the traits of physical or cyber spaces guide people into certain behaviors. Relevant to our discussion, autonomous systems can be engineered to nudge or force people to use them in a safe way. In the examples above, NHTSA's endorsement of the way Tesla uses technology to try to keep users focused on the road—through micro-torque sensors in the steering wheel measuring whether a user's hands were touching the wheel, autopilot lockouts after repeated distraction, and other warnings—shows how regulators are likely to come to rely on autonomous systems' interface design to push people into safe uses of those systems.[13] As discussed in Chapter 2, analyzing interactions between a new technology and its surrounding

[13] *Infra* Section 4.3.3.

users or environment is a key facet of the systems engineering perspective. Systems engineering practices are necessary to identify the need for these kinds of architectural regulations, design and implement them, and test their efficacy.

Together, these four principles describe the basic concerns of scholars when they talk about the challenges of regulating emerging technologies.[14] For autonomous systems—complex engineered products which pose risks of product failure, human-machine breakdowns, and societal disruption—the foregoing arguments suggest that a regulatory approach based in systems engineering starts to solve the major issues. Therefore, this thesis recommends a systems engineering approach to regulating autonomous systems.

What should a systems engineering approach to regulating autonomous systems look like? The details of the regulatory schemes adopted for different industries will necessarily be different, due to the diverse nature of innovation across sectors, the risks associated with different types of products, and other outside forces like jurisdictional limits. However, several elements will always be key to getting the full value out of systems engineering to the benefit of regulators' policy goals. Figure 4 illustrates these elements.

---

[14] *Infra* Chapter 1.

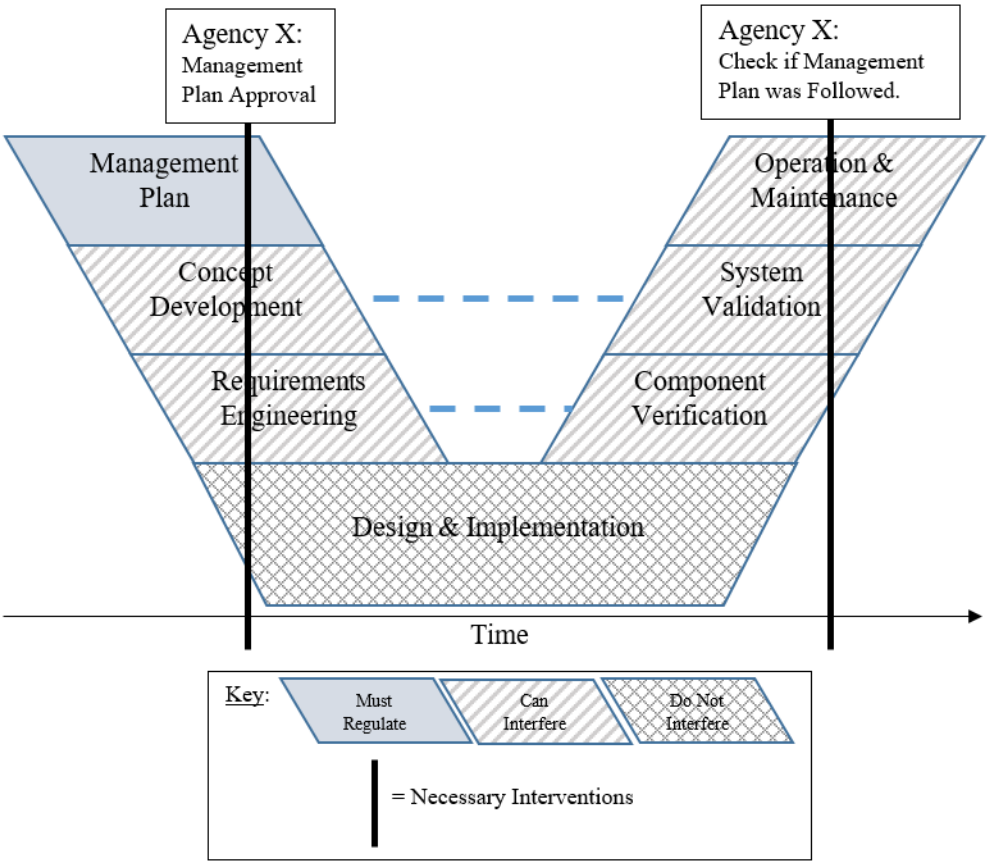

Figure 4: Systems Engineering Approach to Regulating Autonomous Systems

First, rules compelling the creation of a systems engineering management plan are essential. Requiring a management plan makes companies consider systems engineering principles from the beginning of the product life-cycle, guiding them to follow the steps through their entire development process. The full benefit of systems engineering's interconnected phases and activities cannot be achieved without a plan at the outset to rigorously adhere to best practices. By definition, then, a systems engineering approach to regulating autonomous systems includes requiring companies to make a plan to follow systems engineering practices.

