GOVERNMENT-TO-ROBOT ENFORCEMENT
Forthcoming, 2019 UNIV. ILLINOIS L. REV. __.
Draft. Please contact author for permission to cite.

Susan C. Morse[*]

INTRODUCTION

Law is usually enforced on specific actors. A plaintiff claims that a defendant has violated the law. A court or other forum determines who is right and imposes and enforces penalties. This individualized approach to enforcement holds true for government regulation too. Countless regulatory schemes touch every aspect of contemporary life. That is, enforcing compliance with regulatory schemes generally proceeds on an individualized

---

[*] Angus G. Wynne Sr. Professor of Civil Jurisprudence, University of Texas School of Law. smorse@law.utexas.edu. Thanks to …

model. These schemes touch every aspect of contemporary life. They include, for instance, the federal income tax law, wage and hour reporting requirements, and environmental regulation. In other words, an enforcer calls an individual taxpayer, or employer, or polluter, to answer for its actions.

This historic idea that law is enforced on specific actors is fading. There is a sea change underway that turns the longstanding tradition of individualized enforcement on its head. The change involves two converging forces – the automation and centralization of law. These will change the locus of disputes so that disputes are no longer individualized. Instead, disputes will occur between, on the one hand, government regulators; and, on the other hand, the makers of automated, centralized legal systems. Prior work treats cyberspace as a phenomenon whose architecture law should seek to influence. In contrast, this Article treats automated, centralized law systems as instruments that can be used to enforce independently existing law.[1]

Machines already implement the law. They evaluate bank solvency,[2] determine welfare benefits,[3] respond to copyright-based Internet takedown requests,[4] and avoid liability for car accidents.[5] A human programmer might state a legal rule – such as a rule that following drivers are always responsible for rear-end collisions – and direct a robot to implement the rule through a logical algorithm.[6] Or, a computer might itself use machine learning techniques to derive a legal decision from a database of primary sources, such as cases, regulations, and empirical correlations between driving behavior and car accidents.[7]

Centralization means that a single action – for instance, writing a bit of software code – produces legal decisions for many individuals at once. Say that

---

[1] *See* LAWRENCE LESSIG, CODE AND OTHER LAWS OF CYBERSPACE (1999) (arguing that democratic mechanisms should oversee and edit the "architecture" of cyberspace).

[2] Kenneth A. Bamberger, *Technologies of Compliance: Risk and Regulation in a Digital Age*, 88 TEX. L. REV. 669 (2010) (describing private automated law systems that failed to recognize risks to bank capital reported leading into global financial crisis).

[3] Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1256 (2008) (describing the Colorado Benefits Management System, which generates welfare eligibility decisions).

[4] *See* Maayan Perel & Niva Elkin-Koren, *Accountability in Algorithmic Copyright Enforcement*, 19 STAN. TECH. L. REV. 473, 477 (2016) (describing internet service provider algorithms).

[5] *See, e.g.*, Mark A. Geistfeld, *A Roadmap for Autonomous Vehicles: State Tort Liability, Automobile Insurance, and Federal Safety Regulation*, 105 CAL. L. REV. 1611, 1624-25 (2017) (describing existing vehicle systems such as cruise control and collision prevension systems).

[6] *See* Sarah B. Lawsky, *A Logic for Statutes*, 21 FLA. TAX REV. __ (2017) (analyzing statutes as a formal system of default logic).

[7] *See* Anthony J. Casey & Anthony Niblett, *The Death of Rules and Standards*, 92 IND. L. J. __ (2017) (predicting that big-data-derived "microdirectives" will constitute law in the future).

timekeeping software omits scheduled breaks from the working time recorded for minimum-wage employees at Wal-Mart, McDonalds and so forth.[8] The software's decision means that employers will decrease workers' wages because employers will not count scheduled break time. The single software code decision could reduce the wages of tens of millions of hourly workers nationwide. Enforcing wage and hour law in this context requires government to take enforcement actions that in some way address the algorithms and computer programs through which the employers may be breaking the law.

Government-to-robot enforcement is not just a prediction for the future. It has already arrived. For instance, giant tax preparation software systems like TurboTax stand between individual taxpayers and the government. Taxpayers can buy audit insurance from TurboTax, in which case the software company not only files the taxpayer's returns, but also resolves the taxpayer's legal disputes with state and federal tax authorities.[9] Also, if the government finds errors in TurboTax programs, it reportedly works directly with TurboTax to resolve the errors. TurboTax may alert all affected users to the mistake; may pay for the users to file amended returns; and may even reimburse users for additional taxes. This all amounts to TurboTax and the government determining how the law will apply to the individual taxpayers who are TurboTax customers. It is, already, government-to-robot enforcement.

An automated, centralized law system can deliver accurate and speedy legal determinations. This can help both regulated parties and the government. The invention of TurboTax and related tax preparation systems saves taxpayers and the government billions of dollars every year, in part because taxpayers need not spend as much time doing their taxes.[10]

But no matter how well-designed an automated law system is, errors are inevitable. Prior work has recommended government oversight to reduce mistakes,[11] but this oversight cannot eliminate them. There remains the task of

---

[8] *See* Elizabeth Tippett, Charlotte S. Alexander and Zev J. Eigen, *When Timekeeping Software Undermines Compliance*, 19 YALE J. L. & TECH. 1 (2017).

[9] *See* Jay A. Soled & Kathleen DeLaney Thomas*, Regulating Tax Return Preparation*, 58 B.C. L. REV. 151 (2017).

[10] *See* Marsha Blumenthal & Joel Slemrod, *The Compliance Cost of the U.S. Individual Incoem Tax System: A Second Look After Tax Reform*, 45 NAT'L TAX J. 185, 190 (1992) (giving data re time spent on different elements of tax compliance); Rosemary Marcuss et al., *Income Taxes and Compliance Costs: How Are They Related?* 66 NAT'L TAX J. 833, 845 (2013) (reporting average out of pocket costs).

[11] *See, e.g.*, Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CAL. L. REV. 513 (2015) (recommending a "Federal Robotics Commission"); Rory Van Loo, *Rise of the Digital Regulator*, 66 DUKE L. J. 1267, 1327-28 (recommending oversight by an interdisciplinary "technology meta-agency").

identifying and correcting mistakes in automated centralized law systems.[12] Government-to-robot enforcement is well-suited for the job. This is because it can correct a mistake throughout an entire system – not just in an individual case.

For instance, assume that a Wal-Mart employee routinely works through her daily breaks, consistent with her supervisor's expectations. Her employer's timekeeping software omits this work time from her compensated time, which (let us assume) violates the law. The mistake is not small potatoes. It could shortchange the employee by a thousand dollars a year.[13] If the mistake affected equally 10 million hourly workers in the United States, it would amount to underpayment of these employees by $10 billion annually.

Under the historic, individualized law enforcement model, this legal violation might be corrected through the employee's claim of underpayment against Wal-Mart, or by the government's claim that Wal-Mart had violated wage and hour timekeeping rules.[14] A court or other forum would decide whether the employee had been underpaid. It might award her thousands of dollars in back pay, and perhaps punitive damages. But it almost certainly would not assess damages as high as $10 billion.

Under the government-to-robot enforcement model, the government would file a claim against the maker of the timekeeping software. And the remedy could assess damages for *all* employees at *all* employers who use the software – rather than only addressing the single dispute between the individual and Wal-Mart. If this problem arose for 10 million minimum wage workers, annual damages could equal $10 billion. The maker of the software would be liable for the damages, and would pass on the cost to its users in the form of higher software license fees.

Indeed, the maker of the software would already have charged its users in exchange for the software firm's assumption of the risk of violating the law. If the software took legally risky positions, it would cost more for Walmart and others to use the software. The software maker could put aside the additional user fees and self-insure by establishing a reserve fund to cover the possibility of a judgment like the $10 billion award described above. Or, the software firm

---

[12] *Cf.* Henry Smith, *Fusing the Equitable Function in Private Law* (forthcoming in PRIVATE LAW IN THE 21ST CENTURY (Kit Barker, Karen Fairweather, and Ross Grantham eds.) Harvard Public Law Working Paper No. 16-27 (arguing that errors are an inevitable feature of law and that equity can be understood as a mechanism to correct them).

[13] Two 15-minute breaks, 200 working days per year and $10 minimum wage amounts to a wage loss of $1000, or about 6% of a minimum-wage worker's salary.

[14] *See* Daniel V. Dorris, *Fair Labor Standards Act Pre-emption of State Wage-and-Hour Law Claims*, 76 U. CHI. L. REV. 1251, 1252 (2009) (noting private and public causes of action).

might purchase insurance.[15]

The biggest advantage of government-to-robot enforcement is that it can solve underdetection and underenforcement, which are classic vulnerabilities that prevent legal compliance from functioning as it should.[16] Under the individualized enforcement model, many violations will go undetected and unpunished because of resource constraints. When government does find a compliance violation, it may underpunish, because it may reduce penalties to account for intent or fault or other mitigating factors. It also generally assesses penalties based on the harm caused by the violation alone, and does not multiply penalties to account for the low likelihood of detection.[17] The result is that regulated parties do not pay for the social cost of all violations. They do not internalize the negative externalities of extra pollution, underpayment of wages, or tax avoidance.

Government-to-robot enforcement can solve these chronic problems of underdetection and underenforcement. Government-to-robot enforcement is cost-effective. It is cheaper to find and pursue a single claim against an automated centralized system which has violated the law, compared to pursuing many claims against many individuals.

Also, government-to-robot enforcement would support strict liability. A strict liability approach would not reduce penalties to account for intent or fault. This would help right-size enforcement and correctly regulate the problem of noncompliance, because it would charge a more accurate price for violations of law.[18]

A government-to-robot enforcement model also is an excellent fit for a damages multiplier, which can further address the problem of

---

[15] This centralization of enforcement and liability is the opposite of the prediction of peer-to-peer insurance and "radical financial disintermediation" suggested elsewhere. *See* Michael Abramowicz, *Cryptoinsurance*, 50 WAKE FOREST L. REV. 671, 673 (2015).

[16] *See, e.g.*, Jennifer Arlen & Reinier Kraakman, *Controlling Corporate Misconduct: An Analysis of Corporate Liability Regimes*, 72 N.Y.U. L. REV. 687, 757 (1997) (explaining that corporate employees' incentives to violate the law follow from less than 100% probability of detection and enforcement).

[17] *See, e.g.*, Michael J. Graetz & Louis L. Wilde, *The Economics of Tax Compliance:. Fact and Fantasy*, 38 NAT'L TAX J. 355, 358 (1985) (arguing that "moral and political constraints" make the idea of a damages multiplier "irrelevant as a policy matter").

