

ARTICLE

Can Robots Be Lawyers? *Computers, Lawyers, and the Practice of Law*

DANA REMUS* & FRANK LEVY†

ABSTRACT

We assess frequently advanced arguments that automation will soon replace much of the work currently performed by lawyers. In doing so, we address three weaknesses in the existing literature: (i) an insufficient understanding of current and emerging legal technologies; (ii) an absence of data on how lawyers divide their time among tasks; and (iii) inadequate attention to whether computerized approaches to a task conform to the values, ideals, and challenges of the legal profession. Combining a detailed technical analysis with a unique data set on time allocation in large law firms, we estimate that automation has a measurable impact on the demand for lawyers' time, but one that is less significant than popular accounts suggest. We then look ahead to future developments through a series of three questions. First, what is the likely path of technical innovation and diffusion in an unregulated market? Second, what are the benefits and adverse consequences of such a path? Third, to what extent can regulation reduce the adverse consequences of new technologies without reducing their benefits? Throughout the discussion, we ask how computers are changing—not simply replacing—the work of lawyers.

* Professor of Law, University of North Carolina School of Law. © 2017, Dana Remus & Frank Levy.

† Professor Emeritus, Department of Urban Studies and Planning, Massachusetts Institute of Technology and Research Associate, Department of Health Care Policy, Harvard Medical School.

The authors would like to thank: Jeanne Anderson; Andrew Arruda; David Autor; Kate Bartlett; Regina Barzilay; Joseph Blocher; Bernie Burke; Jill Burstein; David Curle; Randy Davis; Stuart Elliott; Bruce Elvin; Kathleen Engel; Bill Freeman; Ron Friedmann; Ken Grady; Maura Grossman; Dan Gutfreund; Eddie Hartman; Nathalie Hofman; Dan Jackson; Jim Jones; Jonathan Kash; Boris Katz; Scott Krowitz; Ruddick Lawrence; Nick Lemann; Maggie Lemos; Matt Levitt; Alex Levy; Marin Levy; Ryan McClead; Matthew McCubbins; Kyla Moran; Jean O'Grady; Steve Obenski; Ross Pascal; Christina Patterson; Joe Procopio; Bob Rowe; Dan Rubins; Craig Smith; Jim Spohrer; Robin Toone; Doug Ventola; Jim Wagner; Noah Waisberg; Heathcoate Wales; Vern R. Walker; Ed Walters; Jessica Watts; Brad Wendel; Keith Winstein; Yuan Zhang; Pat Ziegler; and all participants of the 2015 Stanford Legal Ethics Schmooze, the William & Mary Faculty Law School Workshop, the Wake Forest Law School Faculty Workshop, and the MIT CSAIL/Economics lunch series and the MIT Sloan School IWER Seminar. We are particularly grateful to Consilio's Sky Analytics for their generous provision of data and to the Spencer Foundation for supporting Frank Levy's work.

TABLE OF CONTENTS

INTRODUCTION	502
I. EMPLOYMENT EFFECTS	506
A. THE DATA	506
B. AUTOMATING LEGAL WORK	508
C. MACHINE COMPLEXITY VERSUS TASK COMPLEXITY.	531
D. ESTIMATING EMPLOYMENT IMPACTS	533
II. LEGAL PROFESSIONALISM IN THE DIGITAL AGE	537
A. THE MARKET FOR NEW TECHNOLOGIES	537
B. CURRENT APPROACHES TO REGULATION	541
C. THE VALUE OF REGULATION	545
CONCLUSION	556
APPENDIX	557

INTRODUCTION

On March 14, 2011, a *New York Times* headline read: “Armies of Expensive Lawyers, Replaced by Cheaper Software.”¹ In the article, *Times* technology reporter John Markoff described how computers, capable of identifying relevant words and phrases, were displacing large numbers of lawyers in discovery practice. The article posed a warning to lawyers as well as to other professionals: computers could replace humans in a highly educated, white-collar occupation.

The warning has become common wisdom. Scholars, lawyers, and commentators alike are now predicting the end of the legal profession, citing specific examples of computers successfully performing lawyers’ jobs.² Predictive

1. John Markoff, *Armies of Expensive Lawyers, Replaced by Cheaper Software*, N.Y. TIMES, Mar. 4, 2011, at A1.

2. Richard and Daniel Susskind argue that lawyers, among other professionals, face a future in which “increasingly capable machines, autonomously or with non-specialist users, will take on many of the tasks that currently are the realm of the professions.” RICHARD SUSSKIND & DANIEL SUSSKIND, *THE FUTURE OF THE PROFESSIONS: HOW TECHNOLOGY WILL TRANSFORM THE WORK OF HUMAN EXPERTS* 231 (2015); *see also* RICHARD SUSSKIND, *TOMORROW’S LAWYERS: AN INTRODUCTION TO YOUR FUTURE* (2013); RICHARD SUSSKIND, *THE END OF LAWYERS: RETHINKING THE NATURE OF LEGAL SERVICES* (2010). Law professors John McGinnis and Russ Pearce contend that “the disruptive effect of machine intelligence will trigger the end of lawyers’ monopoly.” John O.

coding, the subject of Markoff's article, is a machine learning application that automates document classification in discovery practice.³ Ross Intelligence, a legal application of IBM's Watson, advertises the ability to provide concise answers to natural language legal questions.⁴ LegalZoom, RocketLawyer, and other online legal service providers produce basic wills, divorce agreements, contracts, and incorporation papers without a lawyer's involvement.⁵ These technologies challenge the traditionalist view that lawyering is irreducibly human, and force us to recognize that computers are changing the way law is practiced.

From one perspective, the dramatic impact of technology on legal practice is nothing new. The Internet, email, and legal research databases like Westlaw and Lexis have been impacting and altering legal practice for decades.⁶ But from another perspective, we may be on the precipice of a more fundamental shift.

McGinnis & Russell G. Pearce, *The Great Disruption: How Machine Intelligence Will Transform the Role of Lawyers in the Delivery of Legal Services*, 82 FORDHAM L. REV. 3065 (2014). Other commentators predict that "[i]n the not-too-distant future, artificial intelligence systems will have the ability to reduce answering a legal question to the simplicity of performing a search," Josh Blackman, *The Path of Big Data and the Law* 1 (S. Tex. Coll. of Law Houston Working Paper Grp., 2013). Furthermore, scholars have stated, "[o]nce we have fully artificial intelligence enhanced programs like LegalZoom, there will be no need for lawyers, aside from the highly specialized and expensive large-law-firm variety." Paul Lippe & Daniel Martin Katz, *10 Predictions About How IBM's Watson Will Impact the Legal Profession*, A.B.A. J. (Oct. 2, 2015), http://www.abajournal.com/legalrebels/article/10_predictions_about_how_ibms_watson_will_impact/ [<https://perma.cc/UHY4-TPLKJ>]; Professor Dr Robot QC, *ECONOMIST* (Oct. 15, 2015), <http://www.economist.com/news/business/21674779-once-regarded-safe-havens-professions-are-now-eye-storm-professor-dr-robot> [<https://perma.cc/YN7A-8578>]; Debra Cassens Weiss, *Will Newbie Associates be Replaced by Watson?*, A.B.A. J. (Oct. 26, 2015), http://www.abajournal.com/news/article/will_associates_be_replaced_by_watson_computing_35_percent_of_law_firm_lead/?utm_source=maestro&utm_medium=email&utm_campaign=tech_monthly [<https://perma.cc/NHU9-LX3A>] [hereinafter Weiss, *Newbie Associates*]; see also MARTIN FORD, *RISE OF THE ROBOTS: TECHNOLOGY AND THE THREAT OF A JOBLESS FUTURE* (2015); JERRY KAPLAN, *HUMANS NEED NOT APPLY: A GUIDE TO WEALTH AND WORK IN THE AGE OF ARTIFICIAL INTELLIGENCE* (2015).

3. See Markoff, *supra* note 1.

4. See, e.g., Weiss, *Newbie Associates*, *supra* note 2. Ross Intelligence is developing one such application, which it describes as your "brand new Super Intelligent Attorney." See ROSS INTELLIGENCE, <http://www.rossintelligence.com/> [<https://perma.cc/DN24-DV4M>] (last visited Apr. 3, 2017) ("You ask your questions in plain English, as you would a colleague, and ROSS then reads through the entire body of law and returns a cited answer and topical readings from legislation, case law and secondary sources to get you up-to-speed quickly.").

5. See, e.g., BENJAMIN BARTON, *GLASS HALF FULL: THE DECLINE AND REBIRTH OF THE LEGAL PROFESSION* (2015) [hereinafter BARTON, *GLASS HALF FULL*]; Benjamin Barton, *The Lawyer's Monopoly: What Goes and What Stays*, 82 FORDHAM L. REV. 3068 (2014) [hereinafter Barton, *The Lawyer's Monopoly*]; Benjamin Barton, *Lessons from the Rise of LegalZoom*, BLOOMBERG BNA, at A1 (June 18, 2015), <https://bol.bna.com/lessons-from-the-rise-of-legalzoom/> [<https://perma.cc/F9UV-LSY3>].

6. Debra Cassens Weiss, *Will Technology Create a Lawyer 'Jobs-pocalypse'?*, A.B.A. J. (Jan. 5, 2016), http://www.abajournal.com/news/article/does_technology_presage_a_lawyer_jobs_pocalypse_naysayers_overstate_impact/ [<https://perma.cc/22R3-HRK8>] ("Word processing revolutionized document drafting. The Internet permitted rapid document transmission and video conferencing; accelerated the breakdown of law firms' information monopoly on rates, services, and clients; and increased clients' ability to spread legal work among multiple law firms. Email increased the speed and ease of communication both among lawyers and between lawyers and clients, and expanded the number of associates a single partner could supervise and so has facilitated the growth of large law firms.").

Machine learning applications appear poised to displace lawyers, to make inroads on the profession's monopoly, and to open new ways of addressing the access to justice gap.

In this Article, we examine prevalent claims and predictions surrounding new legal technologies, including that they are triggering the imminent and widespread displacement of lawyers by computers. In doing so, we seek to add depth and nuance to the conversation in three ways. First, we engage with technical details. We appreciate why much existing work does not—specifics blur the headlines and may be uninteresting to lay readers. But the details are critical for understanding the kinds of lawyering tasks that computers can and cannot perform. The details explain, for example, why document review in discovery practice is more amenable to automation than in corporate due diligence work, and why the automation of Associated Press sports stories and short memos on questions of law do not suggest the imminent automation of legal brief-writing.⁷ The details also offer useful insights on what is likely to be automated in the foreseeable future. We therefore offer a detailed review of salient legal technologies based on a set of unstructured interviews over an eighteen-month period with computer scientists, legal technology developers, and practicing lawyers. Second, we ground our analysis in lawyer time usage data provided by Sky Analytics, a division of Consilio.com.⁸ Lacking such data, existing employment predictions remain mere speculation. For example, scholars suggest that the automation of document review is displacing large numbers of junior associates without reference to the amount of time junior associates previously spent on document review.⁹ Our data cast doubt on these predictions.

Third, we grapple with the intersection of technological advances, access to justice, and professionalism. Many scholars maintain that “professionalism” is a mere cover for lawyer protectionism, and that the public interest is best served by commoditizing and computerizing as many legal services as possible.¹⁰ Doing

7. See *infra* notes 67–68 and accompanying text.

8. Sky Analytics assists corporate clients in monitoring and analyzing the clients' legal expenditures. See *Sky Analytics*, CONSILIO, <http://www.consilio.com/technology/sky-analytics/> [<https://perma.cc/6L9K-TMWD>] (last visited Apr. 3, 2017). The data suffers from a number of limitations, discussed below, but is nevertheless useful in providing a general picture of how lawyers spend and bill their time.

9. See, e.g., McGinnis & Pearce, *supra* note 2, at 3047–48.

10. Critics have long maintained that the professional form inures to the benefit of lawyers themselves at the expense of society by facilitating protectionism. See, e.g., SUSSKIND & SUSSKIND, *supra* note 2, at ch. 1; RICHARD L. ABEL, *AMERICAN LAWYERS* 20 (1989); ANDREW ABBOTT, *THE SYSTEM OF PROFESSIONS: AN ESSAY ON THE DIVISION OF EXPERT LABOR* 184–86 (1988); JEROLD S. AUERBACH, *UNEQUAL JUSTICE: LAWYERS AND SOCIAL CHANGE IN MODERN AMERICA* 88, 92, 99–102 (1976); MAGALI SARFATTI LARSON, *THE RISE OF PROFESSIONALISM: A SOCIOLOGICAL ANALYSIS* xvii (1977). Certainly, individual lawyers and the organized bar have at times acted in protectionist and self-serving ways. But to conclude that the professional form is therefore undesirable is to ignore the many ways in which the profession is constitutive of the state and critical to the rule of law. See Robert W. Gordon & William H. Simon, *The Redemption of Professionalism?*, in *LAWYERS' IDEALS/LAWYERS' PRACTICES: TRANSFORMATIONS IN THE AMERICAN LEGAL PROFESSION* 235 (Robert L. Nelson et al. eds., 1992); see

so, they contend, will lower costs and increase access.¹¹ Unquestionably, the profession acts in self-interested and troubling ways at times; undoubtedly, new technologies are opening promising paths for addressing the access to justice gap. But we believe that the requisite analysis is much more complex than existing accounts acknowledge.

Our discussion proceeds in two parts. In Part I, we address the extent of computer displacement of lawyer labor, seeking a more nuanced understanding than is offered in the existing literature. We use data from Consilio's Sky Analytics to test two pieces of conventional wisdom—that the overall employment impacts of computers on lawyers are significant, and that the effects are the greatest among junior associates. After reviewing near-term capabilities of computers to automate various categories of lawyering tasks, we argue that there is no strong relationship between computers' employment effects and positions within a firm. Even where automation has made significant progress, its impact has been less than the headlines would have us believe.¹²

In Part II, we explore the longer-term evolution of legal technologies by reference to three core lines of inquiry. First, we ask how legal technologies would likely develop in an unregulated market. Next, we consider the approach of existing regulatory structures, and argue that such structures unnecessarily impede the development and adoption of new technologies. Finally, we argue for the ongoing value of professional norms and regulation, notwithstanding significant problems with existing approaches. The challenge, we conclude, is to design regulatory structures that protect professional values without impeding the advance of new legal technologies.

Throughout this discussion, we focus on the ways in which computers are changing—not simply replacing—the work of lawyers. We argue that the relevant evaluative and normative inquiries must begin with an understanding of how computers perform various lawyering tasks differently than humans, and the ways in which those differences impact not only individual clients, but also the legal system writ large.

also Dana A. Remus, *Reconstructing Professionalism* (Apr. 30, 2015), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2676094 [<https://perma.cc/J359-LT4B>] [hereinafter Remus, *Reconstructing Professionalism*].

11. See, e.g., McGinnis & Pearce, *supra* note 2, at 3054–55; James E. Cabral et al., *Using Technology to Enhance Access to Justice*, 26 HARV. J.L. & TECH. 241, 246–56 (2012); William E. Hornsby, Jr., *Gaming the System: Approaching 100% Access to Legal Services Through Online Games*, 88 CHI-KENT L. REV. 917, 931–34 (2013); Ronald W. Staudt, *All the Wild Possibilities: Technology that Attacks Barriers to Access to Justice*, 42 LOY. L.A. L. REV. 1117, 1128–34 (2009); Michael J. Wolf, *Collaborative Technology Improves Access to Justice*, 15 N.Y.U. J. LEGIS. & PUB. POL'Y 759, 773–85 (2012).

12. As we explain below, the loss of junior associate jobs has been occurring over time. It is likely driven by significant weakness in the market for lawyers that accelerated with the 2008 financial collapse. By that time, much of “traditional” junior associate work—for example, document review in discovery—had already been farmed out to contract lawyers. See *infra* Part I.

I. EMPLOYMENT EFFECTS

In this Part, we first present data on how much lawyer time is devoted to various categories of lawyering work. We then review a set of basic ideas in artificial intelligence and use the ideas to explain computers' varying capacities to automate these work categories. Finally, we translate the extent of automation into a rough picture of how much lawyer time is being displaced by computers.

We anchor the discussion in the current and foreseeable trajectory of these technologies in the present and mid-term future (roughly the next decade). The resulting analysis is admittedly linear, risking that we underestimate the likelihood and impact of radical future innovation. But those who predict radical innovation have some responsibility to explain their reasoning. Simply invoking Moore's Law or pointing to an undefined future¹³ creates an argument that defies proof or refutation, and that therefore fails to inform the debate in a meaningful way.

A. THE DATA

Our data on time usage comes from Consilio's Sky Analytics of Framingham, Massachusetts,¹⁴ a consulting firm that provides corporate clients with aggregation and analysis of invoices billed by law firms. Typically, each invoice covers a small increment of time and describes the work the lawyer performed by reference to a task code from the ABA's Uniform Task-Based Management System (UTBMS).¹⁵ The UTBMS consists of 114 distinct task codes, which we have aggregated into thirteen categories for purposes of identifying patterns.¹⁶ Sky Analytics supplements the invoice with information on the submitting lawyer, including their status within the firm (associate or partner) and how many years they have been practicing.

For purposes of this project, Sky Analytics provided us with a blinded data set of invoices for 2012 through early 2015 that allowed us to construct the following information¹⁷:

13. For example, Susskind and Susskind argue the post-professional society will be reached "in the fully fledged, technology-based Internet society." SUSSKIND & SUSSKIND, *supra* note 2, at 232.

14. *Sky Analytics*, *supra* note 8.

15. See *Uniform Task-Based Management System*, AM. BAR ASS'N SECTION OF LITIG., http://www.americanbar.org/groups/litigation/resources/uniform_task_based_management_system.html [https://perma.cc/FS3P-7RCM] (last visited Apr. 3, 2017) (a set of billing codes designed to standardize the categorization of tasks).

16. Our thirteen aggregated tasks are: Advising Clients; Other Communications/Interactions; Case Administration and Management; Court Appearances; Document Drafting; Document Management; Document Review; Due Diligence; Fact Investigation; Legal Analysis and Strategy; Legal Research; Legal Writing; and Negotiation.