Second, the management plan should be reviewed and revised in early collaboration between regulators and regulated companies. Early contact between agencies and industry allows for clarification and relative certainty about what steps the agency expects companies to follow to manage the risks of new systems, limiting the potential burden of investment into noncompliant development. Reviewing the management plan makes clear that regulators are serious about the benefits of systems engineering practices.

Third, regulators should avoid direct interference with the bottom block of the V-Model, instead leaving the Design and Implementation stage to companies to push the bounds of technology. Flexibility in that stage is necessary to allow firms in an industry to compete on innovation and design, and thereby make contributions towards achieving the societal benefits offered by autonomous systems. As shown in Table 1 in Subchapter 4.4, current regulatory schemes understand the importance of this freedom. Other blocks of the V-Model can be implicated directly by regulation when needed to achieve policy objectives.

Lastly, a systems engineering approach to regulating autonomous systems requires agency review towards the end of the V-Model to determine whether the management plan was followed satisfactorily. Across industries, this review can take different forms. For the highest-risk systems, direct agency involvement in the verification and validation activities laid out in the management plan may be justified.

For others, presentation of documents demonstrating that the mutually-developed systems engineering plan was followed through all steps might suffice. Agency certification that systems engineering practices were followed ensures that companies followed the risk management principles best suited to handle the risks of their new autonomous systems.

In application, these four elements of a systems engineering approach must be molded and complemented to meet the needs of a given agency. But, in general, this approach to regulation captures the truth that autonomous systems pose risks as complex products that are best mitigated by following the practices of systems engineering—practices themselves developed by engineering firms seeking to create new, cutting-edge technologies. Regulations based in systems engineering could drive society safely into the future.

References

14 C.F.R. subchapter C (Airworthiness Standards).

49 C.F.R. § 571 (Federal Motor Vehicle Safety Standards)

Adam B. Jaffe et al., "Environmental Regulation and the Competitiveness of U.S. Manufacturing: What does the Evidence Tell Us?", *Journal of Economic Literature* (March 1995) http://www.ucl.ac.uk/cserge/Jeffe%20et%20al%201995.pdf.

Alberto Galasso and Hong Luo, "Tort Reform and Innovation" Harvard Business School Working Paper 16-093, 2016. http://www.hbs.edu/faculty/Publication%20Files/16-093_14c952bf-4842-4ed7-b785-f4b8ae39875b.pdf

Alexander Kossiakoff and William N. Sweet, "Systems Engineering: Principles and Practice", Wiley 2003, 6.

Andrea Renda, et al., "Selecting and Designing European ICT Innovation Policies," *EU Science Hub*, 2016. ("The nature of the ICT ecosystem determines a growing need for flexible, adaptive regulation.").

Anita Kim et al., "Review of Federal Motor Vehicle Safety Standards(FMVSS) for Automated Vehicles", NHTSA, 2016, available at https://ntl.bts.gov/lib/57000/57000/57076/Review_FMVSS_AV_Scan.pdf

Eric Schatzberg, "History of Technology", lectures Spring 2013, Madison, WI.

Executive Office of the President, "Preparing for the Future of AI", at Introduction, 2016. (Obama Adminsitration) https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf

FAA, "The FAA and Industry Guide to Product Certification", 2nd Edition, Sept. 2004.

FDA, "Early Collaboration Meetings Under the FDA Modernization Act (FDAMA): Final Guidance for Industry and for CDRH Staff" http://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm073604.htm

Gary Marchant, "The Growing Gap Between Emerging Technologies and the Law", in *Growing gap between emerging technologies and legal-ethical oversight*, Springer 2011.