[18] *See* John C.P. Goldberg & Benjamin C. Zipursky, *The Strict Liability in Fault and the Fault in Strict Liability*, 85 FORDHAM L. REV. 743, 745 (2016) (explaining that strict liability supports "licensing-based liability," distinct from "liability imposed on the basis of wrongdoing"). *See also* Alex Stein, *The Domain of Torts* 117 COLUM. L. REV. 535, 594 (2017) (explaining that strict liability supports a "public mechanism of accident regulation").

underenforcement.[19] Penalties assessed against automated, centralized robots could increase damages to reflect fully the problem of underdetection. The cost of the multiplied damages would be appropriately spread among the users of the system who engaged in the penalized, illegal activity, such as pollution, underpayment of wages, or tax avoidance. As a result, government-to-robot enforcement could fulfill the objective of forcing regulated parties to internalize the negative externalities of noncompliance.

This internalization of costs in turn could produce market differentiation that would allow users to choose their desired approach to reporting and compliance.[20] A system that encouraged or required users to pursue safer aggressive positions would cost less, because it would cost less to insure against the chance that the position turns out to be illegal. A system that took more aggressive reporting positions -- like rounding time to the hour even if workers are never late, but sometimes early -- would cost users more, because the insurance cost would be higher. The higher insurance cost for riskier reporting positions better tracks the greater likelihood that a riskier reporting position will break the law.

Because government-to-robot enforcement solves problems of underdetection and underenforcement, it can minimize tax avoidance, result in fairer payments for hourly workers, curtail illegal pollution, and improve compliance in many other areas. But government-to-robot enforcement also carries costs. Government-to-robot enforcement may encourage automated, centralized systems to work closely with the government, raising problems of capture of robots by the government, or vice versa. Individuals will be less able to assert their own legal claims against the government. And there will be winners and losers as a result of the government's improved ability to find and penalize certain kinds of noncompliance.

The problem of capture is a concern that government might be unduly influenced by the makers of centralized automated systems.[21] For instance, if a self-driving car company persuaded the government to create a presumption that self-driving cars obey speed limits, that could be an example of capture.

---

[19] *See* Gary Becker, *Crime and Punishment: An Economic Approach*, 76 J. POL. & ECON. 169 (1968) (recommending damages multiplier).

[20] *See, e.g.*, Steven Shavell, *Strict Liability Versus Negligence*, 2 J. LEG. STUD. 1, 3 (198_) (explaining that strict liability is appropriate for cases of "accidents between sellers and strangers" because if sellers are forced to pay for harm to strangers, market forces will adjust the prices charged to customers until the outcome is efficient").

[21] *See, e.g.*, James Q. Wilson, *The Politics of Regulation*, *in* SOCIAL RESPONSIBILITY AND THE BUSINESS PREDICAMENT 135 (James W. McKie ed. 1974) (noting ability of a small homogenous group to "make substantial gains by imposing unobtrusive costs on large numbers of others).

The law might favor self-driving cars because lawmakers wanted to ingratiate themselves to the makers of the self-driving cars.

When automated, centralized systems work closely with government, capture can also go in the reverse direction.[22] In other words, the robots might be unduly influenced by government. The self-driving car software engineers might decide, for instance, to prohibit self-driving cars from making right turns at red lights, even if the law allows it. The engineers' goal might be to ingratiate themselves to government regulators by making their cars extra safe. The no-turn-on-red bit of software code might not rest on an independent policy that balances the costs and benefits of, say, traffic safety as against driver time.

Another problem with government-to-robot enforcement is that as it sidelines the individualized model of legal disputes and legal development, it reduces the number of possible litigants. Initial decisions are less likely to be questioned, and the law could become less responsive and flexible as a result. There will be less diversity in legal outcomes. The development of the law may slow. For instance, consider the landmark Constitutional case of *United States v. Windsor*, which concluded that federal law could not refuse to respect a same-sex marriage that was valid under state law.[23] *Windsor* was a tax case. As government-to-robot enforcement occupies the field of tax compliance, it makes it more difficult for plaintiffs like Edith Windsor to make a claim.

Government-to-robot enforcement also raises novel questions about winners and losers. If it is easier for the government to find and penalize legal violations committed through automated, centralized law systems, who will win and who will lose? In the case of tax preparation software, it may be that the losers will be risk-seeking taxpayers, perhaps owners of capital, who use centralized software to claim aggressive tax positions. Some of the winners would be risk-averse taxpayers, perhaps wage earners, who cannot claim aggressive positions. This seems desirable from the perspective of distributive justice, if risk-seeking individuals are usually richer, and risk-averse individuals are usually poorer.[24] But other winners would be taxpayers who do not use automated, centralized systems like TurboTax at all, but rather are wealthy enough to use individualized or bespoke systems of tax preparation.

---

[22] *Cf.* Jody Freeman, *The Private Role in Public Governance*, 75 N.Y.U. L. REV. 543, 671-72 (2000) (noting that "interdependence among private and public actors" could increase government power or influence).

[23] *See* United States v. Windsor, 570 U.S. __ (2015) (holding that Windsor could claim the surviving spouse estate tax exemption under federal law upon the death of her wife and ordering the IRS to issue a tax refund).

[24] *Cf.* John Rawls, *Some Reasons for the Maximin Criterion*, 64 AM. ECON. REV. 141, 144 (1974) (arguing that "considerable normal risk-aversion" supports the maximin principle).

Part I of this Article explains the phenomenon of regulatory compliance and illustrates its challenges of legal uncertainty, underdetection and underenforcement. As Part I explains, compliance involves judgment calls that add up to big shifts in resources from one group to another. A central enforcement problem is that noncompliance often goes undetected and/or underenforced. Groups other than regulated parties – such as employees whose employers underreport wages, or general-public taxpayers who do not take aggressive reporting positions – bear the burden, while the noncompliant reap the benefits.

Part II outlines how a broader regime of direct liability for errors in automated centralized law systems could work. Part II describes a regime with strict liability and a damages multiplier, so that the dispute about an individual's compliance would be a referendum on all similar positions. It considers the problem of subrogation, meaning that the automated law system would represent its users in enforcement actions, and users might sometimes dislike the result.

Part III identifies advantages of government-to-robot enforcement: better and less costly compliance, market differentiation, and risk pricing and insurance. Part IV identifies disadvantages: capture of government by regulated parties, capture of regulated parties by government, and the reduced importance of individuals' claims. Part V discusses the problem of winners and losers, using the example of self-driving cars.

I. THE PROBLEM OF COMPLIANCE

A. Uncertainty

Automated law is widely used for so-called "compliance" purposes.[25] "Compliance" here means the act or process of following and conforming to law, including statutes, regulations, and administrative guidance.[26] It might

---

[25] *See, e.g.*, Tom C. W. Lin, *Compliance, Technology, and Modern Finance*, 11 BROOK. J. CORP. FIN. COMM. L. 159, 167-68, 178-79 (2016) (noting that "between 2012 and 2014, JP Morgan alone invested billions of dollars and added 13,000 new employees to respond to regulatory compliance requirements and explaining the "governance, risk and compliance" or "GRC" systems that "allow compliance departments to automate and analyze large volumes of information").

[26] *See* GEOFFREY PARSONS MILLER, THE LAW OF GOVERNANCE, RISK MANAGEMENT AND COMPLIANCE 2-3 (2014) ("'Compliance' refers to the processes by which an organization polices its own behavior to ensure that it conforms to applicable rules and regulations."). This differs a bit from the dictionary definition of compliance as "the act or process of complying to a desire, demand, proposal, or regimen or to coercion," https://www.merriam-

seem that compliance just involves looking up the tax treatment or the environmental reporting requirement and applying it, as one might look up an internet address and open up the right browser window. But it is not so simple.

Compliance is neither boring nor straightforward. The meaning of law is often unclear, and compliance involves judgment calls about what can be labeled, or reported, as compliant. Sometimes the uncertainty can arise from sources of law that fall outside the rules and regulations that appear to govern the conduct at issue, like the Constitution. Consider, for instance, the question of promoting a drug for "'off-label use,' that is, for a purpose not approved by the U.S. Food and Drug Administration." The applicable food and drug law criminalizes the intentional promotion of off-label use.[27] But at least one Court of Appeals has blocked a criminal proceeding involving a salesperson pushing Xyrem, also known as the "date rape drug," on First Amendment grounds.[28]

Of course, regulatory uncertainty also comes in more mundane – but perhaps no less important -- packages. Consider the question of compliance[29] with the Fair Labor Standards Act (FLSA), which requires employers to "make, keep and preserve" wage and hour records.[30] Timekeeping software is used to keep records for a large fraction of the 80 million hourly workers in the United States. Yet the repeated issue of how to calculate hourly worker's time raises legal questions as to which the answer is unclear.

One area of uncertainty has to do with how to calculate time worked under federal law when an employee arrives early or leaves late. The question is one of rounding. FLSA regulations date back to the days of paper records. They accept the practice of rounding "starting and stopping time … to the nearest quarter of an hour" so long as it does not cause "a failure to compensate the employees properly for all the time they have actually worked."[31] While this

---

webster.com/dictionary/compliance (last visited December 24, 2017). It is not far from the converse of the dictionary definition of noncompliance, the "failure or refusal to comply with something (such as a rule or regulation)." https://www.merriam-webster.com/dictionary/noncompliance (last visited December 24, 2017).

[27] 21 U.S.C. § 331(a) (prohibiting transaction in a "misbranded" drug in interstate commerce).

[28] *See* United States v. Caronia, 703 F.3d 149 (2012); *Recent Case*, 127 HARV. L. REV. 795, 800 (2013) ("Although *Caronia* is defensible as a matter of Constitutional doctrine, it is undesirable as a matter of policy.").

[29] The descriptions of errors and features of electronic timekeeping systems is based on a qualitiative empirical examination of thirteen such systems. *See* Elizabeth Tippett, Charlotte S. Alexander and Zev J. Eigen, *When Timekeeping Software Undermines Compliance*, 19 YALE J. L. & TECH. 1 (2017).

[30] 29 USC § 211(c).

[31] 29 C.F.R. § 785.48.

regulation concretely states that rounding time within seven minutes of the hour is generally acceptable, it subjects this rule to a mushy exception – the rounding will not be acceptable if it fails to "compensate the employees properly." That is the uncertain part, or the part that requires a judgment call. What if an employer's policy effectively prevents tardiness, so that employees are sometimes early, but never late, and therefore they are systematically undercompensated by the rounding rule?

Rounding seven minutes to zero minutes may seem like a small thing. Why not let the machine resolve this minor regulatory uncertainty in favor of employers? Indeed it is not unusual for specific compliance decisions to present as unimportant, small and tedious – just the sort of decision one would want a machine to take care of.

But consider this: In 2016, there were about 80 million hourly workers in the U.S., representing almost 60% of all workers paid a wage or salary.[32] Say half of those workers lose six minutes of pay each working day because of the rounding error, and that the six minutes would have been paid at $10 per hour. That amounts to the underpayment of 40 million workers at $1 per working day, or about $200 underpayment per worker per year, or $8 billion total annually. When an automated law system resolves a seemingly small uncertainty, it can add up to a big shift in resources or wealth. Compliance matters.

B. *Underdetection*

Underdetection and underenforcement also complicate compliance. Underdetection means that the government simply does not know about most violations of law. Underenforcement means that even when government knows about violations, it does not always challenge them or, when it does challenge the violations, it does not always impose full penalties.