17. Because Sky Analytics' customer base changed from year to year, the multiple years of data were not suitable for examining trends.

- Distribution of hours billed by task.
- Distribution of hours billed by task further disaggregated by law firm size in five “Tiers” (Tier 1 > 1000 lawyers through Tier 5 < 25 lawyers).
- Distribution of hours billed by task and law firm size, further disaggregated by position in the firm (Associate ≤ 2 years; Associate > 2 years; Partner).

These data have a number of limitations. First, the original UTBMS codes (and hence our thirteen aggregated codes) allow lawyers significant discretion in how they record their time, painting at best a rough picture of time usage.¹⁸ Second, the data provide no information on the work patterns of solo practitioners, who comprise about forty percent of all practicing lawyers,¹⁹ or contract attorneys, whether hired by the law firm or the client.²⁰ It therefore focuses our analysis on law firm lawyers, primarily in the corporate hemisphere. Finally, because the invoices come from corporate clients, SkyAnalytics cannot provide a complete set of invoices billed by a single or several law firms.

Nevertheless, the data set is quite large—2014 invoices alone totaled \$2.31B—and we (and Sky Analytics) believe that a pooled sample of all billing from firms with 1000 or more lawyers (Tier 1 firms) provides a rough approximation of the distribution of hours billed to each task by junior associates (two years or less), senior associates, and partners in a typical large law firm.²¹ The data suggest that time-on-task among smaller sample firms (Tiers 2 through 5) follow a similar distribution.

Specifically, Table 1 lists the thirteen aggregated task codes with two distributions of hours spent on task: the 2012–2015 distribution of time-on-task billed by all Tier 1 firms (> 1000 lawyers) and the 2012–2015 distribution of time-on-task for all Tier 2–5 firms (all other firms in the Sky Analytics sample). We list the tasks in order of difficulty to automate—what we describe as machine complexity—which we analyze below.

18. Errors can entail both mislabeled time and inaccurately recorded amounts of time. Billing partners may also revise time allocations prior to sending an invoice to the client. For example, one interviewee explained that clients do not like to see large amounts of time invoiced to legal research, so billing lawyers might reallocate that time to the task the research is associated with, the broad category of “legal analysis and strategy,” or a category of unbilled time. Telephone Interview with Jean P. O’Grady, Author of the Dewey B Strategic Blog & Senior Dir. of Research Servs. at an Am Law 100 law firm (July 22, 2015).

19. CLARA N. CARSON & JEEYOON PARK, *THE LAWYER STATISTICAL REPORT: THE U.S. LEGAL PROFESSION IN 2005*, at 6 (2012).

20. Nor does the data account for time billed to business development or other internal matters not billed to clients.

21. The distribution of hours billed is not precisely the same as the distribution of tasks in the firm because firms bill less than 100% of junior associates’ hours.

TABLE 1
PERCENT OF INVOICED HOURS SPENT ON VARIOUS TASKS—2012–2015

Task	Tier One Firms	Tier Two–Five Firms
Document Management	0.4%	0.7%
Case Administration and Management	3.7%	5.6%
Document Review	4.1%	3.6%
Due Diligence	2.0%	3.4%
Document Drafting	5.0%	4.0%
Legal Writing	11.4%	17.7%
Legal Research	0.5%	0.4%
Legal Analysis and Strategy	28.5%	27.0%
Fact Investigation	9.2%	9.6%
Advising Clients	9.3%	3.2%
Negotiation	3.0%	5.0%
Other Communications/Interactions	8.8%	5.0%
Court Appearances and Preparation	13.9%	14.5%
Totals**	99.8%	99.7%

** Percentages may not sum to 100% due to rounding.

B. AUTOMATING LEGAL WORK

Translating Table 1 into employment effects requires an analysis of each category of legal work to understand the current and near-term potential for automation. In preparation for that analysis, we review here a set of basic ideas from artificial intelligence that undergird all legal software.

1. MODELING INTELLIGENCE

We begin with two observations: (i) virtually all of a lawyer's tasks²² involve the processing of information²³ and (ii) a computer processes information by executing instructions. It follows that for a computer to automate a lawyer's task, it must be possible to model the lawyer's information processing in a set of

22. We follow the economists' convention of describing a job as a collection of tasks because many computer applications automate a part of a job rather than the whole job. See David H. Autor et al., *The Skill Content of Recent Technological Change: An Empirical Exploration*, 118 Q.J. ECON. 1279, 1279–1333 (2003).

23. For example, a lawyer processes information about family relationships and assets into a will, or transaction information into a contract. Information processing, which we define broadly as changes in the form, organization, storage, or use of information, is central to virtually all human work.

instructions. In other words, computers can automate those lawyer’s tasks that are “structured” or “routine.”

The tasks are modeled using both *deductive instructions* and *data-driven instructions*. Deductive instructions model information processing where the structure is readily apparent—searching a legal database for opinions from a particular judge or court, or populating fields in a legal form with relevant names or other information.²⁴

Data-driven instructions arise where the structure of information processing is not apparent—the way in which an individual makes a decision as to what she will eat for lunch. In some cases, it is possible to approximate this information processing by estimating a statistical model that relates the information output to the information inputs, treating the intervening steps as a black box. Data-driven instructions are the estimated equations of such a statistical model.

Consider the problem of predicting how a judge might rule in a legal malpractice case. The information *inputs* include the facts of the case and the elements of the cause of action; the information *output* is the judge’s decision. The relationships between inputs and output are often complex and opaque, but can nevertheless be approximated by a statistical model based on a set of the judge’s prior decisions in similar cases. The model can be sketched as follows:

$$\text{Equation 1: } Y_i = \beta_1 X_{1i} + \beta_2 X_{2i} \dots \dots \dots + \mu_i$$

Where: $Y_i = 1$ if the judge decides in favor of the plaintiff in the i ’th case;
 $= 0$ if the judge decides in favor of the defendant in the i ’th case;
 $X_{1i}, X_{2i} \dots$ are case characteristics drawn from the record of the i ’th case;
 $\beta_1, \beta_2 \dots$ are the estimated coefficients of the case characteristics, including the facts of the case and elements of the cause of action; and
 μ_i is a stochastic error term for the i ’th judicial decision.

This estimation process is called “training” or “supervised (machine) learning”—supervised because the estimation requires the parameters to align with the judge’s prior decisions; learning because the estimation process can be seen as learning the relationship (summarized in β ’s) between the case characteristics and the judge’s decisions.²⁵ Once estimated, Equation 1 becomes a data-driven instruction—an instruction that can be applied to characteristics of a new case to predict judge’s decision.

24. In the early days of artificial intelligence, it was assumed most tasks could be described in deductive instructions (also called rules-based logic), but this proved incorrect.

25. The estimation process is also described as training or as a form of pattern recognition, as the computer searches for the pattern of application information that best predicts a default.

Data-driven instructions also can arise from “unsupervised (machine) learning,” which encompasses techniques to uncover patterns in a data set that can form the basis for subsequent analysis. Latent semantic analysis (LSA), an example of unsupervised learning, plays multiple roles in legal software.²⁶ For example, it creates a basis for determining whether two pieces of text are close in meaning—a measure that is useful in automating whether a particular passage is responsive to a question, or whether clauses in two contracts are conceptually similar. Testing for similar meaning involves more than asking whether the same words appear in two texts because different words can be used to convey the same meaning—for example, “automobile” and “vehicle.” Conversely, two documents can refer to different topics even while using the same word—computer chip, paint chip, chocolate chip.

LSA exploits the insight that a word’s meaning in a piece of text is partially established by its context. Beginning with a set of documents (or a number of pieces of text), the analysis first constructs a term-document matrix (Figure 1), in which each cell contains the number of times a particular term or word²⁷ appears in a particular document. Using unsupervised learning, LSA software uses correlations among words in a document to identify “word clusters”—words that usually appear in the same document when they appear in the set of documents.

FIGURE 1
TERM-DOCUMENT MATRIX

	Document #1	Document #2	Document #3, etc.
Term #1	T ₁₁	T ₁₂	T ₁₃
Term #2	T ₂₁	T ₂₂	T ₂₃
Term #3	Etc.		
Term #4			

Consider a set of 1000 documents. Suppose that the term “automobile” appears in 100 of these documents, 90 of which also include the words “safety,” “braking,” and “distance.” If the words “safety,” “braking,” “distance,” and “vehicle” (but not “automobile”), appear in another 120 documents in the set, the chances are reasonable that (a) “automobile” and “vehicle” represent the same concept in these documents and that (b) the two sets of documents involving “safety,” “braking,” and “distance” are invoking similar meanings. LSA identi-

26. In recent years, much of what LSA does is also being accomplished through probabilistic language models that use neural nets (i.e., “deep learning”). See *Vector Representations of Words*, TENSORFLOW, <https://www.tensorflow.org/versions/r0.11/tutorials/word2vec/index.html> [<https://perma.cc/KH96-WYVZ>] (last updated Dec. 20, 2016). We focus on LSA here because the logic is more intuitive.

27. “Words” in this description exclude “the,” “and,” “this,” and similar words that typically appear in every piece of text. In natural language processing, these words are called “stop words.”

fies all clusters in the set of documents and then mathematically represents each document in terms of the clusters it contains. A pair of documents, represented in this way, can then be measured for their similarity of meaning.²⁸

Machine learning models—both supervised and unsupervised—offer useful perspective on the argument that the work of lawyers (and other professionals) is more routine than we recognize.²⁹ Return to the modeling of a judge's decision—if the model fits the data, it is equivalent to saying that the judge reaches his decisions through a tacit mental protocol that ensures the same result for all cases with the same characteristics. Stated otherwise, the judge's decision-making process is “structured” or routine. Because the mental protocol is tacit—and not easily articulated—the judge may not experience his decisions as routine, but the machine learning model makes the tacit protocol explicit as a mathematical combination of characteristics taken from the case (Equation 1), which can then be used to predict future judicial decisions. In this way, machine learning unravels a part of Michael Polanyi's paradox that “[w]e know more than we can tell.”³⁰

There are, however, limits to machine learning's ability to reveal and formalize routine work. Most importantly, the task being modeled must have underlying, if unrecognized, structure—it must *actually* be routine. If, over time, the judge makes different decisions when faced with the same case characteristics, then the model will not fit the data³¹ and it will have limited predictive power (as befits an unpredictable judge).

In addition, the model's predictive ability is restricted to cases that are generally similar to the judge's past cases on which the model was estimated. If the past cases all involved female plaintiffs, the model may not correctly predict the judge's decision in cases with a male plaintiff. More generally, machine learning models—estimated statistical models—have difficulty processing contingencies that lie outside the data on which they were trained. Thus, just as computer-based question-answering systems have problems confronting questions that differ sharply from the questions on which they were trained, autonomous vehicles have problems navigating road hazards not included in their training data, and so on.³² We return to this point below.

28. The basic LSA reference is Scott Deerwester et al., *Indexing by Latent Semantic Analysis*, 41 J. AM. SOC'Y FOR INFO. SCIENCE 391, 391–407 (1990). The representation of documents in terms of clusters is roughly equivalent to principal components analysis.

29. See SUSSKIND & SUSSKIND, *supra* note 2.

30. MICHAEL POLANYI, *THE TACIT DIMENSION* 4 (1966).

31. In a well-estimated model, the coefficients (the β 's) should be statistically significant and the equation as a whole should have a reasonably high squared correlation coefficient (R²).

32. Google, Uber, Ford Motor Company, and others are currently engaged in large-scale data collection efforts for autonomous vehicles, with particular emphasis on collecting rare but dangerous events. See, e.g., Johana Bhuiyan, *Uber's Autonomous Cars Drove 20,354 Miles and Had to Be Taken Over at Every Mile, According to Documents*, RECODE (Mar. 16, 2017), <https://www.recode.net/2017/3/16/14938116/uber-travis-kalanick-self-driving-internal-metrics-slow-progress> [<https://perma.cc/CMT9-T5E9>].

There are, finally, a significant number of legal tasks that are too complex to be modeled by any set of instructions (at least at the present time). Unscripted human interaction falls into this category because it often depends on formulating responses to unanticipated questions and statements. This, in turn, requires recognizing the broader context in which words are being used—not only the surrounding words (as in LSA), but the identity and motivation of the speaker and the purpose of the communication.

Understanding context frequently requires recognizing the affect of the person making the statement. Certainly, progress has been made in the field of “affective computing,”³³ enabling computers to recognize a user’s affect by measuring physiological states and facial expressions.³⁴ But as a leader of the field explains, it is one thing to differentiate between “user is frustrated” and “user is not frustrated,” or even to differentiate between basic emotional states such as anger, fear, sadness, and love.³⁵ It is quite another, and much more difficult, for a computer to recognize and label the infinite array of more complex emotional states that we ourselves can rarely label, but that we nevertheless navigate using the tacit skills of emotional intelligence.³⁶ Such tasks lack sufficient structure to be modeled as a set of deductive or data-driven instructions and cannot be automated at this time.

2. SPECIFIC APPLICATIONS

With this background in mind, we turn to a more specific discussion of the current and likely automation of the categories of legal work. After describing the extent of automation, we label each category as subject to light, moderate, or heavy employment effects. These are, by necessity, “eyeball” classifications but they rest on examination of current and near-term technology. For example, predictive coding software clearly has a “heavy” effect on employment in document classification. In contrast, contract review software automates the search for problematic obligations in a potential acquisition, but this is only a part of the due diligence process and we classify this software as having a moderate employment effect. By adopting a partial equilibrium approach—assuming that the demand for legal work is constant so that the automation of any task results in reduced employment—we focus narrowly on automation’s current impact (while

33. See ROSALIND W. PICARD, *AFFECTIVE COMPUTING* (2000).

34. Juan Martinez-Miranda & Arantza Aldea, *Emotions in Human and Artificial Intelligence*, 21 *COMPUTERS IN HUM. BEHAV.* 323 (2005); *Special Issue on Affective Computing*, 59 *INT’L J. HUM.-COMPUTER STUD.* 1 (2003); see Aditya Khosla et al., *Modifying the Memorability of Face Photographs* (2013), http://people.csail.mit.edu/khosla/papers/iccv2013_khosla.pdf [<https://perma.cc/PHG4-ZQ6D>] (paper presented at the International Conference on Computer Vision (ICCV)).

35. Rosalind W. Picard, *Affective Computing: Challenges*, 59 *INT’L J. HUM.-COMPUTER STUD.* 55, 58 (2003).

36. See Xiaobai Li et al., *Reading Hidden Emotions: Spontaneous Micro-expression Spotting and Recognition*, CORNELL UNIV. LIBR. (Nov. 2, 2015), <http://arxiv.org/abs/1511.00423v1> [<https://perma.cc/6HV2-V6HR>].

recognizing that automation is only one among several factors shaping the market for legal services³⁷). In the long run, this is an unrealistic assumption,³⁸ but in recent years, it has proven plausible.³⁹ Moreover, it offers a transparent basis on which to begin to estimate automation's effects.

3. DOCUMENT MANAGEMENT (LIGHT EMPLOYMENT IMPACT)

The first category in Table 1, document management, entails applications designed to increase workflow efficiency, including by creating, populating, and maintaining databases and filing systems. Some aspects of this work have long been automated by networked computers and servers, and by software that can sort and search files. For decades, large firms have been using document management software that centralizes, stores, and organizes all of a firm's files, allowing all lawyers within the firm to search for and retrieve particular documents. More recent products have expanded to include automated templating, entry and billing of lawyers' hours, and the tracking of trust accounts.⁴⁰ With a few exceptions—most notably, optical character recognition to process scanned documents—these products are built using deductive instructions. We refer to them as productivity applications.

For two reasons, we believe document management productivity applications will have only a light impact on lawyer employment. First, many of these products have existed for years such that any impact would have taken effect long ago. Second, many of these tasks were previously performed by paralegals or

37. Consider, for example, that the introduction of Automatic Teller Machines (ATM) was expected to reduce the number of teller jobs per 1000 population by lowering the number of tellers per branch bank. In practice, fewer tellers per branch bank meant branch banks were cheaper to operate and for a time, the number of teller jobs per 1000 population increased, because many bank corporations began to compete by opening large numbers of branches.

38. By many estimates, more than seventy-five percent of civil legal need in the country goes unmet. See Deborah L. Rhode, *Whatever Happened to Access to Justice?*, 42 *LOY. L.A. L. REV.* 869, 869–70 (2009). The automation of lawyering tasks may address this latent market rather than replacing existing lawyer labor. Alternatively, it may push lawyers to serve this latent market as a means of finding new work.

39. See CTR. FOR THE STUDY OF THE LEGAL PROFESSION, GEORGETOWN UNIV. LAW CTR. & THOMSON REUTERS PEER MONITOR, 2016 REPORT ON THE STATE OF THE LEGAL MARKET 4, chart 3 (2016), https://peermonitor.thomsonreuters.com/wp-content/uploads/2016/01/2016_PM_GT_Final-Report.pdf [<https://perma.cc/74DD-76WW>] [hereinafter 2016 LEGAL MARKET REPORT], which shows very little growth in billings since 2010 for a sample of Am Law 100 and 200 firms. Similarly, data from the U.S. Department of Commerce "National Income and Product Accounts" on the value of legal services show that between 1990 and 2007, the value of legal services (adjusted for inflation) grew at an average rate of 12.8 percent. Between 2007 and 2013, the growth rate was –0.6 percent while the number of lawyers in the country continued to grow (by an annual 1.9 percent). U.S. Dep't of Commerce: Bureau of Economic Analysis, *Interactive Data*, tbls.6.1 B–C, <http://www.bea.gov/iTable/iTable.cfm?ReqID=9&step=1#reqid=9&step=1&isuri=1> [<https://perma.cc/UU2V-GTS7>] (last visited Apr. 2, 2017).