Geoffrey M. Hand, "COMMENTS: Should Juries Decide Aircraft Design? Cleveland v. Piper Aircraft Corp. and Federal Preemption of State Tort Law", 29 U.S.F. L. Rev. 741, 754-756

James Flaherty, Jr., *Defending Substantial Equivalence: An Argument for the Continuing Validity of the 510(k) Premarket Notification Process*, 63 Food Drug L.J. 901. *FDA Mission Statement*, http://www.fda.gov/downloads/aboutfda/reportsmanualsforms/reports/budgetreports/ucm298331.pdf . (last accessed Feb. 9, 2017).

Jason Millar and Ian Kerr, "Delegation, Relinquishment and Responsibility: The Prospect of Expert Robots", *Robot Law* (eds. Calo, Froomkin, Keer) 2016.

John Graham and Johnathan Wiener, "Confronting Risk Tradeoffs", *Risk v. Risk* (1995).

Johnathan Wiener, "The regulation of technology, and the technology of regulation", *Technology in Society* (2004), available at: http://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=1960&context=faculty_scholarship.

Jonathan B. Wiener, "Precautionary Principle," chapter in *Principles of Environmental Law* (Ludwig Krämer and Emanuela Orlando, eds.) of the *Encyclopedia of Environmental Law* (Michael Faure, ed.) (IUCN and Edward Elgar, forthcoming 2017)

Josh Makower, Aabed Meer, & Lyn Denend, FDA Impact on U.S. Medical Technology Innovation: A Survey of Over 200 Medical Technology Companies (Nov. 2010), 23 (available at http://advamed.org/res.download/30).

Laurence Lessig, "The Law of Horse: What Cyberlaw Might Teach", Harvard Law Review, Fall 1999.

Madeline Elish, "Moral Crumple Zones", We Robot 2016.

Mary Cummings and Jason Ryan, "Who Is in Charge? The Promises and Pitfalls of Driverless Cars." Transportation Research Board 2014.

Mary Cummings, "Systems Engineering", lectures Fall 2015, Durham, NC.

Mary Cummings, "The Brave New World of Driverless Cars", TR News (in press), available at http://hal.pratt.duke.edu/publications.

Mica Endsley, "Automation and Situation Awareness," Automation and human performance: Theory and applications, R. Parasuraman & M. Mouloua, eds. (1996).

MITRE Systems Engineering Guidebook, https://www.mitre.org/publications/systems-engineering-guide/about-the-seg

National Conference of State Legislatures, "Autonomous Vehicles Legislative Database", http://www.ncsl.org/research/transportation/autonomous-vehicles-legislative-database.aspx.

NHTSA Enforcement Bulletin 2016–02: Safety-Related Defects and Automated Safety Technologies, Federal Register Vol. 85, No. 185 (9/23/2016), 65705-9., 65707.

NHTSA Strategic Plan 2016-2020, https://www.nhtsa.gov/about-nhtsa

NHTSA, "Federal Automated Vehicles Policy", Sept. 2016, https://www.transportation.gov/sites/dot.gov/files/docs/AV%20policy%20guidance%20PDF.pdf.

Nicholas A. Ashford and Ralph P. Hall, "The Importance of Regulation-Induced Innovation for Sustainable Development" *Sustainability*, 2011, http://www.mdpi.com/2071-1050/3/1/270/pdf.

Office of Defect Investigations "Tesla Autopilot Investigation Report", NHTSA, Jan. 2017, https://static.nhtsa.gov/odi/inv/2016/INCLA-PE16007-7876.PDF.

Office of Vehicle Safety Compliance, "Compliance Testing Program", NHTSA, https://one.nhtsa.gov/cars/testing/comply/Mission/1_ovsc_1.html

Paul N. Otto and Annie I. Anton, "Addressing Legal Requirements in Requirements Engineering", *15th IEEE International Requiremetns Engineering Conference* (2007).

Ryan Calo, "Open Robotics", Maryland L. Rev. 2011.

Ryan Calo, "Robotics and the Lessons of Cyberlaw", Cal. L.R., 2015 http://www.californialawreview.org/wp-content/uploads/2015/07/Calo_Robots-Cyberlaw.pdf

Ryan Calo, "The Case for a Federal Robotics Commission," Brookings, 2104, https://www.brookings.edu/research/the-case-for-a-federal-robotics-commission/