Underdetection causes people to act as if no one is watching. Consider, for instance, the 2 million bank accounts wrongfully opened by Wells Fargo employees without customer authorization between 2011 and 2015. Why did the employees do this? Because it would increase their bonuses, and because they thought they would not get caught.[33] Or consider Volkswagen's modifications to vehicles' emission measurement systems, crafted specifically to evade environmental regulations, but create the appearance of

---

[32] *See* U.S. Bureau of Labor Statistics, Characteristics of Minimum Wage Workers 1 (April 2017) (providing statistics).

[33] *See* Emily Glazer, *Wells Fargo to Pay $185 Million Fine Over Account Openings*, WALL ST. J., Sept. 8, 2016.

noncompliance. Why did VW do this? Because it would increase their profit, and because they thought they would not get caught.[34]

Taxpayers experience this underdetection incentive as the so-called "audit lottery." In tax, the "audit lottery" simply means that the IRS usually does not check tax returns.[35] In other words, audit rarely happens. Perhaps two percent of individual returns are audited.[36] This is such a small number a taxpayer who takes a chance at the so-called "audit lottery" will often win. That is, a taxpayer who takes an aggressive position will usually get away with it, because the government will never notice.[37] Taking a chance on an aggressive reporting position is not thought of as bad behavior. It is simply how the game is played.[38]

C. *Underenforcement*

Of course, in the cases that make it to the newspapers, the gamble on underdetection did not pay off. The Wells Fargo account-opening fraud was detected. So was the Volkswagen effort to evade emissions regulations. But in these cases, the problem of underenforcement still exists.

Despite a less-than-100% probability of detection, penalty and enforcement practice could still effectively deter regulated parties and prevent noncompliance. For instance, penalties might be subject to a damages multiplier. They might equal the harm of the violation, multiplied by the inverse of the probability of detection.[39] This idea has been applied in

---

[34] *See* Danny Hakim, Aaron M. Kessler & Jack Ewing, *As Volkswagen Pushed to Be No. 1, Ambitions Fueled a Scandal*, N.Y. TIMES, Sept. 16, 2015.

[35] To clarify: Automatic computer-matching audits regularly check that figures reported by third parties like employers or banks, such as wages, interest and dividends, match the figures on individual returns. In-depth audit that looks more closely at a taxpayer's return rarely happens.

[36] The number varies somewhat based on the income level of taxpayers, so that a larger percent of higher-income individual returns are audited. But in any case, it is a small number.

[37] *See* Linda M. Beale, *Tax Advice Before the Return: The Case for Raising Standards and Denying Evidentiary Privileges*, 25 VA. TAX REV. 583, 609-10 (2006) (explaining the audit lottery and related questions such as the attorney ethics of describing the odds of getting caught).

[38] *See* Joel Slemrod, The Economics of Corporate Tax Selfishness, 57 NAT'L TAX J. 877, 882 (2004) (stating that taking a position is not about honesty or dishonesty, but rather about "rational calculat[ion] of what is in [the taxpayers'] best interest).

[39] *See* Gary Becker, *Crime and Punishment: An Economic Approach*, 76 J. POL. & ECON 169 (1968).

regulatory scholarship, for example in scholarship on tax avoidance or evasion[40] and in scholarship on corporate misconduct.[41]

But contrary to the idea of a penalty multiplier, regulated parties who commit legal violations can sometimes escape liability for reasons that are unrelated to the harm of the violation. If a regulated party can show that it did not intend a certain harm, it can sometimes escape liability.[42] For example, a class action lawsuit against Wells Fargo in the unauthorized account scandal must show that the board negligently or intentionally allowed the fraud to continue.[43]

If a regulated party has certain procedures in place, it may also escape liability in some circumstances.[44] The securities law includes a safe harbor that protects against supervisory liability for employees who commit legal violations if certain institutional procedures are in place.[45] Some suggest that the existence of a corporate compliance program should be a mitigating factor more generally that should reduce enforcement or reduce applicable penalties.[46]

Law also sometimes assigns different penalties to a legal violation depending on whether there was a reasonable ex ante argument that the action in question did not violate the law. Rules governing so-called "reporting positions" in tax articulate this view. The idea of tax reporting positions is to

---

[40] *See, e.g.*, Michael G. Allingham & Agnar Sandmo, *Income Tax Evasion: A Theoretical Analysis*, 1 J. PUB. ECON. 323 (1972).

[41] *See, e.g.*, Jennifer Arlen & Reinier Kraakman, *Controlling Corporate Misconduct: An Analysis of Corporate Liability Regimes*, 72 N.Y.U. L. REV. 687, 757 (1997) (explaining that corporate employees' incentives to violate the law follow from less than 100% probability of detection and enforcement).

[42] The question of what an institutional regulated party "intends" is tricky. *Compare* United States v. Bank of New England, 821 F. 2d 844 (1st Cir 1987) (finding "collective intent" where head tellers at bank allowed transactions apparently structured to avoid currency transaction reporting requirements) *with* United States v. SAIC, 626 F.3d 1257, 1274 (D.C. Cir 2010) (holding that "collective knowledge" would not support knowledge under False Claims Act, 31 U.S.C. § 3729(b)(1)(A)).

[43] *See* Jon Hill, *Judge Lets Most of Wells Fargo Derivative Suit Proceed*, Law360, Oct. 5, 2017, *available at* https://0-www-law360-com.tallons.law.utexas.edu/articles/971688; Evan Weinberger, *Judge Allows Suit Against Wells Fargo's Board to Press On*, Law360, May 5, 2017, available at https://0-www-law360-com.tallons.law.utexas.edu/articles/920791.

[44] *Cf.* Geoffrey Miller, *Catastrophic Financial Failures: Enron and More* 89 CORNELL L. REV. 423 (2004) (noting that Enron and other companies had textbook corporate governance systems on paper, but catastrophic failure in reality).

[45] *See* 15 U.S.C. § 78o(b)(4)(E) (describing "procedures … which would reasonably be expected to prevent and detect, insofar as practicable, any such violation by such other person").

[46] *See, e.g.* Ellen S. Podgor, *A New Corporate World Mandates a Good Faith Affirmative Defense*, 44 AM. CRIM. L. REV. 1537 (2007).

attribute increased liability to taxpayers and tax preparers in the case of more aggressive tax positions. Taxpayer "substantial understatement" penalties, for instance, equal up to 20% of underreported liability.[47] Tax preparer liability can include penalties such as a prohibition on future practice before the IRS, imposed by an IRS ethics office.[48] But these extra penalties do not apply to a position for which there is "substantial authority," meanings that the balance of the authorities support the position.[49]

In other words, the tax law imposes extra penalties on positions that taxpayers and tax advisers should know are wrong ex ante. But it does not impose extra penalties on plausible positions that turn out to be wrong when later tested. If the law is uncertain, so that taxpayers and tax advisers could reasonably have believed the position could be correct ex ante, then penalties do not apply.

D. *Law Creep: A Product of Underdetection and Underenforcement*

Underdetection and underenforcement encourage regulated parties to take positions that they think will not be detected, and/or that they think will not be punished too harshly if they are detected. Sometimes this pushes the development of the law too far in favor of reporting parties. Frequent flyer miles and loyalty programs present one example of a tax reporting position pushed to an extreme.

When an employee receives loyalty points on a purchase made for work, those points should be taxable as income under generally accepted income tax principles.[50] But employees (and employers) rely on a 2002 Announcement stating that the IRS will not press the issue on benefits "attributable to … business travel."[51] Everyone takes the position that loyalty points received in connection with employment are never taxable – even though such programs have spread far beyond travel, even though that many employees may derive

---

[47] I.R.C. § 6662.

[48] I.R.C. § 6694; Circular 230.

[49] Treas. Reg. § 1.6662-4(d)(3) (outlining approach to determining whether substantial authority is present). Extra penalties also do not apply to a weaker "reasonable basis" position that is disclosed.

[50] *See* Dominic L. Daher, *The Proposed Federal Taxation of Frequent-Flyer Miles Received From Employers: Good Tax Policy But Bad Politics*, 16 AKRON TAX J. 1 (2001) (detailing 1995 IRS attempt to tax miles, court cases dealing with miles, and technical doctrinal pathways to the taxation of miles).

[51] IRS Announcement 2002-18 (stating that "the IRS will not assert that any taxpayer has understated his federal tax liability by reason of the receipt or personal use of frequent flyer miles or other in-kind promotional benefits attributable to the taxpayer's business or official travel").

substantial value from them, even though points often can be used to purchase a broad range of consumer goods, and even though the sale of points can produce taxable income.[52] There is little chance that the government would have the stomach to challenge taxpayers' reporting positions on loyalty points now. The law on loyalty points has crept too far away from the theoretically correct answer of inclusion in income.

Other agencies with enforcement responsibilities that exceed their grasp likely fall into a similar pattern. Violations go undetected because resource constraints prevent the agency from auditing everyone. The agency is likely to treat gently a position based on a plausible legal interpretation. Law creep happens, meaning that aggressive positions are embraced by the market and accepted by the administrative agency. The result is that many transgressions go uncorrected, to the benefit of the regulated party that claims the benefit – i.e., the person that reported a lower tax liability, or smaller pollution emissions, or greater amount of Tier One capital, than the law really requires.

E. Negative Externalities

Inappropriate and uncorrected benefits claimed by regulated parties who take aggressive reporting positions raise the question of who bears the burden, or detriment, of these changes.[53] The answer is that third parties bear the burden of undetected errors.[54] It is a problem of negative externalities. Underreporting taxable income for a user taxpayer – like omitting the value of loyalty points earned on employment-related purchases -- results in higher taxes for other taxpayers.[55] Wage underreporting undercompensates hourly employees. Underreporting environmental emissions produces more pollution for the general public and the environment.

 To correct the problem of negative externalities, regulated parties who take aggressive reporting positions ought to be forced to internalize them. This is a classic application of the "single owner" principle.[56] If the party that

---

[52] Charley v. Commissioner, 91 F.3d 72 (9th Cir. 1996).

[53] This assumes that the law is correct, i.e., that error-free compliance would properly measure time worked, impose tax liability, set environmental pollutants and so forth. Analyzing the normative correctness of any particular existing regulatory regime falls outside the scope of this article.

[54] *Cf.* Oren Bracha & Frank Pasquale, *Federal Search Commission? Access, Fairness, and Accountability in the Law of Search*, 93 CORNELL L. REV. 1149, 1185-86 (2008) (noting the problem of possible negative externalities resulting from the filtering and organizing of search results by algorithms).

[55] *See* Linda Sugin, *Invisible Taxpayers*, 69 TAX L. REV. 617 (2016).

[56] *See, e.g.*, Richard A. Epstein, *Holdouts, Externalities and the Single Owner: One More Salute to Ronald Coase*, 36 J.L. & ECON. 553 (1993).

benefits from an action (like pollution) also must pay the costs of that action (like polluting in violation of environmental regulations) then an efficient amount of pollution – or at least, the efficient amount of pollution according to the regulations – should result.[57]

The development of automated, centralized law systems presents the perfect opportunity to force regulated parties to internalize such negative externalities. They could allow a solution to the pervasive problem of underdetection and underenforcement. This is the enforcement opportunity which is the subject of the next Part of this Article.