40. See, e.g., LAWBASE, <http://lawbase.com/> [<https://perma.cc/4NUN-RF5P>] (last visited Apr. 2, 2017); ROCKET MATTER, <https://www.rocketmatter.com/legal-billing-software/> [<https://perma.cc/4F5N-6S2Y>] (last visited Apr. 2, 2017); MYCASE, www.mycase.com [<https://perma.cc/D8ZZ-7ZBA>] (last visited Apr. 2, 2017).

clerical staff, such that they, and not lawyers, would likely feel the impact of any continuing labor displacement.

4. CASE ADMINISTRATION AND MANAGEMENT (MODERATE IMPACT)

Case administration and management encompasses tasks such as budgeting, billing, assigning and monitoring workflow, retaining experts, and managing contracts. Some of these tasks, like those that comprise document management, have been successfully automated by the umbrella productivity applications just described, which primarily impact paralegals and legal assistants.⁴¹ Some emerging products, however, use machine learning to perform aspects of contract management currently performed by lawyers. For example, KM Standards advertises software that reviews all of a company's contracts in a particular area, extracts the common provisions, and creates a basic template.⁴² The software also highlights discrepancies between the template and contracts proposed by other parties. Both tasks involve identifying similar meaning between pieces of texts and are likely accomplished using LSA or a probabilistic language model.⁴³

KM Standards joins other companies whose products, if frequently used, could have a meaningful impact on the demand for lawyers' work in the corporate sphere. Kira Systems offers software that pulls analogous provisions from different contracts into summary charts and compares particular provisions or entire documents, highlighting different contracting strategies.⁴⁴ Kira Systems' software also facilitates team and task management by organizing and monitoring task assignment and by keeping track of which documents and provisions have been reviewed and which have not. Software by the British company, Ravn, performs similar tasks and has also made strides in grouping documents by meaning, as well as in searching for particular pieces of information in an organization's files.⁴⁵

Other aspects of case management, in contrast, lie well beyond the current capacity of computers. Tasks such as monitoring junior lawyers' work or dealing with parties who fail to honor contractual obligations require unstructured human interaction of a kind that computers cannot currently perform.⁴⁶ This will no doubt change, particularly as technologies like Kira Systems increasingly

41. See *supra* note 40 and accompanying text.

42. KM STANDARDS, <http://kmstandards.com> [<https://perma.cc/Y46S-KVUB>] (last visited Apr. 2, 2017).

43. This is speculation on our part. As noted earlier, some processing previously done using LSA is now using probabilistic language models and neural networks.

44. See, e.g., KIRA, <https://kirasystems.com/> [<https://perma.cc/XP4F-XNZF>] (last visited Mar. 2, 2017); CONTRACTASSISTANT, www.contractassistant.com [<https://perma.cc/J9AX-9V2G>] (last visited Apr. 2, 2017). These and other applications encompass such tasks as filing documents, identifying differences between successive drafts of contracts, and issuing alerts on due dates of contractual obligations.

45. RAVN SYSTEMS, <https://www.ravn.co.uk> [<https://perma.cc/L5JQ-35ET>] (last visited Apr. 2, 2017).

46. In theory, software could allocate a payment if a contract provision was breached, but this would require agreement that a breach had actually occurred.

facilitate task management. For now, however, we combine a potentially significant impact of automation on demand for lawyers in contract management with the unlikely impact in these other areas, and conclude that automation is having a moderate overall employment impact on the tasks of case management and administration.

5. DOCUMENT REVIEW (STRONG IMPACT)

The essence of document review—which we define as reviewing documents for purposes of discovery in litigation or government investigations—is the lawyer’s judgment that the content of a given document is or is not responsive to an opposing party’s requests for information. Lawyers have been automating aspects of this work since the 1990s, when the explosion of electronically stored information⁴⁷ demanded some means of culling through massive electronic data sets. Early attempts, which relied on deductive instructions to search documents for keywords or combinations of keywords⁴⁸ that suggested responsiveness,⁴⁹ were highly flawed. As noted in our discussion of LSA, particular meanings and content do not necessarily correlate with specific words. As a result, searching only for specific words produces results that are both under-inclusive (risking that important documents were being overlooked) and over-inclusive (raising the costs of review by returning large quantities of non-responsive documents).⁵⁰

More recently, a number of vendors have begun marketing predictive coding technologies that use machine learning to model more accurately the basis of the human judgment regarding responsiveness. “Predictive coding” is an umbrella term encompassing significant variations and multiple products.⁵¹ Under early versions, supervising lawyers⁵² would review a “training sample” of documents (perhaps one or two thousand) from among the full data set (likely hundreds of thousands or millions of documents), classifying each document as responsive or not.⁵³ The software would then scan the training sample and estimate a supervised learning model similar in spirit to Equation 1, above. In this model, the outcome variable—“Y”—is a (0,1) variable describing the lawyers’ classification of the document as “responsive” or “not responsive.” The information inputs—“X’s”—capture characteristics of the document such as frequency of

47. Dana Remus, *The Uncertain Promise of Predictive Coding*, 99 IOWA L. REV. 1691, 1698 (2014) [hereinafter Remus, *Predictive Coding*].

48. A typical keyword search rule involving the competitive behavior of a corporation might be: Select Document if: [Price] is within fifteen words of [“customer”] or [“competitor”]. The program would then return all documents meeting the search criteria.

49. Remus, *Predictive Coding*, *supra* note 47, at 1698.

50. *Id.*

51. *Id.*

52. Initially, the training sample was coded by partners or senior associates familiar with the case. Anecdotally, the task has already been pushed down to more junior lawyers.

53. Remus, *Predictive Coding*, *supra* note 47, at 1702.

keywords, “n-grams” (three or four word sequences), word clusters derived from LSA, or other document characteristics.⁵⁴ The estimated model was then used to classify each of the remaining documents pursuant to an ex ante probability of relevancy.

Since 2012, when a federal judge first issued an opinion blessing predictive coding as an acceptable means of meeting discovery obligations,⁵⁵ use of the software has steadily increased and a number of variations have entered the market.⁵⁶ Pursuant to the most effective currently available protocol, called “continuous active learning,” supervising lawyers start by using keyword searching to select an initial set of potentially relevant documents, which they rank for relevancy.⁵⁷ These documents form the seed set and are then used to create a statistical model designed to predict responsiveness.⁵⁸

Studies show that many predictive coding technologies consistently achieve higher rates of recall⁵⁹ and precision⁶⁰ in document review than human lawyers, leading to increased use and unquestionable impacts on the demand for lawyer labor.⁶¹ Because of this, we characterize predictive coding technologies as having a strong employment effect on discovery practice.

Nevertheless, we note that predictive coding cannot completely displace lawyer labor in discovery practice for several reasons. First, attorneys must still classify a sample of documents and train the system’s parameters, leading to up-front costs that render the system inefficient for cases that do not entail large volumes of documents. Second, lawyers who have an understanding of the case, the implicated document sets, and the variety of available predictive coding

54. These characteristics are called “features” of the document. As noted above, this estimation might now be done using a probabilistic language model and a neural net that would automatically extract relevant features as part of the estimation. In practice, the lawyers test the model on subsequent sample sets in an iterative process that continues until the lawyers are satisfied that the program is appropriately classifying documents.

55. *See Moore v. Publicis Groupe*, 287 F.R.D. 182, 193 (S.D.N.Y. 2012), *adopted sub nom. Moore v. Publicis Groupe SA*, No. 11 Civ. 1279(ALC)(AJP), 2012 WL 1446534 (S.D.N.Y. Apr. 26, 2012) (“Computer-assisted review now can be considered judicially-approved for use in appropriate cases.”).

56. *See, e.g.,* KCURA, <https://www.kcura.com/about-us/our-company/> [<https://perma.cc/6KRX-6FXM>] (last visited Apr. 2, 2017).

57. Telephone conversation with Maura Grossman & Gordon Cormack (Jan. 13, 2016).

58. That statistical model, or algorithm, then ranks each document in the complete set for the likelihood that it is responsive. The top-ranking documents are “skimmed” off the top and coded by the supervising lawyers. The algorithm is then retrained using all documents that have been coded, and re-applied to the entire document set, less those that have already been set aside as responsive. This process continues until so many of the responsive documents have been identified and set aside that the highest scoring documents returned by the algorithm no longer appear responsive.

59. Recall is the fraction of all responsive documents that the algorithm identifies as responsive.

60. Precision is the fraction of all documents that the algorithm identifies as responsive that are actually responsive.

61. *See, e.g.,* Maura R. Grossman & Gordon V. Cormack, *Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient than Exhaustive Manual Review*, 17 RICH. J.L. & TECH. 1, 3 (2011); Herbert L. Roitblat et al., *Document Categorization in Legal Electronic Discovery: Computer Classification vs. Manual Review*, 61 J. AM. SOC’Y FOR INFO. SCI. & TECH. 70, 70, 74–75 (2010).

technologies are still needed to select the most appropriate products and protocols given the implicated data sets, and, if need be, to defend those choices in court.⁶² Finally, the typical algorithm assigns each document an *ex ante* probability of responsiveness, requiring lawyers to hand-classify those documents with intermediate probabilities.⁶³

6. DUE DILIGENCE (MODERATE EMPLOYMENT IMPACT)

Due diligence entails investigating and reviewing a particular client, entity, or situation to ensure comprehensive understanding of all factual and legal issues relevant to a proposed deal or transaction. Part of this, which we address here, entails reviewing documents; part, which we address below, entails investigating implicated facts and interviewing relevant parties.

The document review of due diligence differs in critical respects from the document review of discovery practice, addressed above. The former is a structured task in which a single pattern of linguistic features is used to classify an entire set of documents. The latter encompasses a structured component (locating and, where possible, analyzing the contractual obligations of a potential partner or acquisition) and an unstructured component (searching for unexpected or surprising information from a diverse set of documents).

Technology firms have worked, with some success, to automate the structured component of due diligence. Apogee Legal⁶⁴ and Kira Systems⁶⁵ have developed software that crawls a company's network to identify vendor and sourcing contracts, customer agreements, software licenses, and leases. Notably, these programs are only effective if they can be trained on a sufficient volume of similar documents. Seeking to overcome this limitation, Kira Systems has also developed a platform that flags particular clauses (for example, assignment clauses) in a diverse array of contracts and other documents.⁶⁶ The software contains a standard list of target clauses (each in multiple wordings), and additional target clauses can be specified by the user. The software uses machine learning techniques, similar in purpose to LSA, to automate the judgment that two sets of words have similar meaning.⁶⁷

Other aspects of due diligence review resist automation because they involve searching for unanticipated information—for example, a contractual relationship

62. Some versions classify documents primarily by reference to words; some by reference to word fragments (i.e., four character combinations); some by reference to metadata; and some by reference to a combination. Moreover, and as just described, some protocols begin with a random sample of documents, while others focus on clearly relevant ones.

63. We return to this point *infra* Part II.C.

64. APOGEE LEGAL, <https://apogeelegal.com/> [<https://perma.cc/3MGG-7L28>] (last visited Apr. 2, 2017).

65. KIRA, *supra* note 44.

66. Interview with Noah Waisberg & Steve Obenski, Kira Systems (Jan. 13, 2016).

67. See *How Kira Works*, KIRA, <https://kirasystems.com/how-it-works/contract-analysis> [<https://perma.cc/7NVW-JM2Z>] (last visited Apr. 2, 2017).

with a party that might violate a provision of law.⁶⁸ This limitation points to a broader issue. The human mind can draw correct inferences from very limited information,⁶⁹ allowing a human lawyer to use context, analogies, and common sense to identify a contractual reference as a problem even if the reference was unanticipated. By contrast, current machine learning software will identify the reference as problematic only if something related to the problematic language was anticipated and included in the training data.

Over time, machine learning software can improve through use as incorrect responses to previously unseen questions are corrected. This is particularly true if the software is marketed as a service sold to multiple customers⁷⁰ so that all corrections can be pooled in a single version of the software. For the present, however, these limitations remain. We therefore characterize due diligence as being subject to moderate employment effects.

7. DOCUMENT DRAFTING (MODERATE EMPLOYMENT IMPACT)

Document drafting is the development of legal documents such as deeds, contracts, wills, and trusts, that reflect the intent and agreement of the parties as accurately and unambiguously as possible. Showing that this task is, at base, structured, lawyers have long used templates in drafting these documents. Since the advent of personal computing, they have been storing these templates, often referred to as forms, on desktop computers.

More recently, a number of applications have enabled automated customization of basic forms.⁷¹ For example, a lawyer will enter information about a client's wishes regarding disposition of her estate and the computer will produce a customized will for the lawyer to review. These programs can certainly increase lawyers' efficiency, but given lawyers' prior and longstanding reliance on forms in legal drafting, we estimate that new software will only have a moderate impact on lawyer labor within firms.⁷²

A more distinct innovation, which may have a more distinct impact on lawyer employment (though outside of large firms), is the business model of online service providers that market templates directly to consumers. LegalZoom, for example, allows a consumer to obtain a number of legal documents from its website (including wills, powers of attorney, business filings, and bankruptcy or

68. Telephone Interview with Nathalie Hofman, Managing Dir., Huron Legal, Huron Consulting Grp. Inc. (July 21, 2015).

69. See, e.g., Linda Smith—*The Visual Side of Early Object Name Learning (1 to 2 year old toddlers)*, YOUTUBE (Oct. 5, 2015), <https://www.youtube.com/watch?v=zDZlpqJkNe8> [<https://perma.cc/7DHH-PS6Y>].

70. The alternative is that each customer operates a free-standing software package.

71. Automated document drafting programs are frequently incorporated into document management software used by law firms. See *supra* note 40.

72. See *infra* note 150 and accompanying text for discussion of another frequently discussed innovation, blockchain contracts.

divorce petitions)⁷³ by indicating the document in which he or she is interested and answering a series of document-specific questions. Based on the consumer's answers, the website produces a completed and customized document.

We return to online service providers below. For now, we simply note that any employment impacts on lawyers, which we characterize as moderate, will be felt among solo practitioners and lawyers in small firms.

8. LEGAL WRITING (WEAK EMPLOYMENT IMPACT)

Legal writing, as distinct from legal document drafting, is the development of written work that characterizes the state of the law and/or its application to a particular factual situation. Whether objective or persuasive, legal writing is very difficult to automate. Commentators cite automated Associated Press summaries of baseball games and corporate earnings reports to argue that this will soon change⁷⁴ and that automated legal briefs are right around the corner, but the analogy does not hold. Extracting and summarizing relevant information about a baseball game or a company's reported financial situation is a structured task—a baseball game can be largely reconstructed from the pitch-by-pitch game feed, and earnings reports have relatively structured and standard formats. Once the game (or the earnings report) has been reconstructed, writing the summary involves a structured selection and listing of prewritten phrases with insertion of particular proper nouns (e.g., players' names).⁷⁵

Notwithstanding recent innovations by Ross Intelligence, discussed below, to respond to a legal query with a short memo, the vast majority of legal writing is insufficiently structured to be automated in this way. Whereas a sports writer covers a settled game structure and a final definitive score, a lawyer often writes amidst indeterminacy. Certainly, parts of a brief are standard and predictable—for example, the preliminary and concluding material, and the statement and explanation of relevant standards of review. But much legal writing requires conceptual creativity and flexibility that computers do not currently exhibit. The analysis section of a legal brief requires a complex interplay between law and fact, in which the law that governs is determined by the facts, while the relevant facts are determined by the governing law. The use of precedent, while second nature for a lawyer, is exceedingly difficult (and currently impossible) to model

73. *About Us*, LEGALZOOM, <https://www.legalzoom.com/about-us> [<https://perma.cc/9KF5-684B>] (last visited Apr. 2, 2017).

74. McGinnis & Pearce, *supra* note 2, at 3051.

75. Note also that these short articles are usually directed at readers with limited information demands: a New York newspaper will contain an extensive, non-automated article on a Yankees game while using Associated Press summaries to report on out-of-town games. For example, Automated Insights's website describes producing stories in four steps: "1. Add your data; 2. Write your template; 3. Preview your stories; 4. Publish your stories." *Here's How Wordsmith Works*, AUTOMATED INSIGHTS, <https://automatedinsights.com> [<https://perma.cc/3PHK-TBHV>] (last visited Apr. 2, 2017).

for a computer. A single case can be used to support two opposing positions; arguing for one as opposed to the other requires an ability to contextualize the case in a line of precedent, and to distinguish between a binding holding and non-binding dicta. Often, an effective legal argument also requires the ability to transplant concepts from one area of law to another in order to argue for a novel legal theory or change in the law. These unstructured and opaque conceptual tasks lay beyond the current capacity of computers. We therefore categorize legal writing as currently subject to weak employment impacts.

9. LEGAL RESEARCH (MODERATE IMPACT)

Vern R. Walker, co-organizer of the first conference on argument mining, offers a useful perspective on the evolution of automated legal research:

The ultimate goal of legal research by lawyers and decision makers is to find arguments and reasoning reported in the past, so that they can evaluate the likelihood of success of those and similar arguments, and can generate new arguments to use in future cases. I think that this suite of tasks is also the ultimate goal of software analytics, such as Westlaw, FastCase, Ross, etc. That's the direction in which we are all headed with automating legal research.⁷⁶

From an artificial intelligence perspective, legal research is a problem in information extraction, where the critical design element involves linking a user's search query to the best available answers. Early innovations in automated legal research came from computerized legal databases such as Westlaw and Lexis. A user of Westlaw and Lexis could begin with a key word search, but keyword searching is frequently both under and over-inclusive.⁷⁷ The innovation of both services was to offer an indexing tool—an improved link between query and legal case. Constructing these indexing tools involves a significant amount of human processing, however, and is therefore expensive.⁷⁸ Because of the difficulty of automating text summarization, discussed below, humans write a summary headnote for each case filed in the system.⁷⁹

For Lexis, humans then use these headnotes to classify cases into the Lexis Topic system; for Westlaw, a machine learning algorithm links the headnotes to

76. Interview with Vern R. Walker, Professor of Law, Dir., Research Lab. for Law, Logic, & Tech., Hofstra Law Sch. (June 29, 2016).

77. A user could refine the keyword-searched database so as to examine only cases in a relevant jurisdiction or time period.

78. In economic terms, the indexing is a large fixed cost, which explains why the system is constructed by a small number of providers and sold as a service.