II. AN ENFORCEMENT OPPORTUNITY

A. *Automated Centralized Law Systems Make Mistakes*

A machine operates as an automated or centralized law system (for purposes of the analysis here) if it produces a legal determination.[58] For instance, an automated law system might produce the legal determination that a party is compliant with a regulatory requirement.[59] Automated law systems produce wage and hour records,[60] tax returns,[61] environmental reports.[62] They respond to copyright-based takedown requests.[63] These compliance cases are

---

[57] *See generally* WARD FARNSWORTH, THE LEGAL ANALYST 37-46 (2007) (describing the single owner principle).

[58] Securities law compliance software appears to still focus on editing and compilation. In other words, it has not crossed the line into automated law territory yet. A securities law examdple is the NovaWorks GoFiler suite of products. *See* http://www.novaworkssoftware.com/index.php?page=agents/gofilercomplete.html (last visited August 14, 2017) (promising superior editing capabilities for XBRL, which is the format in which many EDGAR filings are made).

[59] Rather than the law responding to the existence of a machine or other technology, the goal is to determine how machines can help the enforcement and making of law. *Compare, e.g.* Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CAL. L. REV. 513 (2015) (recommending a "Federal Robotics Commission");

[60] *See* Elizabeth Tippett, Charlotte S. Alexander and Zev J. Eigen, *When Timekeeping Software Undermines Compliance*, 19 YALE J. L. & TECH. 1 (2017) (providing qualitiative study of 13 different timekeeping automated law systems).

[61] *See* Rodney P. Mock & Nancy Shurtz, *The TurboTax Defense*, 15 FLA. TAX REV. 443 (2014) (describing tax preparation software).

[62] Perillon Software Inc. provides an environmental law compliance automated law example. *See* http://www.perillon.com/environmental-data (last visited August 14, 2017) (stating, for instance, that "[o]ur customers use our environmental data management module for GHG MRR Reporting for Subpart A, C, D, W reporting requirements including evolving electronic submission standards (e.g. e-GGRT)"). e-GGRT is the EPA's reporting system for greenhouse gases. See https://www.epa.gov/ghgreporting/e-ggrt-news (last visited Aug. 14, 2017).

[63] *See* Maayan Perel & Niva Elkin-Koren, *Accountability in Algorithmic Copyright Enforcement*, 19 STAN. TECH. L. REV. 473, 477 (2016) ("[M]ajor online intermediaries use algorithms to

the focus of the analysis.[64]

An increasing variety of technologies support automated law. These include logical algorithms,[65] machine learning and other artificial intelligence techniques,[66] and computer network approaches such as blockchain, or distributed ledger, systems.[67] All of these technologies are fallible.

Recommendations to improve automated, centralized law systems on an ex ante basis are sensible, but insufficient. This Article is interested in the ex post question. That is, how should the law respond to inevitable mistakes?

---

filter, block, and disable access to allegedly infringing content automatically, with little or no human involvement."). These are often responses to robot-generated Digital Millenium Copyright Act takedown requests. In contrast to most of the other automated law systems considered here, they appear to be proprietary, i.e. developed, owned and used by a firm such as Google or Facebook.

[64] Automated law also includes private law examples, which are beyond this Article's scope. Ethereum project; Kevin D. Werbach & Nicholas Cornell, *Contract Ex Machina*, 67 DUKE L.J. 313 (2017). Smart contracts charge computers with the responsibility of verifying the fulfillment of contract terms, such as delivery of goods, and with the responsibility of executing contract terms, such as transferring funds in payment. *See also* Kiviat, Note, DUKE L.J. 2015 (re digital asset transfers, confirmation of authorship, title transfers, K enforcement)

[65] E.g. Maayan Perel & Niva Elkin-Koren, *Accountability in Algorithmic Copyright Enforcement*, 19 STAN. TECH. L. REV. 472, 477, 488-91 (2016) (explaining that "platforms[] such as Google, Facebook, and Twitter … appl[y] various algorithms to perform qualitative determinations, including the discretion-based assessments of copyright infringement and fair use" in order to respond to robot-generated takedown requests by copyright owners and suggesting that this results in over-enforcement of copyright rights).

[66] *See* Benjamin Alarie, Anthony Niblett and Albert Yoon, *Using Machine Learning to Predict Outcomes in Tax Law* (October 16, 2016), available at https://ssrn.com/abstract=2855977 or http://dx.doi.org/10.2139/ssrn.2855977 (describing AI technique applied to database consisting of the text of hundreds of cases to give answer re: whether worker is an employee or an independent contractor for tax purposes).

[67] The use of distributed ledger or blockchain technology, which also supports the bitcoin currency, has been proposed for use by different computers in several legal capacities. For instance, computers in different jurisdictions might agree on the status of an import/export transaction. *See* Richard T. Ainsworth and Musaad Alwohaibi, *Blockchain, Bitcoin, and VAT in the GCC: The Missing Trader Example* (2017 working paper) (describing blockchain-based information confirmation system proposed in for new VAT system in Middle East Gulf Cooperation Council trading bloc). Blockchain technology might confirm and effect international payments. *See* Marcel T. Rosner & Andrew Kang, Note, *Understanding and Regulating Twenty-First Century Payment Systems: The Ripple Case Study*, 114 MICH. L. REV. 649, 651 (2016) (suggesting that the Federal Reserve would have an interest in this regulatory solution). A proposed system based on blockchain has been built to reduce the cost of administering so-called know-your-customer regulations relevant to anti-money laundering and anti-tax evasion laws. *See* Jose Parra-Moyano & Omri Ross, *KYC Optimization Using Distributed Ledger Technology* (avail on SSRN 2017 paper). *See generally* Carla L. Reyes, *Conceptualizing Cryptolaw*, 96 NEB. L. REV. __ (2017).

Documented errors in automated law systems range from the mundane, such as a miscalculation of depreciation;[68] to the heartbreaking, such as an erroneous denial of food stamp benefits;[69] to the macroeconomic, such as a failure to correctly recognize risks to bank capital on the eve of the global financial crisis.[70]

There are several reasons for error. One is that humans design and build automated law, and people make mistakes. Another reason is that automated law systems only have access to existing, or past-developed, information. The idea that technology can automatically determine whether a worker is an employee or an independent contractor, for instance, depends on the technology's access to a database of worker status determinations.[71] Historical data cannot reliably predict answers in some new situations. The issue of worker classification in the gig economy provides an example of a novel new set of facts.[72]

System designers also have an incentive to favor regulated parties who purchase and use their system. For instance, an automated law system may be intentionally designed to avoid law, or support aggressive reporting positions, or find loopholes. There is an incentive to "redesig[n] behavior for legal advantage."[73] Consider the tax preparation software feature that constantly updates a taxpayer on the status of his or her payment due or refund.[74] Surely this encourages taxpayers to input larger deductions and smaller income items. Automated systems could also favor regulated parties by discovering new

---

[68] *See* Mock & Shurtz at 463 (describing TurboTax errors in 1994 and 1996). *See also* Choe v. Commissioner, T.C. Summ. Op. 2008-90, 2008 WL 2852249, at *1 (July 24, 2008); Rev. Rul. 85-187 (each involving erroneous software depreciation calculations).

[69] Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1256 (2008) (describing a state government automated law system that incorrectly denied benefits to eligible welfare recipient).

[70] Kenneth A. Bamberger, *Technologies of Compliance: Risk and Regulation in a Digital Age*, 88 TEX. L. REV. 669 (2010) (describing private automated law systems that failed to recognize risks to bank capital reported leading into global financial crisis).

[71] *See* Anthony J. Casey & Anthony Niblett, *The Death of Rules and Standards*, 92 IND. L. J. __ (2017) (giving driving example to illustrate how technology might generate "microdirectives").

[72] See, e.g., O'Connor v. Uber Techs., Inc., 82 F. Supp. 3d 1133 (N.D. Cal. 2015) (allowing trial to proceed on classification issue); Cotter v. Lyft, Inc., 60 F. Supp. 3d 1067 (N.D. Cal. 2015) (same). *See generally* Shuyi Oei & Diane Ring, *Can Sharing Be Taxed?* 93 WASH. U. L. REV. 989 (2016) (considering regulatory issues presented by the sharing economy).

[73] Tim Wu, *When Code Isn't Law*, 89 VA. L. REV. 679, 707-08 (2003). Wu describes a "code designer act[ing] like a tax lawyer … look[ing] for loopholes or ambiguities in the operation of the law. *Id.* at 708.

[74] *See* Jay A. Soled & Kathleen DeLaney Thomas*, Regulating Tax Return Preparation*, 58 B.C. L. REV. 151, 200-01 (2017) (recommending prohibition of the "prepayment-position status bar").

evasion strategies by examining primary sources of law.[75]

The tension between the architecture of computer systems and the goals of law or democracy has been explored before. It is at the core of Larry Lessig's work on cyberspace.[76] More specific shortcomings of particular centralized, automated systems have also been explored.[77] Often the recommended remedy is oversight, for instance in the form of a technology agency that regulates the content of automated systems.[78]

Still, errors will exist.[79] Who should bear the responsibility for them? If the systems themselves, and through the systems the users, pay for errors, then the underdetection and underenforcement problems that betray the promise of regulatory schemes will be solved. This is the promise of government-to-robot enforcement.

B. *Automated Systems are at the Center, But Have No Formal Place*

The idea that automated law systems can have liability for their errors would radically change the way in which regulatory law is enforced. Under current law, the government pursues enforcement directly against a regulated party. If an automated law system were directly liable, the locus of disputes would change. The adversarial parties who would develop the future course of the law would be the government, on the one hand; and the firms that make and sell automated law systems, on the other hand.

Another way of putting this is as follows. As a de jure matter, the tax software firm has a contract relationship with the taxpayer and no relationship with the government. As a de facto matter, the tax software firm sits directly in

---

[75] *See* Marcos Pertierra, Sarah Lawsky, Erik Hemberg and Una-May O'Reilly, *Towards Formalizing Statute Law as Default Logic through Automatic Semantic Parsing* (2017 working paper).

[76] *See generally* LAWRENCE LESSIG, CODE AND OTHER LAWS OF CYBERSPACE (1999) (arguing that democratic mechanisms should oversee and edit the "architecture" of cyberspace).

[77] *See, e.g.*, Kenneth A. Bamberger, *Technologies of Compliance: Risk and Regulation in a Digital Age*, 88 TEX. L. REV. 669, 729-30 (2010) (recommending "dynamic model of regulation" to improve private automated law systems); Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1256 (2008) (considering "a reconceived *Mathews* test [that] might permit hearings on flaws in [government] software, and recommending that agencies test and allow public comment on automated law software").

[78] *See, e.g.*, Rory Van Loo, *Rise of the Digital Regulator*, 66 DUKE L. J. 1267, 1327-28 (recommending oversight by an interdisciplinary "technology meta-agency").

[79] *See ,e.g.,* Henry Smith, *Fusing the Equitable Function in Private Law* (forthcoming in PRIVATE LAW IN THE 21ST CENTURY (Kit Barker, Karen Fairweather, and Ross Grantham eds.) Harvard Public Law Working Paper No. 16-27 (arguing that errors are an inevitable feature of law and that equity can be understood as a mechanism to correct them).

between the taxpayer and the government. Government-to-robot enforcement would take advantage of the practical place that tax software firms occupy between taxpayers and the government, by imposing legal liability for error directly on the firms rather than on the taxpayers.

For instance, assume that a maker of an automated law tax preparation program -- like Intuit – sells a program – like TurboTax – to a taxpayer. As a formal legal matter, TurboTax has no relationship with the tax system. It has only a contract with the taxpayer. The contract gives the taxpayer the right to use the tax preparation software in exchange for a fee.

The tax software contract does not change the formal relationship between the taxpayer and the government. The taxpayer remains directly liable for any errors on the tax return. The taxpayer's responsibility for returns prepared with software applies whether the law arises from an error made by the software program, an error made because the taxpayer submitted incorrect facts, or a combination of the two.

There is a thus a wide gap between the de jure exclusion of TurboTax from any formal place in the tax system and the de facto centralization of tax return preparation within the software systems of TurboTax and similar providers. As a formal matter, TurboTax does not even have the status of an advisor, which would allow its advice to support a reporting position; or the status of a tax preparer, which would give it liability if it gave extremely bad advice.[80] In stark contrast, as a practical matter, TurboTax makes legal decisions that determine how 33 million federal income tax returns every year will be filed.

TurboTax is involved with the tax system not just because it files 33 million federal income tax returns, but also because it is an avid consumer of government guidance. There is close relationship between the writing of tax administrative guidance and the writing of tax software programs' code. TurboTax, for instance, directly imports government forms and instructions into its software. It is widely reported that TurboTax confers with the government in informal meetings to check the validity of its software on an annual basis. Those in charge of writing tax laws in Congress readily admit that the audience for their work features two important groups: the IRS employees who draft forms and instructions, and the tax software developers

---

[80] *See, e.g.*, Rodney P. Mock & Nancy E. Shurtz, *The TurboTax Defense*, 15 FLA. TAX REV. 443, 490-505 (2014) (explaining case law that declines to waive penalties based on the "TurboTax defense" that the software facilitated an error and declines to treat software companies as possibly liable tax preparers). [But see Merrill Lynch settlement case]

who translate forms and instructions into computer code.[81]

Private contracts have begun to bridge the gap between the formal exclusion of tax software programs from the tax system and the de facto place that tax software programs hold at the center of the tax system. Audit insurance is one development.[82] Under audit insurance, a taxpayer pays a tax software system an extra fee in exchange for the software company's promise that it will manage any audit that arises and pay applicable taxes and penalties that arise from errors made by the software company.

It also appears that tax software companies have chosen voluntarily to accept the role of fixing and paying for some errors, even if no audit insurance policy requires them to do so. Informal conversations suggest that if a systematic error is discovered in a tax software program, the government may inform the tax software company. Then, the software company may initiate a solution that involves persuading users to file amended returns and paying the costs of the amended filings, including additional tax liability. A quiet settlement of the issue along these lines protects the reputational interests of the company – and puts relevant legal decisions in the hands of the software firm and the government.

The market developments of audit insurance and quiet assumption of liability are not surprising. They are consistent with the incentives presented by a centralized automated compliance system like TurboTax. It is more efficient for a centralized system to assume risk and fix mistakes compared to requiring millions of software users to individually assume risk and fix mistakes. Intuit is larger and less risk averse than taxpayers, so it is cheaper for Intuit to assume risk. Intuit has better resources and benefits from economies of scale, so it is cheaper for Intuit to fix mistakes. In the case of TurboTax, private ordering is well on its way to producing government-to-robot enforcement.

C. *Subrogation*

What if government designed a system of government-to-robot enforcement, rather than leaving the task to private ordering? An automated,

---

[81] *See* Shuyi Oei & Leigh Osofsky, *A Political Economy of Tax Exceptions* (working paper 2017)

[82] *See* Jay A. Soled & Kathleen DeLaney Thomas*, Regulating Tax Return Preparation*, 58 B.C. L. REV. 151, 180 (2017) (describing tax preparation software audit insurance); Rodney P. Mock & Nancy E. Shurtz, *The TurboTax Defense*, 15 FLA. TAX REV. 443, 492-94 & nn. 272-74 (2014) (noting tax preparation software "limited guarantees" of accurate calculations and advice including carveouts for errors due to taxpayer inputs of incorrect information or incorrect classification of information).

centralized law system has the potential to revolutionize regulators' ability to enforce the law.[83] It might work as follows:

1. Contract or regulatory law assigns liability to the maker of an automated, centralized system for compliance errors made by its users, perhaps excepting those proven by the firm to result from incorrect facts given by the user.

2. Administrative agency alleges error, and notifies user and system.

3. Subrogation: System controls dispute, including decisions about settlement, appeal etc.[84]

4. Strict liability: If penalties result from the controversy, the system pays. Issues about the truth of facts input by user resolved between system and user.

5. Damages multiplier: The automated centralized system also pays an additional amount determined by a damages multiplier. The idea is that the firm is settling not only this user's case, but also the liabilities of other users to the extent they arise from the same error.

6. Preclusion: The decision would apply to some group of filings prepared by the automated law system. Perhaps all filings with this particular issue in a particular year.

This approach fundamentally changes the way in which the law operates. No longer is the regulated party-government relationship the central or key compliance relationship. Rather, the regulated party, also the user of an

---

[83] Other have recognized the legal design opportunity presented by centralized machine gatekeepers. *Cf.* Susan Klein & Crystal Flinn, *Social Media Compliance Programs and the War Against Terrorism*, 8 HARV. NAT'L SEC. L. J. 53, 57 (2017) (recommending "criminalizing the failure of social media programs to institute policies that discover [and report] terrorism-related posts"); Yesha Yadav, *The Failure of Liability in Modern Markets*, 102 VA. L. REV. 1031, 1039-40 (considering strict liability and other regimes for harms generated by high-frequency algorithmic trading).

[84] An interesting model for this approach is the TEFRA partnership audit system, under which a partnership's "tax matters partner" is authorized to resolve matters relating to the tax treatment of partnership items with the government, and to bind certain partners to that treatment. *See, e.g.*, DAVID M. RICHARDSON, JEROME BORISON & STEVE JOHNSON, CIVIL TAX PROCEDURE 161-66 (2d ed. 2008) (explaining the powers of the tax matters partner, including the right to seek judicial review and to make a settlement agreement with the IRS that is binding on partners who "have less than a 1% profits interest in a partnership with more than 100 partners").

automated, centralized law system, agrees that the automated law system will have the primary relationship with the government. As a result, the system will be the decisionmaker in terms of what positions to take, which to defend when challenged, how to settle them and so forth. Controversy practice between the firms that offer software and the government becomes a primary avenue for the development of the law.[85]

The automated law system, in other words, would have the right of subrogation. As the party solely liable for legal error, it could step into the shoes of the user to litigate the question of legal error in the taxpayer's case. One issue this raises is a possible conflict of interest between the taxpayer and the system with respect to whether liability proceeded from an error of fact or an error of law. Issues of confidentiality and privacy for users are also raised by subrogation. Government-to-robot enforcement might try to address these problems by shifting them to the automated law system.

D. *Strict Liability*

Sometimes the choice between strict liability and negligence is framed as a choice between no-fault regulation and the fault-based concept of holding a defendant accountable for a wrong. Under this framework, automated law presents a classic case for strict liability.[86] The automated law systems considered here cover matters of public regulation. The action of failing to comply with a regulation is not the kind of direct or targeted harm done one person by another that motivates a wrongs-based system of liability.

This also means that intent or malice is not relevant for automated law liability. Automated law liability does not mean to achieve corrective justice. It means to properly regulate. Success means a bureaucratic exercise that shifts

---

[85] This centralization of enforcement and liability is the opposite of the prediction of peer-to-peer insurance and "radical financial disintermediation" suggested elsewhere. *See* Michael Abramowicz, *Cryptoinsurance*, 50 WAKE FOREST L. REV. 671, 673 (2015). Insurance questions are raised by making automated law systems liable for compliance violations. An automated law system is likely to be better able to pay a judgment compared to an individual user, which is another reason why a damages multiplier is a good fit for such a system. *Cf.* GUIDO CALABRESI, THE COSTS OF ACCIDENTS: A LEGAL AND ECONOMIC ANALYSIS 50-54 (1970) (identifying deep pockets as a possible reason supporting enterprise liability). But a regulating agency might require evidence of the system's creditworthiness before allowing the automated law system to prepare compliance submissions or other legal determinations or filings, and self-insurance, bonding or reinsurance markets might emerge to support the good credit of automated law systems. automated law systems.

[86] *See, e.g.*, Steven Shavell, *Strict Liability Versus Negligence*, [2] J. LEG. STUD. 1, 3 (198_) (explaining that strict liability is appropriate for cases of "accidents between sellers and strangers" because if sellers are forced to pay for harm to strangers, market forces will adjust the prices charged to customers until the outcome is efficient").

costs of error until they fall on the right party – the user who enjoys the benefits of the aggressive legal position. This is consistent with strict liability.[87]

Automated law systems are involved in at least three different kinds of errors: clear errors of law, unclear errors of law, and mixed errors of fact and law.   The systems can be said in some sense to cause each of these kinds of errors. Bringing all of these kinds of errors into government-to-robot enforcement through strict liability would maximize the ability of automated law liability to price, surface and debate questions of law.

To illustrate the error types, let us return to timekeeping software products.[88] These are built to comply with the Fair Labor Standards Act (FLSA), which requires employers to "make, keep and preserve" wage and hour records.[89] They are used to keep records for a large fraction of the 80 million hourly workers in the United States.

Timekeeping software might make clear errors of law. Let us assume that state law unambiguously states that break and/or meal times count toward paid time. The software's mistake on this front would be a clear legal error.

A clear legal error is the kind of error that an advisor, like a lawyer, might be liable for under a malpractice theory or a regulatory gatekeeper scheme. But note that the bounds of a lawyer's liability need not limit the liability of an automated law system. Just as a legal self-help book does not amount to the practice of law, so too the use of a software program probably does not create such a relationship.[90] Perhaps the relationship between automated law systems and users is more similar to the relationship between a bank and its customer. Banks have strict liability, for instance, for certain failures to withhold taxes.[91]

Timekeeping software might also make unclear errors of law. Consider the software's interpretation of the time rounding rule. FLSA regulations accept

---

[87] *See, e.g.*, John C.P. Goldberg & Benjamin C. Zipursky, *The Strict Liability in Fault and the Fault in Strict Liability*, 85 FORDHAM L. REV. 743, 745 (2016) (contrasting licensing-based liability and wrongs-based liability); Alex Stein, *The Domain of Torts* 117 COLUM. L. REV. 535, 594 (2017) (contrasting "public mechanism of accident regulation" and wrongs-based torts laws).