79. See Dan Jurafsky & Christopher Manning, *Natural Language Processing*, STANFORD ONLINE (June 13, 2012), <http://online.stanford.edu/course/natural-language-processing> [https://perma.cc/5UN5-NVDD].

the West Key typology codes.⁸⁰

FastCase, a 2010 entrant to the legal research market, abandoned the index system and instead links a query to cases primarily by a combination of citation frequency and relative strength of the citation.⁸¹ The algorithm functions similarly to Google's algorithm for searching the web,⁸² and like Google's algorithm, presents results ranked in terms of estimated relevance.⁸³ New versions of Lexis and Westlaw are incorporating relevancy rankings as well, based on a combination of features such as past search patterns, document characteristics, and matching terms.

In Walker's framing, Lexis, Westlaw, and FastCase respond to legal queries by retrieving citations to complete cases—full documents that contain legal arguments or that set forth or summarize existing law. More recent software research tools focus on retrieving something closer to the underlying arguments themselves. One application gaining substantial publicity is Ross Intelligence, an IBM Watson-based question and answer (Q/A) system that accepts natural language questions rather than keywords and that retrieves relevant passages rather than entire cases.⁸⁴

Assume a user enters the following question on bankruptcy law: “When can a debtor reject a collective bargaining agreement?” Roughly speaking, a Watson-based Q/A system answers this question in three phases.⁸⁵ A parsing module determines what the question is about (i.e., the entities of interest—debtor, collective bargaining agreement, and the relationship among entities (rejection)). A second module retrieves potentially relevant passages from the system's database. A third module ranks the retrieved passages, assigning each passage the probability that it represents the best answer. The second and third modules rest on LSA and related techniques, discussed above, to measure the similarity of candidate passages to the question.⁸⁶

80. See Susan Nevelow Mart, *The Relevance of Results Generated by Human Indexing and Computer Algorithms: A Study of West's Headnotes and Key Numbers and LexisNexis's Headnotes and Topics*, 102 L. LIBR. J. 221, 223–24 (2010).

81. See *What is Fastcase?*, FASTCASE, <http://www.fastcase.com/whatisfastcase/> [<https://perma.cc/W24H-KQDP>] (last visited Apr. 2, 2017).

82. See *How Does Google Search the Web*, YOUTUBE (Apr. 23, 2012), <https://www.youtube.com/watch?v=KyCYyoGusqs> [<https://perma.cc/6YRH-FWDC>].

83. Initial versions of Westlaw and Lexis listed results in reverse chronological order or frequency of keywords.

84. ROSS INTELLIGENCE, *supra* note 4.

85. “Watson” is actually a suite of applications. The configuration of Watson that won *Jeopardy!* was comprised of sophisticated natural language processing capabilities and an ability to access multiple databases (including Wikipedia) to search for answers. Developers of more recent Watson applications often retain the natural language processing capabilities, but build their own databases.

86. For a more complete description of Watson's question answering architecture, see David Ferrucci et al., *Building Watson: An Overview of the Deep QA Project*, 31 AI MAG. 59 (2010); Dan Jurafsky & James H. Martin, *Speech and Language Processing* (3d ed. draft Jan. 2017), <https://web.stanford.edu/jurafsky/slp3/> [<https://perma.cc/PVN8-3S38>].

The system is modeled using neural nets, statistical models with multiple non-linear interactions among variables. Such models are statistically complex, but at base similar to Equation 1 above. The neural net takes information inputs—the characteristics of a particular passage—which it then processes into an output—the probability that the passage is the best answer to the question.

When a user builds a new system, much of the language-parsing module is prepackaged. But the retrieval and ranking neural nets must be trained through supervised learning, and so, like Westlaw and Lexis, require a substantial initial effort.⁸⁷ The training process begins by populating the system's database with legal documents that have been broken into passages by human experts. The experts essentially annotate the database by attaching to each passage a set of natural language practice questions such that the passage is the correct answer for each of the attached questions. Each practice question must be worded in multiple ways to reduce the likelihood that the system will fail to recognize a user's question as having the same meaning as an already-processed question. The system is then trained using an iterative process of posing a question, noting whether the system's suggested passage is correct or incorrect, and adjusting the neural net's parameters (roughly equivalent to β 's in Equation 1) accordingly. The process is complicated by the fact that questions are often imprecisely expressed. For example, in the question above, does "When" mean the user expects an answer in the form of a time period ("after ninety days . . .") or a set of conditions ("if the collective bargaining agreement contains . . .").

If experts could anticipate precisely how every question would be asked, there would be no need for machine learning to estimate statistical links—each specifically worded question could be tied directly to its answer. Once in operation, however, the system will receive many questions that are more or less similar but not identical to the questions on which it has been trained. Consider the following example: "Can a debtor reject a collective bargaining agreement where debtor is a city that filed for Chapter 9 bankruptcy and previously attempted to negotiate with a private union before rejecting its collective bargaining agreement?"

This question, both specific and complicated, is unlikely to have been part of the system's training. And yet, it has linguistic features in common with the training question discussed above, as well as other linguistic features that may be able to activate other links: the debtor is a city, the city is in Chapter 9 bankruptcy, a time relationship in which negotiations occurred before the collective bargaining agreement was rejected. Using all relevant links, there is a chance that

87. As is the case with Westlaw and Lexis, the large initial cost means such customers access such question/answering systems by purchasing them as a service rather than building their own system. The corollary is that the system must be focused on a set of questions/answers that are of potential interest to a broad set of customers. Current versions of Ross are focused on bankruptcy law.

the system will produce a relevant and responsive set of paragraphs for this question, which it has not previously seen.

Accordingly, an operational Q/A system will likely confront questions it has not seen before. Questions involving abstract concepts and analogies can be particularly complex to analyze. These problems can be reduced over time with continuous training that refines question-answer links, but such training requires use by senior attorneys and not just young associates who may be too inexperienced to spot the system's errors.

As noted above, Ross's Q/A system now offers yet another innovation—an ability to answer certain legal questions with well-organized two-page memos rather than relevant passages.⁸⁸ Ross officials are understandably reluctant to discuss their technology, but Ross's close association with IBM makes it reasonable to speculate that Ross's memos rely on a variant of the answers produced by IBM's Debater System.⁸⁹ The Debater System is part of the developing field of argument mining, the subject of Walker's quote above. Argument mining is an area of natural language processing that draws on a text corpus—for example, Wikipedia entries or sections of *Collier on Bankruptcy*—to create a short essay supporting or opposing a stated proposition. To set up the software, human reviewers first review the corpus to identify specific topics and label three types of passages associated with each topic: background, claims, and evidence.⁹⁰ Setting up the software also requires constructing an essay template in which selected passages will be inserted into background, claim, and evidence fields to form the completed document. When the system is presented with a question, the software selects candidate passages based on topic and ranks them for their responsiveness to the question—a process broadly similar to judging the similarity between pieces of text. The highest ranked background-claim-evidence sets are inserted into template slots to form the essay.⁹¹

Strictly speaking, this is not a fully automated system because it initially requires humans to label the corpus. Humans label the corpus only once, however, regardless of the number of users. Still, in considering the employment impacts of this and all legal research tools, it is important to recognize the substantial remaining human role in defining and directing research. This role leads us to characterize the impact of automation on legal research as moderate

88. Susan Beck, *Inside ROSS: What Artificial Intelligence Means for Your Firm*, LAW.COM (Sept. 28, 2016), <http://www.law.com/sites/almstaff/2016/09/28/inside-ross-what-artificial-intelligence-means-for-your-firm/?sreturn=20170114162046> [https://perma.cc/NHJ4-PB4W].

89. *IBM Debating Technologies*, IBM, http://researcher.ibm.com/researcher/view_group.php?id=5443 [https://perma.cc/EB8B-42AX] (last visited Apr. 2, 2017).

90. See Ehud Aharoni et al., *A Benchmark Dataset for Automatic Detection of Claims and Evidence in the Context of Controversial Topics*, at 64–65 (Proceedings of the First Workshop on Argumentation Mining, Ass'n for Computational Linguistics, June 26, 2014), <http://www.aclweb.org/anthology/W14-2109> [https://perma.cc/AP9U-XFH6].

91. See *id.*

notwithstanding impressive advances in legal research tools. Consider, for example, the nature of a case law search, which frequently begins with an initial set of controlling cases. As Susan Mart writes:

[I]t is rare that the facts of those cases are so close to the facts of the client's case that your research is complete. The second part of the research project then begins—the search for case-specific relevant authority. The researcher needs to find other cases, similar in legal conclusions and more similar factually to the client's case. This search for more specifically relevant primary law can be called “level two research.” The researcher uses the major and controlling cases in the relevant area of the law (however located) as seed documents to link forward through headnotes, key numbers, KeyCite, and Shepard's or backward through headnotes, key numbers, and the cases cited in the seed cases. This type of forward and backward searching from seed documents is instrumental for finding “application cases”—cases that have only marginal value as support for an abstract proposition of law, [but] have great value in their application of the proposition to facts similar or analogous to the facts of your own case.⁹²

Mart describes an iterative process in which a lawyer specifies the parameters for a search, which the software then performs. It is therefore the search that has been automated, not the entire task of researching precedents. A similar logic applies to question and answering systems. They can automate an actual search, often more effectively than Westlaw or Lexis, but they cannot automate the designation of search parameters. That work remains for lawyers—most often, for associates.

10. LEGAL ANALYSIS AND STRATEGY (MODERATE EMPLOYMENT IMPACT)

Legal analysis and strategy entails the exercise of legal judgment in evaluating a situation and planning accordingly. Two advances in automation have made inroads on this work, which was traditionally thought of as immune to automation. The first is prediction. In recent years, software such as Ravel Law and Lex Machina have collected and analyzed massive amounts of data on judges and their decisions, producing data-driven statistical models, similar in structure to Equation 1, that are often more accurate than human prediction.⁹³ Automated prediction of jury decisions has proven far more elusive, however, and even with

92. Mart, *supra* note 80, at 222 (footnotes omitted).

93. For example, a 2015 press release explains: “By analyzing millions of legal documents, Ravel provides strategic insight into an array of factors that affect a judge's decision-making.” *Ravel Law Announces Unprecedented Judge Analytics Offering*, PRWEB (Apr. 16, 2015), <http://www.prweb.com/releases/2015/04/prweb12656883.htm> [<https://perma.cc/SHC9-VB2J>]; see also *What We Do*, LEX MACHINA, <https://lexmachina.com/what-we-do/> [<https://perma.cc/S9PS-K5N5>] (last visited Mar. 3, 2017) (“We mine litigation data, revealing insights never before available about judges, lawyers, parties, and patents, culled from millions of pages of IP litigation information. We call these insights Legal Analytics®, because analytics involves the discovery and communication of meaningful patterns in data.”).

respect to bench trials, a significant human role remains in interpreting the data and formulating advice for clients. Prediction software and data analytics also offer the possibility of law firms getting a better understanding of the risks and costs associated with large cases. While such knowledge may allow a law firm to run with greater efficiency, it has ambiguous employment effects.

A second area of progress in computerized legal reasoning is the development of expert systems. Built on platforms by Neota Logic⁹⁴ and others, expert systems organize and present a specific, narrow legal task as a structured dialog with the user. Once constructed, these systems can be scaled to many users, delivering legal reasoning at a much lower cost than if a human lawyer responded to each user separately. A recent, much discussed example is DoNotPay, an expert system developed by a 19-year-old British student that helped British drivers overturn at least 160,000 parking tickets.⁹⁵

As of now, such systems can only be constructed for repetitive and fairly narrow tasks under specific bodies of law—for example, compliance with the Foreign Corrupt Practices Act.⁹⁶ As the parking ticket example suggests, not all users of expert systems would have otherwise used a lawyer. To the contrary, expert systems are often described as a way for a law firm to offer a low-cost service to potential clients for addressing low-stakes cases. The hope is that the client then turns to the firm to deal with more complex situations. Because of this, we characterize the employment impacts of automation on legal analysis and strategy as moderate.

11. ADVISING CLIENTS AND OTHER COMMUNICATIONS/INTERACTIONS (WEAK EMPLOYMENT IMPACT)

For current purposes, we group two very different sets of tasks—advising clients and communicating and interacting—with all others. Although the work of these two categories is distinct, it requires a significant amount of unstructured human interaction, rendering both categories of work subject to weak employment impacts.

With respect to client advising, computers have made significant progress in two areas. They have made it easier for individuals to access relevant legal information, whether through free online legal databases or issue-specific

94. NEOTA LOGIC, <http://www.neotalogic.com> [<https://perma.cc/FRG3-6DFK>] (last visited Apr. 2, 2017).

95. Samuel Gibbs, *Chatbot Lawyer Overtakes 160,000 Parking Tickets in London and New York*, *GUARDIAN* (June 28, 2016), <https://www.theguardian.com/technology/2016/jun/28/chatbot-ai-lawyer-donotpay-parking-tickets-london-new-york> [<https://perma.cc/46QT-JK6U>].

96. See, e.g., *Foley & Lardner LLP Selects Neota Logic Technology Platform to Power Award-Winning Foley Global Risk Solutions Offering*, NEOTA LOGIC (2016), <https://cdn2.hubspot.net/hubfs/453162/Foley%20Lardner.pdf?submissionGuid=c6c807c3-3154-40c7-af4c-5ba8c6315cd1> [<https://perma.cc/B7MA-58G7>].

web-based applications.⁹⁷ They have also made progress in an area just noted—prediction.

For at least three reasons, however, most client advising remains outside of the current domain of automation. First, legal prediction software programs address only courts and case law, but lawyers must routinely predict many other things, such as how an opponent will react to a settlement offer or how an agency will interpret a regulation. Second, many clients want more than a series of statistical probabilities. They want a lawyer's judgment and assurances as to what course of action will most effectively serve their short and long-term interests. Some clients want this for their own comfort; others want it to reassure affected constituents; still others want it for purposes of a potential advice-of-counsel defense.⁹⁸ Third and most importantly, effective advising encompasses more than prediction. It requires a lawyer to understand a client's situation, goals, and interests;⁹⁹ to think creatively about how best to serve those interests pursuant to law; and sometimes, to push back against a client's proposed course of action and counsel compliance.¹⁰⁰ These are things that frequently require human interaction and emotional intelligence¹⁰¹ and cannot, at least for the time being, be automated.

More broadly, the vast majority of a lawyer's personal interactions—with clients as well as with all others—continue to require spontaneity, unstructured communication, and emotional intelligence. Examples are plentiful: A lawyer may need to push a client to execute a will; spend hours interviewing a criminal defendant to develop enough trust to elicit full information; or read a deponent's facial expressions and body language to determine how to proceed with questioning. Moreover, many individual clients report that a lawyer's trustworthi-

97. See, e.g., Tanina Rostain et al., *Thinking Like a Lawyer; Designing Like an Architect: Preparing Students for the 21st Century Practice*, 88 CHI.-KENT L. REV. 743 (2013) [hereinafter Rostain et al., *Thinking Like a Lawyer*].

98. Many clients may want a lawyer's advice as a means of avoiding what behavioral economists refer to as "regret"—the guilt and responsibility that can accompany a wrong decision in an uncertain situation. Cf. Richard H. Thaler, *Toward a Positive Theory of Consumer Choice*, 1 J. ECON. BEHAV. & ORG. 39, 54 (1980) (observing that one reason doctors look for second opinions is to share responsibility and reduce regret for diagnoses that may turn out to be wrong).

99. See Katherine R. Kruse, *Beyond Cardboard Clients in Legal Ethics*, 23 GEO. J. LEGAL ETHICS 103, 104 (2010) [hereinafter Kruse, *Beyond Cardboard Clients*]; Katherine R. Kruse, *The Promise of Client-Centered Professional Norms*, 12 NEV. L.J. 341, 346 (2012) [hereinafter Kruse, *Client-Centered Norms*]; Marc Lauritsen, *Liberty, Justice, and Legal Automata*, 88 CHI.-KENT L. REV. 945, 954 (2013); William H. Simon, *The Ideology of Advocacy: Procedural Justice and Professional Ethics*, 1978 WIS. L. REV. 29, 53.

100. See generally DEBORAH L. RHODE, *IN THE INTERESTS OF JUSTICE* (2000) [hereinafter RHODE, *INTERESTS OF JUSTICE*]; Robert W. Gordon, *A New Role for Lawyers?: The Corporate Counselor After Enron*, 35 CONN. L. REV. 1185 (2003) [hereinafter Gordon, *New Role for Lawyers?*]; Tanina Rostain, *Ethics Lost: Limitations of Current Approaches to Lawyer Regulation*, 71 S. CAL. L. REV. 1273, 1274–75 (1998) [hereinafter Rostain, *Ethics Lost*]; W. Bradley Wendel, *Professionalism as Interpretation*, 99 NW. U. L. REV. 1167, 1171–72 (2005) [hereinafter Wendel, *Professionalism as Interpretation*].

101. Remus, *Reconstructing Professionalism*, *supra* note 10, at 25–29.

ness and ability to provide a close and personal relationship are among the most important traits they look for from a lawyer.¹⁰² For the time being, therefore, we think the impact of automation on the areas of client counseling and interactions with third parties will remain weak.

12. FACT INVESTIGATION (WEAK EMPLOYMENT IMPACT)

Fact investigation, similarly dependent on unstructured communications, is also subject to weak employment impacts. Some aspects of the task can be automated—for example, software can usefully pull together vast amounts of online data regarding a client or opponent, and some lawyers and legal aid clinics automate initial client intake. For the most part, however, factual investigation resists automation. It frequently entails interviews in which significant amounts of information may be transmitted nonverbally, in ways a computer would have difficulty detecting, at least for now. It also requires flexibility from a lawyer, beyond the capacity of a computer, in adjusting the relevant questions as new information is discovered.