[88] The descriptions of errors and features of electronic timekeeping systems is based on a qualitiative empirical examination of thirteen such systems. *See* Elizabeth Tippett, Charlotte S. Alexander and Zev J. Eigen, *When Timekeeping Software Undermines Compliance*, 19 YALE J. L. & TECH. 1 (2017).

[89] 29 USC § 211(c).

[90] *See* ABA, Task force on the Model Definition of the Practice of Law, Sept. 18 2002 draft ("[C]ourts have held that the publication of legal self-help books is not the practice of law.").

[91] IRC 1441 et seq

the practice of rounding "starting and stopping time … to the nearest quarter of an hour" so long as it does not cause "a failure to compensate the employees properly for all the time they have actually worked."[92] Timekeeping software apparently implements this guidance with a default setting that rounds time to the hour if a punch-in or punch-out time is within seven minutes of an hour.[93] But if employer rules effectively prevent tardiness, so that employees are sometimes early, but never late, then the software's rounding default may systematically reduce the time recorded for an employee. In this case, the software's default rounding rule encourages an employer to take an aggressive, but not clearly illegal, filing position.

The liability of systems for such unclear legal errors, or incorrect judgment calls in grey areas of law, should be strict. That is, liability should not be limited to liability for a negligent or clear error of law, like the failure to research wage and hour law in a particular state. Instead, it should include liability for the close case that happens to come out in favor of the government and to the detriment of all the users who took the position. If a court invalidates the practice of rounding hourly workers' time according to the 7-minute rule, the automated law system should bear that liability even though it was not clear when the return was filed that the 7-minute rounding rule was illegal.

The reason for strict liability goes to the heart of the idea of government-to-robot enforcement. The idea is that legal questions will be priced, and then debated and decided. The interesting questions, those in need of development, are the close ones. This centralized mechanism of discovering and discussing these questions will be of much less use unless it covers these matters. Also, the system controls and makes these legal decisions as much as it makes the decisions that involve clearer legal error. It is still the least cost avoider.[94]

Finally, mixed errors of law and fact may occur in timekeeping software systems. For instance, the software may prompt employers to enter scheduled break and/or meal times for employees, and then automatically deduct that time from paid time. This connects with a legal error if some state laws do not allow break and/or meal times to be deducted from paid time. It connects with a factual error if an employer enters the wrong information. An employer's incorrect data entry usually would seem to be the employer's fault, not the system's fault. But even in this case one can find mixed questions of

---

[92] 29 C.F.R. § 785.48.

[93] *See* Tippett, Alexander & Eigen at 37 ("A common unit of rounding appears to be seven minutes.")

[94] *See* RONALD COASE, THE PROBLEM OF SOCIAL COST (1960).

fact and law. What if the system makes it very hard to change entered time if it turns out that an employee works through a break? This could be cast as an error of law, perhaps as a legal error because the design of the system so strongly suggests that scheduled time worked, not actual time worked, is the relevant input.

The question of whether systems should be liable for mixed questions of law and fact is difficult. It raises the question *of whether the system is liable for design choices that nudge a user to present facts in a certain way. An important premise here is that software programs have the capacity to manipulate or influence human users' responses through their design, as when a tax software system's "refund due" feature encourages a taxpayer to input larger deductions.[95] Including the mixed question of fact and law in the direct liability space encourages the system to stop nudging users toward noncompliance by suggesting an inappropriate legal framework for the relevant facts. But the fact that users also control facts suggests that a system should not be irrevocably liable for mistakes of fact. It ought to be able to shift liability to a user upon proving that the user provided false information.

E. Damages Multiplier

A key piece of the idea of government-to-robot enforcement is the damages multiplier. The reason for the multiplier is that it better accomplishes the goal of requiring internalization of the negative externalities of noncompliance. In other words, penalties should increase according to a damages multiplier designed to account for the error costs incurred across the system, not just for the user whose specific case is discovered.

A damages multiplier is a well-known tool suggested by literature including Gary Becker's foundational economic model of crime.[96] Say a person decides whether to comply by comparing the cost of compliance ("c") with the cost of noncompliance ("nc") multiplied by the probability of detection and liability ("p"). The person considers whether $c < nc * p$. Compliance will be the attractive answer only if nc (in other words, the penalty in the case of noncompliance) is greater than c by a factor of more than $1/p$. In other words, the damages multiplier should be more than $1/p$.

---

[95] *See* Jay A. Soled & Kathleen DeLaney Thomas*, Regulating Tax Return Preparation*, 58 B.C. L. REV. 151, 200-01 (2017) (recommending prohibition of the "prepayment-position status bar").

[96] *See* Gary Becker, *Crime and Punishment: An Economic Approach*, 76 J. POL. & ECON 169 (1968); *see, e.g.*, Michael G. Allingham & Agnar Sandmo, *Income Tax Evasion: A Theoretical Analysis*, 1 J. PUB. ECON. 323 (1972).

There are a number of issues with damages multipliers. But these issues are less problematic for automated law systems. One challenge is that political and rule of law proportionality constraints limit the ability to vastly increase penalties imposed on a single person based on the idea that her transgression was difficult to detect.[97] A second consideration is that a fixed damages multiplier across different offenses fails to account for the variation in probability of detection and in particular for the likelihood that more serious offenses are more likely to be detected.[98] A third issue is that factors other than the cost of compliance influence the magnitude of penalties. These include aggressiveness, culpability and intent.[99] They also include whether the defendant has deep enough pockets to pay the larger penalty.

A damages multiplier for automated law system can sidestep each of these issues.

First, the imposition of the penalty on the centralized system, not the individual violation, reframes the issue of ensuring that the punishment fits the crime. The idea is that the centralized system itself has the responsibility to correctly state the law or pay appropriate damages. The individual user's penalty is only the starting point for measuring the system's total error.

The idea of a damages multiplier based on other users of the system means a fundamentally different method for adjudicating claims of legal error in compliance systems. Now the litigants, effectively, are the firm that makes the automated system and the U.S. government. In addition, this case is no longer just about the audited user. Instead it is a general test case that automatically will determine liability for dozens or hundreds or thousands of returns, not just the user's.

Second, the damages multiplier can be customized. As a starting point, the damages multiplier, calculated as $1/p$, might be based on p equal to a typical audit rate. But an automated law system could prove out of the high damages by presenting its own data to rebut the calculation of the penalty multiplier.

---

[97] *See, e.g.*, Michael J. Graetz & Louis L. Wilde, *The Economics of Tax Compliance:. Fact and Fantasy*, 38 NAT'L TAX J. 355, 358 (1985) ("That an economic model analyzing the expected utility calculation of a would-be tax evader recommends large increases in the applicable sanction in light of the very low probability of its application quickly becomes irrelevant as a policy matter. In this country, at least, legal, moral and political constraints make this necessarily so. Coherence in our criminal law generally demands that 'punishment fit the crime'….").

[98] *See* Richard Craswell, *Deterrence and Damages: The Multiplier Principle and its Alternatives*, 97 MICH. L. REV. 2185, 2192 (1999).

[99] *See, e.g.*, Alex Raskolnikov, *Six Degrees of Graduation: Law and Economics of Variable Sanctions*, 43 FLA. STATE L. REV. 1015 (2016).

Customization of the damages multiplier is available in this case as for relatively few others,[100] because the information about all similar cases should be within the reach of the automated law provider. Different customized factors might be taken into account in particular cases. As an example, if an issue splits circuits, thus raising the *Golsen*[101] rule, a tax software provider should be allowed to argue for damages to be based only on the tax returns filed in such-and-such circuit.

Third, the penalty itself (aside from the multiplier) imposed on automated law systems could be set without reference to aggressiveness or culpability. The automated law liability idea does not mean to use damages as a message that the system wronged or hurt someone. It is not meant to act as a corrective justice tool. It is more like a "public mechanism of accident regulation."[102] Since the goal of the liability regime is to force automated law systems to internalize the costs of legal error, it should be sufficient to set the penalties equal to the cost of legal error without, for instance, an upward adjustment for culpability. Admittedly, this is easier to figure for some automated law systems as opposed to others. The cost of underpaid taxes equals the tax shortfall.[103] In contrast, the cost of environmental noncompliance may be more difficult to calculate.

---

[100] One example of a torts case in which a damages multiplier may have been customized is a case in which the Seventh Circuit upheld an award of punitive damages against a defendant who operated a bedbug-infested 191-room hotel. Two hotel guests sued, and the total damages award was $10,000 in compensatory damages plus $372,000 in punitive damages -- $2000 for every room of the hotel. Mathias v. Accor Economy Lodging, Inc., 347 F.3d 672, 678 (7th Cir. 2003).

[101] The *Golsen* rule provides that the Tax Court follows the law in a taxpayer's circuit of residence. *See* Golsen v. Commissioner, 54 T.C. 713 (1957). The proper damages multiplier in a circuit split situation might be designed to calculate the total cost of the legal error for all tax returns filed for residents in the circuit that gave the pro-government answer. The automated law system could bear the burden of supplying the information necessary to determine its users' residence.

[102] Alex Stein, *The Domain of Torts*, 117 COLUM. L. REV. 535, 594 (2017). A similar idea is of "licensing-based liability," distinct from "liability imposed on the basis of wrongdoing." *See* John C.P. Goldberg & Benjamin C. Zipursky, *The Strict Liability in Fault and the Fault in Strict Liability*, 85 FORDHAM L. REV. 743, 745 (2016). The proposal of automated law liability here stretches beyond the domains of inherently dangerous activities and the like in which common law tort imposes licensing-based liability. *See id.* at 784.

[103] Although the appropriate discount rate might be controversial.

III. ADVANTAGES OF GOVERNMENT-TO-ROBOT ENFORCEMENT

A. Internalize Social Cost

Government-to-robot enforcement can the costs of noncompliance where they belong – on the regulated parties who violate the law, rather than on the innocent public. The centralized feature of government-to-robot enforcement allows government to more effectively enforce the law with the limited resources available. It enables a strict liability system that seeks to correctly allocate costs and benefits, rather than seeking to allege blame and assign fault. It permits the implementation of a damages multiplier to correct for underdetection. It internalizes social cost.

The most important contribution of government-to-robot enforcement is that it addresses chronic problems of underdetection and underenforcement. Government-to-robot enforcement uses the centralized feature of automated law systems to further the goals of regulation and law by holding these centralized systems liable for all violations, regardless of the ex ante clarity of the law. This is a promising way to implement regulatory schemes. It shifts the negative externalities of violations of the law away from the taxpayer base, away from hourly workers, and away from the air-breathing public, and requires noncompliant regulated parties to bear the full cost of violations instead.

B. Market Differentiation

Government-to-robot enforcement could also support market differentiation. That is, it could allow automated system users to choose their desired level of noncompliance risk. A user could push the boundaries of a law by taking an aggressive reporting position. Or, a user could take a cautious position. If government-to-robot enforcement is in place, automated systems will charge more for the aggressive position and less for the cautious position. Compliance risk will be correctly priced, and the market will help differentiate between risk-seeking and risk-averse automated law products.

To illustrate, assume two tax software products. One, TaxDragon, takes aggressive positions. Another, CleanTax, takes conservative positions. TaxDragon will face a higher likelihood of liability for the errors made on returns it prepares. It will cost more to insure those errors. A TaxDragon customer will receive the benefit of a lower tax bill because of the aggressive positions, but will have to pay more for the product because of the high cost of insuring against the prospect of liability under the TaxDragon system. In contrast, a CleanTax customer will have a higher tax bill, but the product will

cost less because of the lower cost of insuring against automated law liability.[104]

An additional twist is that CleanTax might be likely to check with the government about the validity of its product. This may mean that it would be easier for the government to discover errors in the CleanTax software, because of its familiarity with CleanTax. If so, then the damages multiplier for CleanTax might be lower than the damages multiplier for TaxDragon, because the damages multiplier is supposed to be inversely related to the probability of detection.

IV. DISADVANTAGES OF GOVERNMENT-TO-ROBOT ENFORCEMENT

A. Capture

Under government-to-robot enforcement, makers of automated law systems and government will be encouraged to negotiate and debate the meaning of the law. In some situations, robots and government may jointly develop compliance solutions. In other situations, they may openly dispute legal determinations. But in any case, they most likely will communicate about the decisions, and they will in some sense work together.

A possible result is capture. That is, the makers of an automated law system may persuade regulators to create guidance that favors the users of the system. Public choice theory[105] and interest group theory[106] both predict capture.

For instance, consider the automated law systems that generate consumer credit reports. These reports are produced by firms and the reports contain information about consumer credit, such as whether an individual has paid bills late or missed payments entirely. The reports are sold by the consumer credit reporting firms to lenders, and lenders use the reports to decide, for instance, whether to loan money to an individual consumer.

---

[104] There could also be differentiation within a software product if an automated law provider charged different amounts of insurance based on different positions. An "audit insurance cost" bar, like the "refund due" bar might show a taxpayer how a decrease in tax liability related to an increase in audit insurance cost.

[105] *See, e.g.*, DANIEL A. FARBER & PHILIP P. FRICKEY, LAW AND PUBLIC CHOICE: A CRITICAL INTRODUCTION 144 (1965) (explaining the idea that economic rent-seeking motivates bureaucrats' actions).

[106] *See, e.g.*, James Q. Wilson, *The Politics of Regulation*, in SOCIAL RESPONSIBILITY AND THE BUSINESS PREDICAMENT 135 (James W. McKie ed. 1974) (noting the "political situation in which a small, relatively homogeneous beneficiary group can make substantial gains by imposing unobtrusive costs on large numbers of others").

A central question is the content of the algorithms used to generate these reports. Lenders would rather have "looser algorithms," which give them more negative information about consumers applying for credit. Consumers concerned about false negative information would rather have stricter algorithms. A consumer would prefer the immediate removal of a negative bit of information following that consumer compliant. Consumers in general might prefer a system that required exact matches of name, social security number and other identifying information before incorporating information into a record.[107] Consumers in general might want a consumer complaint to require a credit reporting firm to look for the alleged mistake across its entire database.

The Consumer Financial Protection Board, or CFPB, has the authority to regulate the business of credit reporting.[108] Capture of government by industry would encourage the CFPB to favor credit reporting firms by endorsing "looser" algorithms. A "looser" algorithm would favor lenders, who are the credit reporting firm's users, by overincluding information. For instance, a looser algorithm might allow inexact matches of name or Social Security number before an account is merged.

Capture can cause the law to veer off its intended course in favor of influential interest groups. The credit reporting law, let us assume, considers the costs and benefits of requiring credit reporting and strikes an appropriate balance. If the automated law system that implements the law disrupts this balance by systematically disadvantaging one group – the individuals whose credit reports are generated – to inappropriately save money for another group – the credit reporting agencies who pay for access to the rights – then the policy on the ground is out of whack.

Government enforcement is one possible answer to the problem of capture. So is the private enforcement of public rights.[109] Private attorney

---

[107] Angela Littwin, *Escaping Battered Credit: A Proposal for Repairing Credit Reports Damaged by Domestic Violence*, 161 U. PENN. L. REV. 381-83 (2013) (noting that credit reporting agencies' "main revenue source is the [lenders] who purchase credit reports").

[108] A central CFPB tool is the requirement of "internal compliance management systems," which require monitoring of company compliance and prompt response to consumer complaints. *See* Rachel E. Barkow, *Insulating Agencies: Avoiding Capture Through Institutional Design*, 89 TEX. L. REV. 15, 72-78 (2010) (outlining the institutional framework" of the CFPB). The CFPB has certain independence and other measures designed to guard against capture, but still is exposed to the problem that credit reporting agencies will seek to influence guidance so as to cause it to be more favorable for their paying customers, i.e., to cause credit reporting guidance to favor lenders by endorsing "looser algorithms." Jean Braucher & Angela Littwin, *Examination as a Method of Consumer Protection*, 87 TEMPLE L. REV. 807, 840-43 (2015) (explaining monitoring approach outlined in CFPB guidance).

[109] The availability of private enforcement of public rights varies by subject matter and

general enforcement in environmental law and securities law provides one model. Class action lawsuits might provide another. Each of these solutions, of course, have costs and benefits of their own. A full analysis is beyond the scope of this Article.

B. *Reverse Capture*

Another possible outcome, when automated law systems and government start working together more closely, is that automated law systems might favor the government. The government, in other words, might capture the automated law system. If a tax filing system can charge the same amount to a taxpayer even if the software gives pro-government answers in a grey area, then the automated law system will take the pro-government position, because it costs less to insure. This incentive might encourage an automated law system to work so closely with the government that it builds the government's views into its software,[110] instead of encouraging its users to take more aggressive reporting positions.

Reverse capture can also cause the law to veer off its intended course. The law that automated systems implement, let us assume, is supposed to achieve a certain policy or equilibrium. For instance, it balances (correctly, let us assume) the costs and benefits of an activity, such as a set of driving decisions or a level of pollution or the payment of taxes or the posting of provocative content on the Internet. If an automated law system requires cars to drive too cautiously, or bans too much pollution, or causes taxpayers to overpay the government, or takes down too much Internet content, then it upsets the equilibrium. This is inconsistent with the policy underlying the law in the first place.

Automated, centralized law systems should not be encouraged to lean in favor of the government. They should be encouraged to accurately price the risk of different positions and they should be used to properly allocate the costs of noncompliance. The goal is to debate and resolve controversies, not to eliminate them by resolving every close question in favor of the government.

If an automated law system is too pro-government, then the problem is

---

over time. One view is that it has been systematically cut back over half-century between 1964 and 2014. *See* STEPHEN B. BURBANK & SEAN FARHANG, RIGHTS AND RETRENCHMENT (2017).

[110] *Cf, e.g.*, Joshua D. Blank & Leigh Osofsky, *Simplexity: Plain Language and the Tax Law*, 66 EMORY L.J. 189 (2017) (describing IRS guidance that outlines conservative or safe harbor guidance).

the converse of the problem that most of this paper has focused on.[111] The system will cause undetected errors which the general public should pay to correct. But instead, these undetected errors will remain on the shoulders of, and burden, the users of the system.

For instance, a reverse capture problem was presented by the automated system used by Colorado to determine eligibility for welfare benefits.[112] The system denied benefits to claimants who were in fact eligible under the law. When benefits were incorrectly denied, Colorado taxpayers did not have to spend as much on welfare. But other individuals were hurt, because they did not get the benefits they deserved. One part of the story of the Colorado welfare system is simply the misinformed and tactless character of the software that the state used. But it could also be consistent with reverse capture, especially if the software systematically denied beneficiaries' claims.

Government enforcement is a poor answer to the problem of reverse capture. Reverse capture presents the problem of how to protect the rights of private individuals, like a welfare recipient denied benefits[113] or a taxpayer whose software blocks her from arguing that she is entitled to a deduction. A full analysis of this problem is beyond the scope of this Article, but there are some possible solutions that could be explored in future work. Perhaps the ability to opt out of an automated system altogether to make an individual compliance decision and litigate it directly with the government will help. Or perhaps in some circumstances a group of plaintiffs, or a public interest group representing them, should have the right to complain that an automated system disadvantaged them.[114] There is some precedent for class action lawsuits against a centralized government system that failed to provide deserved benefits.[115]

---

[111] The operation of government-run automated law systems presents the same issue of reverse capture. *Cf.* Dennis Ventry, *Americans Don't Hate Taxes, They Hate* Paying *Taxes*, 44 U.B.C. L. REV. 835, 873-74 (2011) (noting government "competitive advantage" with respect to taxpayer information).

[112] *See* Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1256 (2008) (describing Colorado's automated system for determination of welfare eligibility, which required beneficiaries to ask inappropriate and legally incorrect questions and reached incorrect legal results including improper denial of benefits in some cases).

[113] As in Mathews v Eldridge, 424 U.S. 319 (1976) (holding that due process protected the right to Social Security benefits but did not require a pre-termination hearing); Goldberg v Kelly, 397 U.S. 254 (1970) (holding that termination of welfare benefits implicated due process rights and required a pre-termination hearing).

[114] For instance, employees might claim that wage-and-hour software favors employers at their expense; or individuals covered by credit reporting might claim that credit reporting agency algorithms favor the agencies at the individuals' expense.

[115] *See, e.g.*, Craig Schneider, *Food Stamp Lawsuit Settlement: Feds to Pay Georgians $22M*, ATLANTA JOURNAL-CONSTITUTION (Jan. 11, 2016) (attributing incorrect denial of food stamp

C. The Decline of Individual Claims

The promise of government-to-robot enforcement results partly from the fact that an automated, centralized law system can resolve and accept liability for violations of law across the system, not just for the specific issue raised. This is a key feature of government-to-robot enforcement. It addresses the problems of underdetection and underenforcement and can require users to internalize fully the negative externalities of legal violations. This benefit will be achieved if government-to-robot enforcement includes a damages multiplier and allows the automated, centralized system to settle many users' disputes with the government.

It is a significant departure from existing law to suggest that the resolution of a liability associated with Taxpayer A's filing would also resolve and finalize a liability associated with Taxpayer B's filing. Yet this is what government-to-robot enforcement would mean. If a damages multiplier applied, it would overcount the cost of error to again impose a penalty on a second user, like Taxpayer B, for the same error that has already been accounted for by multiplying the penalties for Taxpayer A's error. Thus neither Taxpayer B nor the system should be held liable for the same legal error.

The question of what the "same" legal error is should tie back to the damages multiplier. In the case of tax software, for instance, the legal error might be present in the version of the software sold for a particular tax year. If the damages multiplier is accordingly set on an annual basis – for instance, by using the inverse of the annual rate of audit for the particular year – then no further liability should be imposed for errors resulting from the same legal errors in the same tax year.