13. NEGOTIATION (WEAK EMPLOYMENT IMPACT)

Traditionally, negotiation also required personal interaction and effective use of emotion. Negotiation experts have long theorized that skill in reading an opponent's emotions allows a negotiator to achieve greater understanding of the opponent's interests and concerns, to assess risk more accurately, and to deploy negotiation tactics more effectively.¹⁰³ Online dispute resolution programs are rendering these human skills unnecessary in a small but growing category of cases, however.¹⁰⁴ An example is Modria, a California firm that markets online dispute resolution to e-commerce companies.¹⁰⁵ Its website describes that it

102. See COREY S. SHDAIMAH, *NEGOTIATING JUSTICE: PROGRESSIVE LAWYERING, LOW-INCOME CLIENTS, AND THE QUEST FOR SOCIAL CHANGE* (2009) (citing interviews of clients expressing that friendship and trust were at the forefront of what they wanted from lawyers); Marcus T. Boccaccini et al., *Client-Relations Skills in Effective Lawyering: Attitudes of Criminal Defense Attorneys and Experienced Clients*, 26 *LAW & PSYCHOL. REV.* 97, 110–11 (2002) (citing a poll in which clients ranked obtaining clients' opinions, spending time with clients before court, and keeping clients informed of their cases as among the things they cared most about in a lawyer; also citing evidence that inmates cared more about a lawyer who cared about them, would be honest, and would spend time with them before their court date than about the lawyer's skills); Marcus T. Boccaccini & Stanley L. Brodsky, *Characteristics of the Ideal Criminal Defense Attorney from the Client's Perspective: Empirical Findings and Implications for Legal Practice*, 25 *LAW & PSYCHOL. REV.* 81, 98 (2001); Anne E. Thar, *What Do Clients Really Want? It's Time You Found Out*, 87 *ILL. B.J.* 331 (1999).

103. See Kimberlyn Leary et al., *Negotiating with Emotion*, *HARV. BUS. REV.* (Jan.–Feb. 2013), <https://hbr.org/2013/01/negotiating-with-emotion> [<https://perma.cc/F7CL-PUL4>].

104. See Richard Susskind & Matthew Levy, *Likely Developments in ODR*, *CTS. & TRIBUNALS JUDICIARY* (Feb. 16, 2015), <https://www.judiciary.gov.uk/publications/likely-developments-in-odr/> [<https://perma.cc/U4V5-68S8>]; see also Hornsby, *supra* note 11.

105. *About Us*, MODRIA, <http://www.modria.com/about-us/> [<https://perma.cc/MD4D-PFVR>] (last visited Apr. 2, 2017).

gathers relevant information regarding the dispute, summarizes areas of agreement and disagreement, and makes suggestions for resolving the issue.¹⁰⁶ It does so through deductive instructions, rendering negotiation (as lawyers understand the task) unnecessary.¹⁰⁷

Currently, the approach is used primarily for small disputes, but Modria is expanding into larger and more complicated types of disputes.¹⁰⁸ A number of other companies are developing similar products,¹⁰⁹ while legal reform groups are encouraging courts to increase efficiency and manage dockets through use of such products.¹¹⁰

Additional new technologies are emerging to aid lawyers in negotiating by, for example, analyzing and representing the overlap between two parties' preferences.¹¹¹ Such programs address one or at most two issues, and their resolution is constrained by the parties' stated initial preferences. They nevertheless suggest that computers may eventually play a larger role in aiding, if not replacing, lawyers' negotiating work.

In theory, online dispute resolution and expert systems could also fall into the category of heavy employment effects given that when used, they entirely replace lawyers (and in the case of online dispute resolution, judges as well). Their impact on lawyer employment may be significant in the future, but we estimate it is minimal at present. The disputes that these systems resolve are generally small stakes e-commerce issues, for which it would not be economically feasible to hire

106. *How it Works*, MODRIA, <http://www.modria.com/how-it-works> [<https://perma.cc/4LRF-RH9H>] (last visited Apr. 2, 2017).

107. For example, a deductive instruction could read: ("If (Customer is Low Risk) and (Dispute Amount is less than ten dollars) and (Customer Disputes Filed Account Lifetime is (0)) then (Authorize Full Refund) and (Close Case).").

108. Eric Johnson, *Modria Wants You to Settle Your Workplace Problems (and Even Patent Disputes) Online*, ALL THINGS D (Nov. 24, 2012), <http://allthingsd.com/20121124/modria-wants-you-to-settle-your-workplace-problems-and-even-patent-disputes-online/> [<https://perma.cc/DC8L-Q6GX>]; see also *Modria—The Operating System for ODR (video extract)*, CTS. & TRIBUNALS JUDICIARY (Feb. 16, 2015), <https://www.judiciary.gov.uk/publications/modria-the-operating-system-for-odr-video-extract/> [<https://perma.cc/6UZP-VP3J>] ("[O]ur goal is to be the operating system for online dispute resolution. So any kind of dispute, no matter how complicated or how simple, how high volume or low volume, we can use these building blocks at Modria to build an appropriate resolution path for that dispute. So that's our objective.").

109. See, e.g., Graham Ross & Beth Silver, *Case Studies*, CTS. & TRIBUNALS JUDICIARY (Feb. 16, 2015), <https://www.judiciary.gov.uk/publications/case-studies/> [<https://perma.cc/K87F-RFQE>].

110. See, e.g., Civil Justice Council Online Dispute Resolution Advisory Grp., *Online Dispute Resolution for Low Value Civil Claims*, at 3 (Feb. 2015), <https://www.judiciary.gov.uk/wp-content/uploads/2015/02/Online-Dispute-Resolution-Final-Web-Version1.pdf> [<https://perma.cc/KJP6-LJKB>] (recommending a new Internet-based court service, which would offer three services: Online Evaluation, which would help users to understand and evaluate their potential claims; Online Facilitation, which would facilitate early resolution of disputes without the involvement of a judge; and Online Judges, who would decide parts or all of cases through structured online pleading).

111. Keith Winstein, for example, has shown that some telecom-related negotiations on access prices could be solved by an online auction. See *Keith Winstein's Homepage*, STANFORD COMPUT. SCI., <http://cs.stanford.edu/keithw/> [<https://perma.cc/5JLW-TM3J>] (last visited Apr. 2, 2017).

a lawyer and litigate (such that lawyer labor is not being replaced).¹¹² Similarly, expert systems are generally directed at reaching new markets rather than improving efficiency in existing tasks. For example, a law firm might offer subscriptions to an expert system covering aspects of tax compliance to clients who otherwise might not consult a lawyer,¹¹³ increasing the firm's business but not displacing any lawyers.¹¹⁴ For the time being, therefore, we characterize the impact of these services on lawyer employment as weak.

14. COURT APPEARANCES AND PREPARATION (WEAK EMPLOYMENT IMPACT)

A final category of work, courtroom advocacy, is distinct from the others insofar as even the most fervent technology advocates are not predicting near-term automation. In part, this is because the policies and restrictions of unauthorized practice of law rules operate at their strongest in the courtroom. More fundamentally, it is because effective advocacy requires emotional engagement with the decision-maker.¹¹⁵ As two experienced advocates explain:

An inexperienced trial lawyer's dull and confusing closing argument in a complex business dispute will create negative feelings of boredom and frustration in the minds of the jurors . . . an accomplished advocate can communicate to the juror the facts of the identical dispute in a way that will evoke positive emotions about justice and fairness in the marketplace.¹¹⁶

It is not only in arguments to a jury that emotion is critical. The emotions a lawyer deploys to persuade a judge may differ from those designed to persuade a jury, but emotion is a critical spur to all action and decision-making.¹¹⁷ And yet, the field of affective computing is nowhere near enabling computers to foster, recognize, and respond to the full range of human emotions.

* * *

Table 2 summarizes the foregoing discussion by restating the time usage data of Table 1, while also indicating the employment effects on lawyers of automation of each task.

112. See *supra* notes 93–96 and accompanying text.

113. The subscription might include a limited amount of access to the firm's lawyers on questions the system cannot answer and the firm would keep track of such questions in order to update the system.

114. The existence of markets for such systems points to the relationship between automation and proportionality, discussed below. See *infra* Part II.C.5.

115. See, e.g., LANE COOPER, *THE RHETORIC OF ARISTOTLE* (1st ed. 1960); RICHARD DU CANN, *THE ART OF THE ADVOCATE: AN INCISIVE EXAMINATION OF THE ROLE OF THE ADVOCATE IN TODAY'S LEGAL SYSTEM* 156 (1980); PAM WRIGHT & PETE WRIGHT, *FROM EMOTIONS TO ADVOCACY* (2003).

116. John C. Shepherd & Jordan B. Cherrick, *Advocacy and Emotion*, 3 J. ASS'N LEGAL WRITING DIRECTORS 152, 152 (2006) (originally published as 138 F.R.D. 619 (1991)).

117. *Id.* at 153.

TABLE 2
PERCENT OF INVOICED HOURS SPENT ON VARIOUS TASKS, GROUPED BY
ESTIMATED EXTENT OF COMPUTER PENETRATION

Task	Tier One Firms	Tier Two–Five Firms
Strong Employment Effects	4.1%	3.6%
Document Review	4.1%	3.6%
Moderate Employment Effects	39.7%	40.4%
Case Administration and Management	3.7%	5.6%
Document Drafting	5.0%	4.0%
Due Diligence	2.0%	3.4%
Legal Research	0.5%	0.4%
Legal Analysis and Strategy	28.5%	27.0%
Light Employment Effects	56.0%	55.7%
Document Management	0.4%	0.7%
Fact Investigation	9.2%	9.6%
Legal Writing	11.4%	17.7%
Advising Clients	9.3%	3.2%
Other Communications/ Interactions	8.8%	5.0%
Court Appearances and Preparation	13.9%	14.5%
Negotiation	3.0%	5.0%
Totals**	99.8%	99.7%

** Percentages may not sum to 100% due to rounding.

Note that only 4.1 percent of lawyers' time at Tier 1 firms, and 3.6 percent of time at Tier 2–5 firms was billed to tasks where automation potentially has strong employment effects. One could argue that these low percentages reflect the impact that predictive coding has already had in automating document review. However, predictive coding was not widely used until it was officially blessed by a federal judge in 2012, and in Sky Analytics data for 2012, lawyers at Tier 1 firms billed only six percent of their time to document review.

More likely, the low percentage is explained by two factors. First, document review in our typology covers only discovery practice, not due diligence (which, as described above, is harder to automate). Accordingly, associates in departments other than litigation would not devote any of their time to the task. Second, as noted earlier, clients have been pressuring law firms for over a decade to hold down litigation costs through outsourcing, offshoring, and using contract

attorneys to perform document review.¹¹⁸ These pressures intensified following the 2008 financial collapse, when a shift in supply versus demand empowered clients to insist on cost-cutting measures, including outsourcing and the exclusion of junior associates from their matters.¹¹⁹ Thus, the task may already have been pushed out of the domain of firm lawyers' work by 2012.¹²⁰ This would be repeating a pattern seen in other settings where the most routine tasks are initially outsourced and eventually automated.¹²¹

C. MACHINE COMPLEXITY VERSUS TASK COMPLEXITY

To develop a better sense of the relationship between the difficulty to automate a task and the difficulty for a human lawyer to perform the task, we show in Table 3 the distribution of hours spent on tasks in large law firms (employment = 1000+). Large law firms employ only a small fraction of all lawyers, but as with Adam Smith's pin factory, a large firm allows lawyers in different positions to specialize in different tasks (whereas a solo practitioner or small firm lawyer must perform all tasks). The economist's assumption of profit maximization suggests the law firm will assign a task to the least expensive lawyer who can perform it at an acceptable level. Thus, assignment of tasks within the large firm provides insight on how law firms rank the complexity of tasks, with the simplest tasks performed by the least experienced lawyers and the most complicated tasks performed by the most experienced ones.

TABLE 3
DISTRIBUTION OF TIME ON TASKS BY TENURE IN TIER ONE FIRMS

	Associates ≤ 2 Years	Associates >2 Years	All Partners	Tier One Total
Strong Employment Effects	8.5%	4.5%	1.1%	4.1%
Document Review	8.5%	4.5%	1.1%	4.1%
Moderate Employment Effects	34.9%	38.5%	44.7%	39.0%
Case Administration/Management	3.4%	2.4%	6.0%	3.7%
Document Drafting	4.4%	5.4%	4.8%	5.0%

118. See *supra* note 12.

119. See *A Less Gilded Future*, *ECONOMIST* (May 5, 2011), <http://www.economist.com/node/18651114> [<https://perma.cc/HQ37-ULHL>]; Eli Wald, *Foreword: The Great Recession and the Legal Profession*, 78 *FORDHAM L. REV.* 2051 (2010).

120. See, e.g., Ashby Jones & Joseph Palazzolo, *What's A First-Year Lawyer Worth?*, *WALL ST. J.*, Oct. 17, 2011 (reporting that twenty percent of corporate legal departments insist that no first- or second-year attorneys work on their matters).

121. For example, transcription of physicians' dictated reports was first done by U.S. secretaries. It later shifted to secretarial services in the Philippines and other offshore locations. It is now largely done by automatic speech recognition.

	Associates ≅ 2 Years	Associates >2 Years	All Partners	Tier One Total
Due Diligence	2.0%	1.6%	2.7%	2.0%
Legal Research	1.6%	0.4%	0.1%	0.5%
Legal Analysis and Strategy	23.5%	28.7%	31.1%	28.5%
Light Employment Effects	56.3%	56.6%	54.2%	56.0%
Document Management	0.7%	0.5%	0.2%	0.4%
Fact Investigation	13.9%	9.2%	6.7%	9.2%
Legal Writing	10.1%	12.5%	9.5%	11.4%
Advising Clients	8.3%	6.2%	14.8%	9.3%
Communications and Interactions	9.0%	11.1%	5.1%	8.8%
Court Appearances	12.0%	14.7%	13.8%	13.9%
Negotiation	2.4%	2.3%	4.2%	3.0%
Totals**	99.7%	99.5%	100%	99.1%
Addendum: % of Hours Billed by Tenure	18.0%	50.0%	32.0%	100.0%

** Percentages may not sum to 100% due to rounding.

Table 3 reveals the absence of a strong association between the ease of automating a task (machine complexity) and whether the task is performed by a junior associate, a senior associate, or a partner (task complexity as viewed by the firm). Some data point toward a connection. Document review is heavily computerized, and when it is performed by firm lawyers (as opposed to contract attorneys), it is largely performed by junior associates. Advising clients is difficult to computerize and much of it is performed by partners. If these were the only two data points, they would suggest that tasks with the lowest machine complexity are assigned to the least experienced lawyers, and tasks with the highest machine complexity are assigned to the most experienced lawyers. This, in turn, would confirm the conventional wisdom that computers are having their greatest impact on the lowest level of lawyers within a firm. But the actual pattern is far less neat. The tasks of fact investigation and communication/ interactions both have minimal computer penetration, and yet junior associates spend a greater percentage of their time on both tasks than do partners.

The factor that undermines a simple relationship between machine complexity and position within a firm is unstructured human interaction, a skill that has so far resisted automation but that is a part of lawyering tasks at every level.¹²² The task of advising clients may require more experience than fact investigation, but both

122. Remus, *Reconstructing Professionalism*, *supra* note 10, at 33–34.

require an ability to conduct unstructured communication with other people—something junior associates and partners can do but computers cannot—which, in turn, illustrates that despite massive amounts of computing power, many tasks that are easy for humans are exceedingly difficult for computers.¹²³

D. ESTIMATING EMPLOYMENT IMPACTS

Estimating employment impacts requires translating “Strong,” “Moderate,” and “Light” employment effects into percentage reductions in lawyers’ hours—an exercise that is imprecise at best.¹²⁴ We nevertheless construct estimates by combining judgments based on interviews with two examples from among a limited set of studies on the effect of automation on other occupations.¹²⁵ These studies, which we describe in more detail in the Appendix, focus on computers’ impacts on employee productivity (output per hour of labor). When the volume of work is assumed constant (the partial equilibrium calculation), a five percent gain in output per hour of labor results in a five percent reduction of work. We also assume that the quality of lawyers’ work remains constant—that lawyers use technology to produce a constant product in less time rather than an improved product with no reduction in time. As noted above, industry surveys and data from the U.S. Bureau of Labor Statistics suggest that at present, a constant volume of work is a realistic assumption.¹²⁶ Nonetheless, given the necessary imprecision of our estimates, some readers may find the calculation unhelpful. We offer it as a step in making more tangible the frequent predictions of computers displacing lawyer labor.

Strong Employment Effects: Only one legal technology systematically falls within this category—automated document review, which was Markoff’s origi-

123. This proposition explains why automation has historically had its greatest effect on “mid-skilled” jobs, such as assembly line and clerical work. It has had much less of an effect on the lowest wage jobs because those jobs involve both unstructured human interaction and unstructured physical movement. See Autor et al., *supra* note 22.

124. Among white-collar occupations, a cause of imprecision is the lack of good output measures that would allow measuring changes in employment holding output constant. Among blue-collar occupations, a cause of imprecision is the overlap between jobs that are being automated and jobs that are being sent offshore. Frank Levy & Richard J. Murnane, *How Computerized Work and Offshoring Shape Human Skill Demands*, in *LEARNING IN THE GLOBAL ERA: INTERNATIONAL PERSPECTIVES ON GLOBALIZATION AND EDUCATION* (Marcelo M. Suarez-Orozco ed., 2007).

125. Autor et al., *supra* note 22, at 432; Sinan Aral et al., *Information, Technology and Information Worker Productivity: Task Level Evidence* (Nat’l Bureau of Econ. Research Working Paper No. 13172, 2007), <http://www.nber.org/papers/w13172.pdf> [<https://perma.cc/EMN6-Q9QD>]; Susan Athey & Scott Stern, *The Impact of Information Technology on Emergency Health Care Outcomes*, 33 *RAND J. ECON.* 399 (2002); Ann Bartel et al., *How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills*, 122 *Q.J. ECON.* 1721 (2007); Julia Adler-Milstein & R. Robert S. Huckman, *The Impact of Electronic Health Record Use on Physician Productivity*, 19 *AM. J. MANAGED CARE* 345 (2013).