But is it acceptable for Taxpayer A's case to set the system's liability for errors in the returns of all Taxpayers B? How costly is it, to Taxpayer B and to the legal system more broadly, to deprive Taxpayer B of an important opportunity to argue that her tax return was correct as submitted?

One issue is that first case to decide an issue does not always get things right. Should there be a second chance to litigate an issue or pursue a similar issue through administrative adjudication? Given the desirability of having several courts consider a problem, it would be wise to leave the door open.

---

benefits in part to the "frustrating failure" of "the agency's new centralized call-in system"). *See also* Michael D. Sant'Ambrogio & Adam S. Zimmerman, *The Agency Class Action*, 112 COLUM. L. REV. 1992 (2012) (recommending agency class action procedures for aggregating claims that draw on Rule 23 process under the Federal Rules of Civil Procedure).

One way to do this is to hold the automated law firm liable not for all similar errors estimated over all returns ever filed, but rather for returns filed in the same accounting period. Another controversy might develop for another year.

Shifting the locus of disputes to a centralized system also means that individual users give up control over contesting the legal determination of his or her rights. Sometimes a user will not be satisfied with the system's resolution of an issue with the government. What of that user's rights?

The interaction between Constitutional issues and the filing of tax returns for same-sex married couples provide an example.[116] In the 1990s, the government and the tax software programs did not allow a same-sex couple married under state law to file a married joint tax return. Instead, TurboTax followed DOMA, the federal law that excluded same-sex couple from the federal law definition of marriage. DOMA turned out to be unconstitutional. The first Supreme Court case to so hold was a tax case, brought by Edith Windsor to claim the surviving spouse estate tax exemption upon the death of her wife, Thea Spyer.[117]

Would a tax software firm have litigated *Windsor* on behalf of its users? Maybe not. Nothing about government-to-robot enforcement requires an automated system to zealously advocate on behalf of an individual user. On the other hand, it is possible that same-sex couples could have used government-to-robot enforcement as an organizing tool. That is, they could have organized or patronized a firm that purposely took and defended tax return positions grounded in the idea that DOMA was unconstitutional.

V. WINNERS, LOSERS, AND THE EXAMPLE OF SELF-DRIVING CARS

A. Government-to-Robot Enforcement for Self-Driving Cars

Some areas of the law already have a strong automated law component that is closely related to government regulation. Some of these are discussed above, including tax preparation systems, wage and hour recordkeeping software, and credit reporting agency algorithms. In contrast, the law that governs car accidents is historically grounded in the common law of tort. It is true that centralized actors – namely, insurance companies – help mediate the assignment of liability and the determination of damages. But the law of car accidents is not yet as automated, or as centralized, or as regulated, as tax law

---

[116] Numerous tax issues arose under DOMA, making it a ripe area for potential litigation. *See generally* Patricia Cain, *DOMA and the Internal Revenue Code*, 84 CHI.-KENT L. REV. 481 (2009).

[117] *See* United States v. Windsor, 570 U.S. __ (2013).

compliance, credit reporting, or some of the other historically regulatory examples.

The development of self-driving cars will change this. It will encourage government-to-robot enforcement for vehicle accidents as well. Others have predicted and/or recommended that the development of self-driving cars will lead to an increasingly important role for product liability. The idea is that the car company, not the driver, will defend against tort claims arising out of accidents involving self-driving cars.[118]

Government-to-robot enforcement goes further than this. It anticipates the end of individualized enforcement. Instead, the task of setting rules and pursuing enforcement for car accidents will rest in the hands of a government agency. A regulatory agency, rather than the court, will say what the responsibilities of the car company are. The government, rather than the injured passenger, will be the plaintiff who brings suit against the maker of the self-driving car when an accident arises.

The advantages of government-to-robot enforcement for self-driving cars follow from the discussion above. It is more efficient for government to pursue centralized cases. Centralized claims are more cost-effective and produce more legal certainty. Government can impose a damages multiplier and thus require the full internalization of negative externalities. Self-driving cars would be required to pay for all of the damage they cause, not just the cost of accidents that happen to be brought to the attention of the court system.

Centralized claims would allow efficient risk-spreading, as all of the users of a particular manufacturer's vehicles will share risks through insurance costs reflected in the price charged to use or buy a vehicle. This in turn supports market differentiation. More risk-seeking drivers might prefer a self-driving car that takes more risks on the road, but under government-to-robot enforcement, the riskier self-driving car will be more likely to produce liability for damages and will cost more to insure and more to buy or use.[119]

The disadvantages of government-to-robot enforcement for self-driving

---

[118] *See, e.g.*, Mark A. Geistfeld, *A Roadmap for Autonomous Vehicles: State Tort Liability, Automobile Insurance, and Federal Safety Regulation*, 105 CAL. L. REV. 1611 (2017).

[119] Note that self-driving cars that are less risk-averse because they operate more cautiously reduce the chance of harm both for the user of the self-driving car and also for other motorists or passengers. This contrasts with the classic example involving large sturdy cars, which protect their occupants but increase risk of harm to other users, and small flimsy cars, which fail to protect their occupants but decrease risk of harm to other users. *See* GUIDO CALABRESI, THE COST OF ACCIDENTS: A LEGAL AND ECONOMIC ANALYSIS (1970).

cars also follow the general set of disadvantages for automated, centralized regulation. This kind of enforcement increases the cost of self-driving cars and could discourage the development of self-driving car technology. Capture could present a problem, since self-driving car makers might persuade government regulators to adopt self-serving rules, such as a presumption that self-driving cars obey speed limits. Reverse capture could also present a problem, since self-driving car makers might opt to overcomply with government regulations, for example by prohibiting their vehicles from turning right at a red light.

Finally, government-to-robot enforcement could slow down law development. Once the government and a car maker have arrived at an answer to a question of liability and damages, they may not be inclined to revisit it. Once the government replaces a large number of potential individual passengers as a plaintiff, the diversity of possible case law decreases.

B. *Winners, Losers and Distributive Justice*

We can use the self-driving car example to consider the winners and losers under government-to-robot enforcement. One set of winners consists of risk-averse users of automated self-driving systems. They will be better off as compared to the current landscape of individualized enforcement, while risk-seeking persons who use the same automated self-driving systems will be worse off. Another set of winners consists of persons who are outside the scope of automated systems altogether. The self-driving car example illustrates both of these dynamics.

An important starting point is the current state of decentralized, individualized enforcement. This produces the chronic regulatory problems of underdetection and underenforcement for aggressive driving behavior. Regulators lack the resources to pursue all claims, and as a result many infractions go undetected and unpunished. Damages multipliers cannot be fully deployed in an individualized enforcement system, and so negative externalities are not fully internalized. What this all adds up to is that bad drivers can get away with it. Bad drivers claim the advantage of shorter travel times or multitasking. Other drivers bear the risk of an accident or the time cost of more careful, defensive driving.

In a self-driving car world with government-to-robot enforcement and market differentiation, a risk-seeking driver would choose a riskier car, meaning a car that took more risks on the road. The government would assess liability against the company that made the riskier car more frequently. The cost to the car company of insuring against the risk of liability would be

greater. The risk-seeking driver would pay more to use the riskier car would increase.

In contrast, a risk-averse driver would choose a safer car, meaning one that took fewer risks on the road. The government would assess liability against the car company less frequently. The cost to the car company of insuring against the risk of liability would be smaller, and the cost to the risk-averse driver of buying or using the safer car would decrease.

Requiring risk-seeking individuals to bear more liability would have the advantage of assigning costs to the place where they arise. In addition, in some cases it is said to be consistent with distributive justice objectives. This is because risk-seeking individuals tend to be wealthier than risk-averse individuals.[120] Risk-seeking and risk-aversion can of course be influenced by a variety of other factors. But there is at least the possibility that government-to-robot enforcement would serve a redistributive function if it is better at identifying and penalizing risk-seeking behavior compared to historic individualized enforcement.

If a certain region embraces self-driving cars, then risk-seeking individuals in that area would see increased costs as the damages caused by their risk-taking were more fully priced in. Risk-averse individuals in that area would benefit. Their costs of driving would be lower (assuming that they choose more risk-averse autonomous cars). Also the risk of accidents for risk-averse individuals would also go down since some other drivers would not prefer risk-seeking enough to choose the more aggressive autonomous car model. Giving an advantage to risk-averse drivers could be desirable for distributive reasons.

There is another distributive justice consideration that is raised by government-to-robot enforcement. This has to do with different treatment of such an enforcement approach of individuals who are within a system versus individuals who are outside a system. Perhaps government-to-robot enforcement could greatly improve the efficiency of enforcement and the pricing of risk among all users of automated, centralized systems. But what about individuals who are outside such systems?

If one pushes the hypothetical further into the future to consider additional technologies, then a different distributive justice question arises. This question involves the ability of the rich to opt out of surface transportation. What if privately owned drones or teleportation devices

---

[120] *Cf.* John Rawls, *Some Reasons for the Maximin Criterion*, 64 AM. ECON. REV. 141, 144 (1974) (arguing that "considerable normal risk-aversion" supports the maximin principle).

became the new mode of transport for the rich, while others are left to the self-driving car surface network? This thought experiment reveals that those with resources might be better able to opt out of the broad, run-of-the-mill automated legal systems. If the rich can opt out of the government-to-robot centralized enforcement system, so that the risk-seeking rich need not internalize the full cost of their risky behavior, then government-to-robot enforcement presents a distributive justice disadvantage. This is because government-to-robot does not touch those outside the boundaries of the automated law system.

Similar distributive justice considerations present not only for self-driving cars, but also for other centralized, automated law systems. For instance, for tax software systems, among the users of the systems, more risk-averse taxpayers win and more risk-seeking taxpayers lose. But also, the taxpayers who are rich enough to avoid TurboTax altogether win, because the enforcement they experience remains at its historically low level.[121]

Whether for self-driving cars, for tax preparation software, or for other kinds of automated, centralized liability, a possible result of government-to-robot enforcement is that it will disadvantage the middle class or the upper middle class, but not the very wealthy. This feature should prompt some consideration of an offsetting enforcement move, such as devoting yet more enforcement resources to the operation of privately owned teleportation devices, or to the auditing of very high-income tax returns.

CONCLUSION

Disputes about regulatory compliance have historically been resolved on a case-by-case basis. This is changing. Automated law systems, like tax filing or wage reporting software, make centralized legal decisions and provide a more efficient avenue for the enforcement of law. This avenue is government-to-robot enforcement. It will move the locus of disputes about the law so that they are between government, on one hand, and automated law systems, on the other hand.

Government-to-robot enforcement could allow government to overcome currently insoluble problems of underdetection and underenforcement. It could correctly require aggressive regulated parties to pay for the negative externalities of noncompliance. But it also has disadvantages, including the risk of capture of government by the market, capture of automated law systems by government, the decline of individual disputes with the government, and

---

[121] Though second-order effects e.g. from freeing up of enforcement resources should also be considered.

problems of winners and losers.