126. See *supra* note 39.

nal example.¹²⁷ Automated document review continues to require senior lawyer time to train the software and review the results, and is not efficient for small classification problems. Nonetheless, to avoid underestimating automation's employment impacts, we assume that automated document review for discovery replaces eighty-five percent of all lawyer hours currently assigned to this task.

In theory, online dispute resolution and expert systems could also fall into the category of heavy employment effects given that when used, they entirely replace lawyers (and in the case of online dispute resolution, judges as well). Their impact on lawyer employment may be significant in the future, but we estimate it is minimal at present. As discussed above, use of online dispute resolution programs is currently limited, and the disputes that these programs resolve are generally small stakes e-commerce issues for which it would not be economically feasible to hire a lawyer and litigate (such that lawyer labor is not being replaced).¹²⁸ Similarly, expert systems are generally directed at reaching new markets rather than improving efficiency in existing tasks. For example, a law firm might offer subscriptions to an expert system covering aspects of tax compliance to clients who otherwise might not consult a lawyer,¹²⁹ increasing the firm's business but not displacing any lawyers.¹³⁰

Moderate Employment Effects: Moderate employment effects arise when a largely unstructured legal task has a significant structured component that can be computerized—for example, a computer-aided precedent search, the structured part of due diligence, or the question answering components of legal research. To calibrate the employment impact of this level of innovation, we refer to a case study of search-related innovation in exceptions processing at a large bank, discussed in the Appendix, and estimate that lawyering tasks in which computers have a Moderate Employment Effect reduce lawyer time devoted to those tasks by nineteen percent.¹³¹

Light Employment Effects: This category encompasses Fact Investigation, Legal Writing (as distinct from Legal Drafting), Advising Clients, Communications/Interactions, Court Appearances, and Negotiation.¹³² These are tasks that entail largely unstructured work with limited room for automation.

To calibrate Light Employment Effects, we use a case study of a limited computer innovation in healthcare, discussed in the Appendix, that concluded that one standard deviation in the use of an EMR increases clinician productivity by five percent. Based on that, we posit that adopting a computer innovation with

127. See Markoff, *supra* note 1.

128. See *supra* notes 106–09 and accompanying text.

129. See *supra* note 113 and accompanying text.

130. The existence of markets for such systems points to the relationship between automation and proportionality, discussed below. See *infra* Part II.C.5.

131. See Autor et al., *supra* note 22, at 437.

132. Light employment effects also arise in tasks relating to document management. Because these tasks are usually performed by clerical staff, automation does not affect lawyer employment per se.

Light Employment Effects would decrease required lawyer employment for a given task by five percent.

This leaves one set of technologies for which even the roughest estimation of employment impacts is exceedingly difficult—document templates sold directly to the public by firms like LegalZoom and Rocket Lawyer. Many commentators believe these templates will fully eliminate many lawyers' jobs.¹³³ There is reason to question such assertions, as it is not at all clear whether these services are tapping into a latent market of previously unserved individuals or taking business away from lawyers. LegalZoom representatives argue that it is overwhelmingly the former—that they serve individuals who would not otherwise have gone to a lawyer¹³⁴—but it is of course in their interests to frame their business model as non-threatening to lawyers.

Indirect evidence of LegalZoom's impact on market share comes from its 2012 decision to table its planned initial public offering after receiving insufficient interest from the markets.¹³⁵ Since that time, there is reason to think that LegalZoom's business has not grown as rapidly as it had projected.¹³⁶ This may be the result of regulatory responses from unauthorized practice of law committees, a topic that we address below. But regardless of cause, the slowed growth offers reason to question sweeping conclusions about massive lawyer displacement. Because of all of these uncertainties, and because any impact will be felt primarily by solo practitioners or small firm lawyers, we do not separately account for the impact of document templates marketed directly to the public. Instead we include it under the general heading of document drafting—a task with moderate computer penetration, for which relevant technologies tend to replace parts but not all of a lawyer's job.

To summarize, our illustrative calculation rests on three estimates:

- Tasks where computer technology has a strong employment effect experience an eighty-five percent reduction in employment.

133. See, e.g., McGinnis & Pearce, *supra* note 2, at 3065–66; Barton, *The Lawyer's Monopoly*, *supra* note 5, at 3068; see also, e.g., BARTON, GLASS HALF FULL, *supra* note 5.

134. Telephone Interview with Eddie Hartman, Chief Prod. Officer, LegalZoom (Sept. 18, 2015).

135. Olivia Oran, *UPDATE 1—LegalZoom IPO Delayed—Source*, REUTERS (Aug. 2, 2012), <http://www.reuters.com/article/2012/08/02/legalzoom-idUSL2E8J2EZF20120802> [<https://perma.cc/RNN7-38D7>]. For LegalZoom's filed Form S-1, see LegalZoom.com, Inc., Registration Statement (Form S-1) (May 10, 2012).

136. See, e.g., Michael Carney, *The \$425M LegalZoom Deal Is a Win for VCs, but Less Exciting for the Company or LA*, PANDO (Jan. 6, 2014), <https://pando.com/2014/01/06/the-legalzoom-deal-is-a-win-for-vc-but-less-exciting-for-the-company-or-la/> [<https://perma.cc/QDD4-W3HY>] (describing a \$200 million investment by private equity firm Permira as “a bit ‘meh,’” and the company's \$425 million valuation in the deal as “a slight downgrade from the \$500 million-plus valuation . . . the company was most recently hoping to attract in the public markets [as part of the proposed IPO]”); *Has LegalZoom Lost its Bloom?*, THEFORMTOOL (Jan. 6, 2013), <http://www.theformtool.com/has-legalzoom-los-its-bloom/> [<https://perma.cc/8N5X-H3FG>] (describing that prior to the proposed IPO, “LegalZoom was already experiencing the chill of a slowing growth rate and tighter margins in its traditional market, legal forms for sale.”).

- Tasks where computer technology has a moderate employment effect experience a nineteen percent reduction in employment.
- Tasks where computer technology has a light employment effect experience a five percent reduction in employment.

To calculate an overall employment impact, we apply these percentages to the lawyers' use of time in 2014 in Tier 1 firms (Table 3).¹³⁷ Table 3 reflects our use of a partial equilibrium calculation, which makes our argument transparent but requires two strong assumptions: (i) that no employment impacts from these technologies have occurred previously, and (ii) that the level of work remains constant. If all the technology above were implemented at one time, it would result in an estimated thirteen percent reduction in hours.¹³⁸ Since law firms have a well-established reputation for slow technology adoption, however, we assume more realistically that the technology is adopted over a period of five years.¹³⁹ Again, assuming a constant volume of legal work, our estimated employment loss spread over five years would indicate that demand for lawyers' hours is decreasing by 2.5 percent per year because lawyer productivity is increasing by 2.5 percent per year.¹⁴⁰ In considering this estimate, we note that labor productivity increasing by 2.5 percent per year is an impressive number; labor productivity across the U.S. non-farm business sector has averaged slightly less than 1.5 percent growth per year for the last ten years.¹⁴¹

To summarize, it is frequently argued in popular writing on artificial intelligence that the automation of legal work causes weakness in the market for lawyers. Our estimates indicate that the argument is overstated and that a more important cause is a basic imbalance between supply and demand. Interviews suggest that by 2004, significant numbers of contract lawyers could be hired, at a much cheaper rate than law firm associates, to classify the huge volume of digital documents that had become part of discovery proceedings. In 2009, National Law Journal (NLJ) 250 firms laid off 5,259 attorneys—about four percent of all

137. Recall that the distribution of time on task in Tier 2–Tier 5 firms is similar to the distribution in Tier 1 firms.

138. Our job loss estimate will be low if solo practitioners spend much of their time on non-adversarial, formulaic issues that could be replaced by templates sold directly to individuals. We may also be underestimating predictive coding's impact on contract lawyer employment, but interviews with industry professionals and corporate counsels in charge of discovery suggest this is not the case. They argue that predictive coding has replaced contract lawyers in some, but not all, parts of the discovery document classification while the volume of discovery work has grown enough to maintain employment levels.

139. The slow adoption of legal technology was emphasized in many of our interviews. We explain why adoption may start growing more rapidly in Part II.

140. Alternatively, the volume of legal services would have to increase by at least 2.5 percent per year to offset automation's impact in reducing demand for lawyers' services.

141. *See, e.g.*, EXEC. OFFICE OF THE PRESIDENT COUNCIL OF ECON. ADVISERS, ECONOMIC REPORT OF THE PRESIDENT app. B, tbl.B-16 (2015).

NLJ 250 attorneys (including 8.7% of NLJ 250 associates).¹⁴² If we consider the beginning of the legal artificial intelligence age the 2012 *Da Silva Moore* decision affirming technology-assisted review, then we can say computerized work has been one of many drags on a generally weak market.

Our calculations rested on a number of assumptions, however, which admittedly narrow the inquiry and simplify reality. In the next part, as we look ahead to future developments, we broaden our focus. We acknowledge that the demand for legal work will not stay constant as new technologies are adopted, that technology may improve the quality as opposed to efficiency of legal work, and that professional regulation will play a key role in steering the direction of technological advances.

II. LEGAL PROFESSIONALISM IN THE DIGITAL AGE

Part I offered estimates of current employment impacts of existing and emerging legal technologies. In this Part, through a series of three questions, we broaden our focus to consider the direction of longer-term technological development and the role of regulation in conditioning that development. First, we ask how legal technologies—specifically, artificial intelligence applications that potentially perform work now performed by lawyers¹⁴³—would likely develop and expand if market forces operate freely. Of course, the market for legal services does not operate freely; it is highly regulated, with significant repercussions for the development of legal technologies. We therefore also examine existing regulatory structures, showing that they are inadequate to address the challenges and opportunities of new technologies. Third and finally, we argue that notwithstanding deep problems with existing approaches to regulation, unimpeded market forces will have deleterious effects while professional norms and regulations have continuing value. We conclude by addressing the challenge of designing regulatory structures that protect professional values without excessively impeding the development and adoption of, and access to, new legal technologies.

A. THE MARKET FOR NEW TECHNOLOGIES

The likely path for legal artificial intelligence will be shaped by two propositions discussed in Part I:

- For a computer to automate a lawyer's task, it must be possible to model the lawyer's information processing in a set of instructions; and

142. Leigh Jones, *So Long, Farewell; There's No Sugar Coating It: This Was the Worst Year Ever for Lawyer Headcount*, NAT'L L.J., Nov. 9, 2009.

143. Law firms also will be involved with other software—in particular, data security and cloud applications—that do not have obvious implications for employment.

- Models estimated by machine learning have difficulty processing contingencies that differ significantly from the data on which the models were trained.

The first proposition limits legal applications to structured tasks—tasks for which information processing follows an underlying pattern (which may be uncovered by machine learning). The second proposition encourages developers to focus the machine learning on relatively narrow tasks—tasks where it is feasible to train the model on most of the contingencies it is likely to confront (though this training may occur over time as the model is used).¹⁴⁴

These restrictions are apparent in the progress of legal artificial intelligence to date: major inroads in document classification in discovery (the subject of Markoff's original article¹⁴⁵) and developing inroads in organizing, drafting, and reviewing contracts for due diligence. Other applications are at more embryonic stages: question answering systems using highly trained databases; improved search and information extraction tools to locate particular information within the enterprise or within a particular body of documents; data-based tools for cost and risk analysis and prediction; and expert systems that offer a platform to provide pre-packaged legal advice.

With the exception of expert systems, these applications all involve language processing—in particular, measuring the similarity of meaning between two documents (or pieces of text). By focusing on comparisons between two documents, the software sidesteps problems in inferring meaning in less structured situations—for example, interpreting a client letter that uses common sense reasoning as part of its language.¹⁴⁶ The difficulty in solving these problems helps to explain why artificial intelligence has not, to this point, penetrated those lawyers' tasks that require unstructured communication.

Despite these limitations, artificial intelligence will both extend existing applications and address other structured tasks. An example of an extension begins with Vern Walker's argument, noted earlier, that an ultimate goal of legal search tools is to locate legal arguments rather than entire cases or text passages.¹⁴⁷ Currently, the advanced applications in this area, such as IBM's Debater System, are based on a document corpus where claims and evidence have been annotated by humans. Improvements in natural language processing

144. Absent from this description is the human ability to draw correct inferences from small amounts of information, a key element in the flexibility of human cognition.

145. See Markoff, *supra* note 1.

146. Common sense refers to the large set of facts that people apply without thinking but computers don't recognize unless they are programmed to do so—for example, a screwdriver dropped from the hand falls down rather than up or sideways. See Ernest Davis, *How to Write Science Questions that Are Easy for People and Hard for Computers*, AI MAG., Spring 2016, at 13.

147. See *supra* note 76 and accompanying text.

are making some progress in identifying claims and evidence automatically.¹⁴⁸

Similarly, several applications now offer advanced proof reading of documents including identifying inconsistent use of terms, improper citation formats, and other features beyond basic spelling and punctuation.¹⁴⁹ As with predictive coding and contract review, these applications have the potential to save lawyer time. One exception to this trend is the much-discussed development of blockchain-based contracts,¹⁵⁰ which appear, at least for now, to be confined to heavily repeated trades of financial instruments where the work involved was previously performed by back office personnel rather than lawyers.¹⁵¹

Improved applications are one way to increase the reach of legal artificial intelligence. An alternative approach is to simplify the task so that computers can perform currently complex tasks with existing software. Kingsley Martin, a developer of contract review software, envisions the development of “auditable contracts”—contracts written in English that are simple enough to be parsed by a computer.¹⁵² Such contracts could, in theory, substantially improve automated contract review. As another example, Modria, the online dispute resolution program, simplifies the task of negotiation by gathering information from each side, summarizing areas of agreement and disagreement, and, based on that information, presenting proposed resolutions.¹⁵³ The program also has mediation and arbitration modules if the parties cannot agree to one of the proposed resolutions, but the company claims that the “vast majority” of claims are settled just by laying out the facts and proposing solutions based on areas of agreement.¹⁵⁴

In some cases, a task can be simplified or standardized without altering its meaning. For example, an individual orders a book from Amazon by clicking on an icon rather than writing a free text email (with potential mistakes), which would be hard to automatically interpret. In other cases, however, part of a task will be lost in the simplification. Recall that early versions of Westlaw and Lexis

148. Ruty Rinott et al., *Show Me Your Evidence—An Automatic Method for Context Dependent Evidence Detection*, at 440 (Ass’n for Computational Linguistics, Proc. of the 2015 Conf. on Empirical Methods in Nat. Language Processing, Sept. 2015), http://aclweb.org/anthology/D15-1050?cm_mc_uid=73807758537314226464001&cm_mc_sid_50200000=1473173993 [<https://perma.cc/YV4T-K3EZ>].

149. See, e.g., *Contract Companion*, MICROSYSTEMS, <http://www.microsystems.com/contractcompanion> [<https://perma.cc/47LR-LZYB>] (last visited Apr. 2, 2017).

150. A blockchain contract rests on a computer protocol that specifies the terms of an agreement—payments to be made when specified conditions are satisfied and so on. By embedding the protocol in a blockchain—a distributed ledger—the contract is protected from manipulation by either party and the parties do not need a bank or other trusted intermediary to enforce contract terms. As noted earlier, to be written as a computer protocol, a contract must involve only well-defined elements.

151. See, e.g., Atsushi Santo et al., *Applicability of Distributed Ledger Technology to Capital Market Infrastructure*, at 5–6 (Japan Exch. Grp., Working Paper Vol. 15, 2016), http://www.jpx.co.jp/english/corporate/research-study/working-paper/b5b4pj000000i468-att/E_JPX_working_paper_No15.pdf [<https://perma.cc/9E32-WT5H>].

152. Telephone Interview with Kingsley Martin (Aug. 2017).

153. See MODRIA, <http://modria.com/> [<https://perma.cc/9EKG-S5JD>] (last visited Mar. 2, 2017).

154. *Id.*

simplified the task of searching for legal concepts and arguments by allowing users to look for particular words or combinations of words. Simplifying the search task to a keyword search distorted the results, leading Westlaw and Lexis to reintroduce broader approaches to legal research, such as through headnotes and indexes.

In addition to impacting technology development, the market will of course influence adoption. Historically, law firms have resisted new technologies of most kinds, but for reasons that now might be changing. As long as clients were willing to pay on the basis of billable hours, the need for technology to increase efficiency was not an imperative. The partnership structure of many law firms further increased resistance because technology costs come directly from partners' profits.¹⁵⁵ A third source of resistance was the well-documented distaste that many lawyers have for technology and "mathiness" of any kind.¹⁵⁶

While much of the legal services market continues to use billable hours, client pressure on hourly rates and total hours is likely to continue to intensify, reflecting the continuing imbalance between the supply of lawyers and the demand for legal services. In Am Law 100 and 200 firms surveyed by the Thomson Reuters *Peer Monitor*, total billable hours have shown virtually no growth since 2010, while the number of lawyers across all U.S. law firms has grown by one to two percent per year.¹⁵⁷ As pressure increases to hold down expenses, the purchase of technology becomes more attractive.

A second factor promoting accelerated technology purchases is the shift of corporate legal work from law firms to the corporation's own legal department. Such legal departments are part of standard, profit-maximizing organizations that put a higher value on efficiency than a partner-based firm.¹⁵⁸

A third but more speculative factor promoting accelerated technology adoption may be Clay Christensen's theory of disruptive innovation.¹⁵⁹ Initially, Big Law

155. We thank Bruce Elvin for this point. In a traditional corporation, the cost of technology is borne by shareholders.

156. See, e.g., Susan Moon, *Moonlighting: Things Not to Say In-House—'I'm Bad at Math'*, ABOVEHELAW (Jan. 6, 2012), <http://abovethelaw.com/2012/01/moonlighting-things-not-to-say-in-house-im-bad-at-math/> [<https://perma.cc/V6R5-RAAN>].

157. See 2016 LEGAL MARKET REPORT, *supra* note 39, at 4, chart 3.

158. See, e.g., Deborah A. DeMott, *The Discrete Roles of General Counsel*, 74 *FORDHAM L. REV.* 955, 958–60 (2005); Robert L. Nelson & Laura B. Nielsen, *Cops, Counsel, and Entrepreneurs: Constructing the Role of Inside Counsel in Large Corporations*, 34 *LAW & SOC'Y REV.* 457, 466–68 (2000).

159. "Disruptive Innovation" is Clayton Christensen's theory of the impact of technological innovation on markets. CLAYTON M. CHRISTENSEN, *THE INNOVATOR'S DILEMMA* (2013). Christensen argues that companies (here, law firms) at the top of the market routinely ignore disruptive technologies in their early days because such technologies focus on the bottom of the market and do not constitute competition. Such companies may even abandon low margin work to these new technologies so as to focus on higher margin work. But the technology producers and vendors that initially targeted the lower end of the market eventually improve their products and services to address higher margin work. Eventually, what appeared to be the best short-term strategy for firms at the top of the market (ignore the new technologies) turns out to enable competition and displacement.

did not appear at all concerned about losing routine work to insourcing, outsourcing, or automation, because it was work for which clients were already (and from firms' perspective, problematically) demanding lowered fees and alternative fee arrangements.¹⁶⁰ But accepting and performing the routine work allowed software developers and legal technology firms to develop approaches and solutions to more complicated and profitable work, which Big Law may have little choice but to adopt under client pressure. Ben Barton and others have predicted that this pattern may repeat on a more dramatic scale with online legal service providers like LegalZoom and Rocket Law, first displacing solo practitioners and small firms, and gradually disrupting the entire law firm model.¹⁶¹ Regardless of whether Christensen's theory plays out in full, it seems likely that technology solutions at the bottom of the market will push change throughout the market.

All of this said, the implications of an acceleration in technology adoption for legal employment are not clear. As we have seen, some kinds of software may primarily address new markets—LegalZoom, for example. Other kinds of software represent new kinds of service that may expand rather than reduce the need for lawyers—for example, expert systems and prediction analytics to measure risk. In addition, the context of the work matters. When applied to pro forma activities—legal tasks that can be completed by applying a fixed amount of effort to anticipated contingencies—artificial intelligence substituting for a lawyer's time will likely reduce the demand for lawyers. Examples include incorporating a new business, writing a will, or ensuring internal compliance with a particular statutory scheme. In adversarial activities, in contrast, a party's effort is generally determined in part by the efforts of the other party, where the other party has strong incentives to develop *unanticipated* contingencies. Here, it is possible that the two parties settle into an arms race in which they use the time that artificial intelligence frees up for other activities.

Accordingly, the likely trajectory of legal technologies in an unregulated market would be determined by factors affecting both development and adoption. We believe the pace of development would depend largely on advances in natural language processing while the pace of adoption would depend on client pressures.

B. CURRENT APPROACHES TO REGULATION

Of course, the market for legal services is far from unregulated. Unauthorized Practice of Law (UPL) rules limit the provision of legal services to individuals

160. See, e.g., *Events: Disruptive Innovation in the Market for Legal Services*, HARV. LAW SCH. CTR. ON THE LEGAL PROF. (Mar. 6, 2014), <https://clp.law.harvard.edu/event-post/disruptive-innovation-in-the-market-for-legal-services/> [<https://perma.cc/2QZY-GHMQ>].

161. Barton acknowledges, however, that bespoke “bet the company” work will remain an exclusive domain of law firms.

who are trained and licensed to practice law.¹⁶² State supreme courts and bar committees then discipline lawyers who fail to adhere to the rules of practice and ethical codes. We see a continuing need for, and value in, professional regulation. However, for at least four reasons, the UPL rules—the principal way in which the profession currently addresses new technologies—provide an unhelpful approach.

First, courts following this approach have posed for themselves an unanswerable question. UPL rules seek to distinguish tasks that can only be performed by trained and licensed lawyers from tasks that lay people, lacking the same training and ethical regulation, can nevertheless provide competently, reliably, and ethically. Courts have applied this framework to new technologies by asking whether a given technology (generally an online service provider) is more similar to a scrivener who completes a form by merely recording the information a customer relays (in which case, the technology would not constitute UPL) or to a service provider who aids in selecting and properly completing a form (in which case, it would constitute UPL).¹⁶³ Neither alternative is ever clearly right or wrong.¹⁶⁴ An online legal forms provider can be viewed as the functional equivalent of a mere scrivener insofar as it is the user him or herself who enters the relevant information via the online questionnaire and completes the form,¹⁶⁵ or the provider can be viewed as the functional equivalent of a human service provider exercising judgment insofar as the software is programmed with deductive rules to ask the user a series of questions and, based on the answers, complete the appropriate document.¹⁶⁶ Accordingly, courts are left making normative decisions with little guidance from the framework.

Second, analogizing to human approaches fails to appreciate that which is unique and distinct about legal technologies. Computers can be trained in ways that avoid human error such that we may be comfortable with a computer performing tasks we would not want performed by an untrained and potentially

162. The principal justification prohibiting unauthorized practice of law is “to protect the public from the consequences of receiving legal services from unqualified persons.” ANNOTATED MODEL RULES OF PROF’L CONDUCT R. 5.5 (2015) [hereinafter ANN. MODEL RULES] (Rationale for Proscription Against Unauthorized Practice of Law) (“The proscriptions are also aimed at facilitating the regulation of the legal profession and protecting the integrity of the judicial system.”).

163. *See, e.g.,* *Janson v. LegalZoom.com, Inc.*, 802 F. Supp. 2d 1053, 1059 (W.D. Mo. 2011) (“Plaintiffs urge the Court to follow the cases . . . which generally involve businesses providing a legal document preparation service for their customers. . . . Defendant LegalZoom argues that its website providing access to online document assembly software is the functional equivalent of [a] ‘do-it-yourself’ divorce kit . . .”).

164. The court in *Janson* even acknowledged this, but declined to revise its analysis accordingly. *See id.* at 1063 (“None of the . . . cases cited by the parties are directly on point, due to the novelty of the technology at issue here.”).

165. *See id.* (noting LegalZoom’s argument that “its customers—rather than LegalZoom itself—complete the standardized legal documents by entering their information via the online questionnaire to fill the document’s blanks”).

166. *See id.* at 1055 (observing that LegalZoom reassures consumers that “we’ll prepare your legal documents,” and that “LegalZoom takes over” once customers “answer a few simple online questions”).

unreliable human.¹⁶⁷ And yet, reducing a lawyering task to a set of computer-implementable rules may over-simplify, ignore complexity, or create opportunities for error that are not immediately apparent.¹⁶⁸ We therefore may not want a computer performing particular tasks in all contexts, notwithstanding effective performance in one context.

A third problem stems from the poor fit between the UPL inquiry and technologies that lawyers use in representing clients (as opposed to those that are marketed directly to the public). Concluding that a non-lawyer cannot competently and reliably perform a particular task does not establish that a computer cannot help a lawyer do so. Perhaps for this reason, some commentators suggest that technologies that lawyers use and oversee are best addressed through the rules of lawyer oversight of non-lawyer service providers.¹⁶⁹ Applied to new technologies, these rules would permit adoption of new technologies where lawyers supervise their use and accept responsibility for their results. At least for now, however, few lawyers are sufficiently knowledgeable to oversee new legal technologies in a meaningful way.¹⁷⁰ Moreover, this approach suggests that computerizing all of a lawyer's functions would be permissible with oversight. But surely some tasks, such as in-court advocacy and settlement or plea negotiations, cannot and should not be delegated to a computer.

A fourth and final problem is that UPL prosecutions often appear to be self-interested efforts by the bar to protect its monopoly.¹⁷¹ Scholars and commentators have long argued that non-lawyers can perform certain aspects of

167. The *Janson* court ignored this, resting entirely on a formalistic UPL analysis. *Id.* at 1064 (“Because those that provide [LegalZoom’s] service are not authorized to practice law in Missouri, there is a clear risk of the public being served in legal matters by ‘incompetent or unreliable persons.’”).

168. *See, e.g., supra* note 77 and accompanying text.

169. These rules, developed to address the outsourcing of work to non-lawyers or offshoring of work to foreign lawyers, provide that such activities are ethically permissible so long as the lawyer supervises the work and retains ultimate responsibility for the result. *See, e.g.,* ABA Comm’n on Ethics & Prof’l Responsibility, Formal Op. 08-451 (2008) (“A lawyer may outsource legal or nonlegal support services provided the lawyer remains ultimately responsible for rendering competent legal services to the client under Model Rule 1.1.”); *see also* Prof’l Ethics of the Fla. Bar, Op. 07-2 (2008) (approving of off-shore outsourcing); Ass’n of the Bar of the City of N.Y. Comm. on Prof’l & Judicial Ethics, Formal Op. 2006-3 (2006) (providing that a lawyer may outsource legal support services to overseas lawyers and non-lawyers if the lawyer supervises the work rigorously).

170. A comment to Model Rule 1.1 advises lawyers of a professional duty to stay abreast of technological advances, *see* MODEL RULES OF PROF’L CONDUCT R. 1.1 cmt. 5 (2016) [hereinafter MODEL RULES], but this provision has little teeth given the vagueness of its standard and its location in the comments rather than in an enforceable rule. Moreover, lawyers’ generally low level of technical competency is reinforced by other provisions of the *Model Rules*, which prescribe a reduced level of required oversight for automated legal work. *See, e.g.,* MODEL RULES R. 5.3 cmt. 4.

171. *See, e.g.,* DEBORAH L. RHODE, *THE TROUBLE WITH LAWYERS* ch. 5 (2015); Benjamin H. Barton, *Why Do We Regulate Lawyers?: An Economic Analysis of the Justifications for Entry and Conduct Regulation*, 33 ARIZ. ST. L.J. 429, 457 (2001) [hereinafter Barton, *Why Regulate Lawyers?*]; Roger C. Cramton, *Delivery of Legal Services to Ordinary Americans*, 44 CASE W. RES. L. REV. 531, 615 (1994); Deborah Rhode, *Access to Justice: Connecting Principles to Practice*, 17 GEO. J. LEGAL ETHICS 369, 371–72, 409 (2004) [hereinafter Rhode, *Access to Justice*]; Deborah L. Rhode, *Professionalism in Perspective: Alternative Approaches to Nonlawyer*

legal practice perfectly well, and that allowing them to do so would dramatically reduce the cost of legal services.¹⁷² This argument applies to technologies as well—if technologies can, in fact, perform aspects of legal practice as well as humans, shouldn't we use them to increase access to justice? And indeed, many courts and legal services organizations are relying on technology to expand their reach. Thus far, they have done so primarily through online filing or intake systems, which simply leverage the power and reach of the Internet.¹⁷³ Commentators advocate more advanced technologies, from automated document assembly to phone apps that give legal advice, as the only workable solution to the access to justice gap.¹⁷⁴

And yet, at least one formulation of this argument, recently adopted by the U.S. Court of Appeals for the Second Circuit, is highly problematic. In a case construing the exemption from over-time pay for individuals “employed in a bona fide . . . professional capacity” under the Fair Labor Standards Act,¹⁷⁵ the plaintiff, a contract attorney, argued that document review, defined as “us[ing] criteria developed by others to simply sort documents into different categories,” did not constitute the practice of law, such that he was not employed in a “professional capacity.”¹⁷⁶ The Second Circuit agreed, reasoning that because these were “services that a machine could have provided,” they could not possibly constitute the practice of law.¹⁷⁷

The Second Circuit's conclusion was based on a high-level generalization that computers can perform document review as well as humans. But neither the Second Circuit, nor scholars and commentators expressing similar reasoning, have undertaken the critical inquiry of whether and how a machine approaches the task differently from a human. The differences have ramifications that extend beyond lowered costs and are central to a meaningful normative and regulatory

Practice, 22 N.Y.U. REV. L. & SOC. CHANGE 701, 709–13 (1996) [hereinafter Rhode, *Professionalism in Perspective*].

172. See, e.g., Barton, *Why Regulate Lawyers?*, supra note 171, at 457; Cramton, supra note 171, at 615; Rhode, *Access to Justice*, supra note 171, at 371–72, 409; Rhode, *Professionalism in Perspective*, supra note 171, at 709–13; Laurel A. Rigertas, *Stratification of the Legal Profession: A Debate in Need of a Public Forum*, 2012 J. PROF. LAW. 79; Cristina L. Underwood, *Balancing Consumer Interests in a Digital Age: A New Approach to Regulating the Unauthorized Practice of Law*, 79 WASH. L. REV. 437, 442 (2004).

173. See, e.g., LAWHELP.ORG, <http://lawhelp.org> [<https://perma.cc/R79Z-MF4Z>] (last visited Mar. 9, 2017).

174. See, e.g., McGinnis & Pearce, supra note 2, at 3066; BARTON, GLASS HALF FULL, supra note 5; Barton, *The Lawyer's Monopoly*, supra note 5, at 3068; Rostain et al., *Thinking Like a Lawyer*, supra note 97.

175. *Lola v. Skadden, Arps, Slate, Meagher & Flom LLP*, No. 14-3845-cv, 2015 WL 4476828, at *2 (2d Cir. July 23, 2015) (reviewing an appeal from an order dismissing plaintiff's putative class action for violation of the Fair Labor Standards Act, 29 U.S.C. 201 et seq., for failing to pay overtime for document review).

176. *Id.* at *6.

177. *Id.* Plaintiff alleged that his work was closely supervised by the Defendants, and his “entire responsibility . . . consisted of (a) looking at documents to see what search terms, if any, appeared in the documents, (b) marking those documents into the categories predetermined by Defendants, and (c) at times drawing black boxes to redact portions of certain documents based on specific protocols that Defendants provided.” *Id.* at *1 (internal quotations omitted).

inquiry.¹⁷⁸ The computer's altered approach is often what makes automation attractive—it may sidestep opportunities for human error, improving accuracy and consistency. But it may also create new opportunities for error or have unintended consequences for legal practice. Access to deeply flawed and error-filled legal services cannot qualify as an acceptable, much less desirable, answer to the access to justice gap.

Accordingly, while we agree with the majority of critics and commentators who view the UPL rules as unreasonably impeding the pace of technological development and use, we do not think the answer is to jump to conclusions based on particular instances of a computer performing a task well. Nor, as we discuss next, do we think the answer is to forego all forms of professional regulation.

C. THE VALUE OF REGULATION

Critics contend that the problem is not only the UPL rules, but also all forms of professional regulation. Professional regulation, they contend, is an exercise in protectionism, limiting competition from computers and human service providers alike.¹⁷⁹ Although this critique is important and powerful, it paints only a partial picture. It fails to consider at least three essential functions of professional regulation, which are implicated by new technologies: (1) protecting consumers in the face of information asymmetries;¹⁸⁰ (2) ensuring that negative externalities do not undermine the integrity of our legal systems; and (3) ensuring universal access to legal services.¹⁸¹ We believe that the values, norms, and structures of the legal profession are necessary to address the challenges new legal technologies pose for all three of these objectives.

178. It is far from established that all legal technologies will lead to lowered costs. See Remus, *Predictive Coding*, *supra* note 47, at 1707.

179. MILTON FRIEDMAN, *CAPITALISM AND FREEDOM* 139 (1962) (arguing that the state-granted power of licensure, which takes the form of monopolistic market power and the right of self-regulation, allows licensees to advance their financial self-interest at the expense of the public interest); see also ABBOTT, *supra* note 10, at 184–86; ABEL, *supra* note 10, at 40–48; SUSSKIND & SUSSKIND, *supra* note 2, at 16–17; Gillian K. Hadfield, *The Cost of Law: Promoting Access to Justice Through the (Un)Corporate Practice of Law*, 38 INT'L REV. L. & ECON. 43, 43 (2014); Barton, *Why Regulate Lawyers?*, *supra* note 171, at 457; Rhode, *Access to Justice*, *supra* note 171, at 371–72, 409.

180. A primary and long-standing justification of professional regulation proceeds as follows: Because of the esoteric and specialized nature of legal expertise, lay people cannot adequately assess or monitor the work of lawyers. Professional licensure is therefore needed to ensure that those who provide legal services have a baseline level of competency; professional regulation is needed to ensure that if clients are harmed there are repercussions.

181. An additional justification of professional regulation addresses the ways in which clients, intentionally or not, may use lawyers to the detriment of opponents or third parties. Thus, for example, the ethical rules prohibit lawyers from aiding clients who choose to perjure themselves and in some limited circumstances allow lawyers to breach a client's confidences to protect third parties.

1. CONSUMER PROTECTION

A common refrain among legal technology advocates is that by eliminating human error, standardizing services, and lowering prices, new technologies have the potential to serve client interests far more effectively than lawyers.¹⁸² This may be true in some contexts, but it is decidedly not true in others. Moreover, many clients and lawyers alike lack sufficient understanding of new legal technologies to determine when use is appropriate and when the risk of harm or error is low or high.

Document classification offers a useful illustration of how a legal technology that eliminates error in some contexts may simultaneously create new risks of error in other contexts. A human lawyer engaging in this task examines a set of documents page-by-page to identify relevant meaning and content. Predictive coding, in contrast, identifies particular combinations of document features pursuant to statistical probabilities of relevance, with no reference to meaning. This is not a problem when the goal is to locate types of data and information that have been well specified in advance, such as in discovery practice.¹⁸³ But as we have seen, machine learning models (which are estimated statistical models) have difficulty processing contingencies that differ significantly from their training data. They therefore have difficulty processing and recognizing the unstructured aspects of due diligence that precede a corporate transaction.¹⁸⁴ As noted above, some firms are making progress in automating tasks encompassed by due diligence, but at least for the time being, their products are effective only in reviewing large volumes of similar documents or identifying similar types of clauses.

Even within discovery practice, predictive coding may create new risks of error by failing to recognize “hot documents”—documents that are highly relevant and damaging to the producing party. Individuals often change their normal writing styles, becoming particularly formalistic or vague, when they explain major decisions or acquire potential liability. As a result, hot documents, which document these decisions or events, frequently employ unusual or coded language or syntax. The most relevant documents in a case may therefore use language and tone that many predictive coding products, trained on a sample of normal documents, will fail to recognize.¹⁸⁵ Some products correct for this problem by training the computer on a sample of responsive documents rather

182. See, e.g., McGinnis & Pearce, *supra* note 2, at 3054–55; Cabral et al., *supra* note 11, at 246–56; Hornsby, *supra* note 11, at 931–34; Staudt, *supra* note 11, at 1128–34; Wolf, *supra* note 11, at 773–85.

183. See, e.g., Grossman & Cormack, *supra* note 61, at 3; Roitblat et al., *supra* note 61, at 70, 74–75.

184. See *supra* notes 68–69 and accompanying text.

185. Hofman Interview, *supra* note 68. Note that for some predictive coding products, the problem may be virtually intractable if the author reverts to coded language (i.e., “I bought two pounds of flounder at the fish market today” to indicate a completed transaction). *Id.*

than random documents¹⁸⁶ and/or by identifying responsiveness by reference to metadata as well as words.¹⁸⁷ But the user must be able to understand and to recognize the importance of selecting the best and most appropriate product. Otherwise, the machine-learning algorithm may fail to recognize a hot document, and, because it will not recognize the existence of a problem in these situations, it will not give the user a warning.

This is one example of the broader challenge posed by machine learning models—the task of determining when and how to notify a user that the computer’s “best” answer is not very good. Some predictive coding applications deal with this problem forthrightly by assigning each document an *ex ante* probability of responsiveness.¹⁸⁸ In a similar approach, software that compares documents can calculate a mathematical index of similarity and a user can specify a cutoff below which documents are judged as not sufficiently similar. But in other cases, the user is at the mercy of the programmer. On different dates, an iPhone “Siri” responded to the question, “Can a dog jump over a house?” with “I’m sorry but I don’t know the answer to that question”¹⁸⁹ or by offering a link to a child’s riddle about a dog jumping over a doghouse and a second link to an ASPCA bulletin on teaching dogs not to jump.¹⁹⁰

Our point is not that predictive coding has no value or should never be used, nor is it that unanticipated contingencies represent insurmountable hurdles to automation. We mean only to emphasize that consumer protection concerns are as salient with respect to legal technologies as they are with respect to legal services generally. In both cases, the expertise in question is complex, esoteric, and specialized; in both cases information asymmetries can quickly lead to market failure.

And yet, the answer cannot be traditional forms of professional regulation, as lawyers themselves frequently lack sufficient understanding of the technologies. As discussed, predictive coding increases the risk for error, in some cases at the same time that it decreases the risk in others, but understanding how and why, and choosing the most appropriate tool and protocol for the context, requires

186. See Grossman & Cormack, *supra* note 61 (“Random training tends to be biased in favor of commonly occurring types of relevant documents, at the expense of rare types. Nonrandom training can counter this bias by uncovering relevant examples of rare types of documents that would be unlikely to appear in a random sample.”).

187. Grossman & Cormack Interview, *supra* note 57. For example, a tool that identifies responsiveness exclusively by reference to words may not identify an email written in coded language, but a tool that also references metadata may identify it by focusing on who is emailing whom and when, and not just on the content of the message.

188. These probabilities are examples of the developing area of “explainable artificial intelligence.” See Marco Tulio Ribeiro et al., “Why Should I Trust You?”: *Explaining the Predictions of Any Classifier*, ACM SIG KDD INT’L CONF. ON KNOWLEDGE DISCOVERY & DATA MINING (KDD) (Aug. 9, 2016).

189. Question posed on Aug. 14, 2015.

190. Question posed on Nov. 27, 2015.

significant technological expertise that many lawyers lack. As we return to below, effective regulation may therefore require technical as well as legal expertise.

2. SYSTEMIC INTERESTS

The value of professional regulation lies not only in protecting clients from lawyers (or legal technologies), but also in protecting society from the ways in which clients may use lawyers (or legal technologies) to the detriment of others, including opponents, third parties, and the legal system itself. Some degree of this is built into our adversarial system—clients hire lawyers precisely in order to gain an advantage over others. But codes of conduct place limits on what lawyers can do for their clients. They ensure, for example, a baseline of fair dealing with an opponent, candor to the court, and respect for the rule of law.¹⁹¹

New legal technologies implicate these interests in important but non-obvious ways. For example, clients may be eager to use, or to have their lawyers use, legal prediction software, given that it often achieves higher levels of accuracy than human prediction.¹⁹² But if such software completely displaces lawyers, the increased accuracy may be accompanied by a number of detrimental consequences. Reducing advice to prediction would eliminate a core function of lawyering—counseling compliance with the law. If a client's only legal advice comes from a computer's prediction of how a court will likely respond, advising will be reduced to calculating what a client can get away with, instantiating the Holmesian Bad Man view of the lawyer.¹⁹³

More broadly, reducing legal advising to legal prediction could threaten to impede the law's development. Predictability and stability are of course critical rule-of-law values, but so too is democratic participation in lawmaking.¹⁹⁴ A core

191. See, e.g., MODEL RULES R. 4.1, 8.4(c); RESTATEMENT (THIRD) OF THE LAW GOVERNING LAWYERS § 98(1) cmt. b (2000).

192. See Theodore W. Ruger et al., *The Supreme Court Forecasting Project: Legal and Political Science Approaches to Predicting Supreme Court Decisionmaking*, 104 COLUM. L. REV. 1150, 1152 (2004) (reporting that a statistical model, which relied on general case characteristics predicted seventy-five percent of the Court's affirm/reverse results correctly, while legal experts collectively got 59.1 percent right); see also Elizabeth Earl, *Law Profs Develop Supreme Court Predictor to Better Understand Court Decisions*, A.B.A. J. (Dec. 1, 2014), http://www.abajournal.com/magazine/article/law_profs_develop_supreme_court_predictor/?utm_source=maestro&utm_medium=email&utm_campaign=tech_monthly [<https://perma.cc/U3U9-RN8J>]. Evidence suggests more broadly that statistical prediction is more accurate than clinical prediction in most contexts. See William M. Grove, *Clinical Versus Statistical Prediction: The Contribution of Paul E. Meehl*, 61 J. CLINICAL PSYCH. 1233 (2005).

193. See Oliver Wendell Holmes, *The Path of the Law*, 10 HARV. L. REV. 457 (1897); Robert Cooter, *Do Good Laws Make Good Citizens? An Economic Analysis of Internalized Norms*, 86 VA. L. REV. 1577, 1591 (2000) (“[T]he ‘bad man’ treats the law as ‘external’ to himself, in the sense that he considers it to lie outside of his own values. Economic models of law typically accept the ‘bad man’ approach and add a rationality element to it: a rational ‘bad man’ decides whether or not to obey the law by calculating his own benefits and costs, including the risk of punishment.”).

194. Jeremy Waldron, *The Concept and the Rule of Law*, 43 GA. L. REV. 1, 5 (2008) (“[O]ur understanding of the Rule of Law should emphasize not only the value of settled, determinate rules and the predictability that

way in which citizens participate is through their lawyers, who translate their interests into persuasive and sometimes novel arguments as to how the law should apply to their clients' circumstances. Lawyers can do so because our legal system is about reasons as well as outcomes—reasons, asserted by lawyers and memorialized in judicial opinions, which provide a continual opportunity through which to debate and potentially change the law.¹⁹⁵ If lawyering is replaced by computer prediction, we will shift to a system that is more about outcomes than reasons—and outcomes that are inescapably “informed by the world as it was in the past, or, at best, as it currently is.”¹⁹⁶

Of course, this may change over time. As natural language processing capabilities advance and computers become more capable of processing concepts and analogies, combinatorial processing may join computer prediction with computer creativity. As lawyers recognize, creativity and novelty in legal arguments generally come from importing legal concepts from one area of law into another, and by combining existing arguments in new and persuasive ways. Indeed, knowledge production in many fields proceeds in this way—by recombining existing ideas in new and innovative ways.¹⁹⁷ Computers cannot currently do this, but their ability to do so will likely increase over time. Much as medical diagnostic programs currently suggest disease hypotheses to physicians based on patient symptoms,¹⁹⁸ legal argument programs may soon be able to suggest new and promising combinations of existing arguments tailored to a client's factual circumstances. For now, computer programs are highly effective in making predictions given the legal system as it currently exists, but far less so in making suggestions for how the legal system could or should evolve.

such rules make possible, but also the importance of the procedural and argumentative aspects of legal practice.”); see also Benjamin Ewing & Douglas A. Kysar, *Prods and Pleas: Limited Government in an Era of Unlimited Harm*, 121 YALE L.J. 350 (2011) (“[T]here is another side to the value of the rule of law that is especially significant in the adversarial American system: law as a structured discourse in which individuals are entitled to articulate their grievances or face their accusers, to stake their claims, and to advance reasons in support of them.”).

195. Frederick Schauer, *Giving Reasons*, 47 STAN. L. REV. 633, 658 (1995) (“That giving reasons is a way of opening a conversation may in fact be an independent basis for a reason-giving requirement.”); Ruger et al., *supra* note 192, at 1193 (noting that the Supreme Court's “role in American society is not merely to process important disputes expeditiously. Rather, the ways in which it addresses those disputes—not merely through outcomes, but through its rationales, its analytical framework, and its language—both gives voice to certain values and influences public understanding of these issues”).

196. Chris Anderson, *The End of Theory: The Data Deluge Makes the Scientific Method Obsolete*, WIRED MAG. (June 23, 2008), <https://www.wired.com/2008/06/pb-theory/> [<https://perma.cc/5K49-B2CD>]; Martin Hilbert, *Big Data for Development: From Information to Knowledge Societies* (SSRN Scholarly Paper No. 2205145, 2013), <http://papers.ssrn.com/abstract=2205145> [<https://perma.cc/ZX2H-284X>].

197. See Martin L. Weitzman, *Recombinant Growth*, 113 Q.J. ECON. 331, 331 (1998) (“Production of new ideas is made a function of newly reconfigured old ideas in the spirit of the way an agricultural research station develops improved plant varieties by cross-pollinating existing plant varieties.”).

198. Mike Orcutt, *Why IBM Just Bought Billions of Medical Images for Watson to Look at*, MIT TECH. REV. (Aug. 11, 2015), <https://www.technologyreview.com/s/540141/why-ibm-just-bought-billions-of-medical-images-for-watson-to-look-at/> [<https://perma.cc/A34D-8SZN>].

Another set of problems created by the nature and process of automated prediction entails a lack of transparency. Like Big Data applications generally, most legal prediction programs give a user results without showing the precise combination of factors that produced those results.¹⁹⁹ Certainly, an application's programmers can view the code of the relevant inductive rule model, but the code is not always interpretable by the programmer, much less a lay person, and will frequently be proprietary, protected as a trade secret. Interpretability could be prioritized such that no outcome would be accepted—whether by a client, a lawyer, or a court—without a full explanation of inputs, but there is reason to doubt that this will happen. Requiring every outcome to be accompanied by a complete explanation of inputs (features that gave rise to the computer's model) would be exceedingly expensive and time consuming. Most users would not be willing to bear that expense. Moreover, interpretability might be a reasonable goal for applications that consider a modest amount of data, but as the universe of data expands to tens of thousands or millions of variables (words, linguistic features, data points), the goal of interpretability becomes more and more difficult, if not impossible.

This lack of transparency threatens a number of consequences over time. If clients increasingly rely on software predictions in determining a course of action—in deciding, for example, whether to file a complaint, to defend a case, or to pursue a particular corporate transaction—the software's predictions, by virtue of their influence over conduct, will influence the law in action. Without anyone realizing it, factors encoded into those predictions—including discriminatory or otherwise problematic factors—could then become encoded into broader swaths of law. For example, a computer might discover a weak correlation between a particular court's decisions and the gender, race, or ethnicity of the litigants. The estimated statistical model would then account for the correlation in predicting success or failure. Because the correlation is weak, the model's results might not immediately alert us to its influence in a way that would allow for accountability. Nevertheless, the discriminatory pattern would inform predictions of the court's decisions, and litigant behavior in the shadow of those decisions. There is also the possibility of the opposite—of technology being used to counter, rather than to entrench, human biases. One could imagine a sophisticated prediction model that produced race-neutral sentencing suggestions based only on the facts of the case. But any such use of prediction software would require coordinated attention and

199. David Martens & Foster Provost, *Explaining Documents' Classification*, at 2 (N.Y.U. Stern Sch. of Bus., Working Paper No. CeDER-11-01, 2011), <http://archive.nyu.edu/handle/2451/29918> [<https://perma.cc/2YZJ-LNQG>] (“Unfortunately, due to the high dimensionality, understanding the decisions made by the document classifiers is very difficult. Previous approaches to gain insight into black-box models do not deal well with high-dimensional data.”); see also Tal Z. Zarsky, *Transparent Predictions*, 2013 U. ILL. L. REV. 1503, 1520 (“Yet interpretability has a flip side as well. Mandating interpretability might render the process less complex and therefore less accurate.”).

action by a broad swath of implicated stakeholders, which would slow the current progress, pushed largely by one particular set of stakeholders—insurance companies and litigation financing firms, eager to gather better information about what cases to back and bring to trial.

3. ACCESS TO JUSTICE

Finally, amidst countless claims that technology alone can solve the access to justice gap,²⁰⁰ we should remain cognizant that without regulation, the development and adoption of legal technologies will be driven by the market—a decidedly ineffective means of ensuring access.

Technology proponents contend that the emergence of services like LegalZoom demonstrates that the market is working better than the profession at providing legal services at the low end of the market. Those who cannot afford a lawyer, they contend, can now access computerized legal services for low or no cost, and surely some form of legal services is better than none.

This is undoubtedly true in some contexts, but it is not in others. For one thing, the computer may encounter an unanticipated contingency but fail to alert the user, creating an error with no notice. For another, the computer cannot exhibit creativity such that, at least for now, it cannot create novel legal arguments that may initiate change in the law. The result could be “a digital divide that institutionalizes a two-tiered system incapable of delivering appropriate justice to low-income persons.”²⁰¹

Technology designed for the top of the market can also pose access challenges. One can easily imagine a dispute or transaction in which one party has access to a particular legal technology while the other party does not. For example, the significant up-front costs of predictive coding, including the licensing fees for patented programs, may be prohibitive for one party but affordable for the other.²⁰² A resource imbalance between parties is nothing new, but the unequal access to technology may introduce new types of unfairness or even abuse. The party who cannot afford predictive coding will likely lack understanding of the technology and therefore be unable to challenge the proposed discovery approaches and predictive coding protocols of the opponent. Meanwhile, the party with predictive coding, aware that the other party lacks access and has limited resources to fund manual review, will have opportunities to hide relevant and damning documents amidst massive document productions.

This does not mean that technology should not or will not play an important role in addressing the access to justice gap but rather, it is to say that the

200. See *supra* note 174.

201. Cabral et al., *supra* note 11, at 257; see Julia R. Gordon, *Legal Services and the Digital Divide*, 12 ALB. L.J. SCI. & TECH. 809 (2002).

202. See Remus, *Predictive Coding*, *supra* note 47.

profession has an important part to play in ensuring that legal technologies are made accessible and used in ways that contribute to, rather than undermine, universal access to the legal system. Segments of the profession are doing just this. For example, the California Administrative Office of the Courts commissioned a study of California legal services providers and self-help center staff to identify potential benefits and barriers that increased use of technology posed for low-income persons.²⁰³ Among other things, the report recommended hybrid legal services systems, which integrate human and automated legal assistance.²⁰⁴ A number of law schools across the country are offering courses in which students design web-based applications that make legal information accessible while explicitly informing the user that she is not receiving legal advice and should contact an attorney with questions (thus, taking the safe approach to unanticipated contingencies).²⁰⁵ Northeastern law school has launched NULaw-Lab, which involves students in a range of projects that use technology to make law more accessible to everyone.²⁰⁶

4. A THOUGHT EXPERIMENT

To further highlight why we believe that some form of professional oversight and regulation of new legal technologies is essential, we offer a simple thought experiment. Suppose that new software can accurately predict the likelihood that an individual will be audited by the Internal Revenue Service (IRS) and, if audited, that the proposed tax treatment of an asset-sheltering trust will be upheld. Suppose further that the software offers each prediction as a numerical probability, and there are no error costs. It is marketed to, and widely adopted by, financial planners who serve wealthy clients interested in minimizing gift and estate taxes.

What will be lost if this software eclipses the advice of tax and estate planning lawyers, such that the values, norms, and structures of the legal profession are cut out of the equation? Answering this question highlights value that lawyers provide and that, at least for the time being, computers cannot:

203. JUD. COUNCIL OF CAL. CT. TECH. ADVISORY COMM., *ADVANCING ACCESS TO JUSTICE THROUGH TECHNOLOGY: GUIDING PRINCIPLES FOR CALIFORNIA JUDICIAL BRANCH INITIATIVES* (2012). The report recognized that “[b]ecause so many cases now involve self-represented parties, technology must be implemented in ways that benefit those with or without legal representation so that all parties have equal access to the courts.” *Id.* at 5.

204. *Id.* at 7–8.

205. Rostain et al., *supra* note 174, at 744.

206. For example, an online game prepares individuals to represent themselves in court, *see, e.g., RePresent: Online Game for Self-Represented Litigants*, NULAWLAB, <http://www.nulawlab.org/view/online-simulation-for-self-represented-parties> [<https://perma.cc/872G-A26D>] (last visited Mar. 10, 2017); a mobile phone app provides underserved women veterans with information about their legal rights and available benefits, *see, e.g., Women Veterans Outreach Tool*, NULAWLAB, <http://www.nulawlab.org/view/women-veterans-outreach-tool> [<https://perma.cc/7DCR-9P6T>] (last visited Mar. 10, 2017); and an automated hotline informs domestic workers in the Boston area of their legal rights, *see, e.g., The Worker App*, NULAWLAB, <http://nulawlab.org/view/the-domestic-worker-app> [<https://perma.cc/U22A-AJBV>] (last visited Mar. 10, 2010).

- Counseling. The tax software can predict how the IRS will act, but it cannot and will not counsel the taxpayer on how to proceed, including on the value of compliance and the possibility of an alternative course of action.²⁰⁷ Nor can it push back against a taxpayer who insists on proceeding with an illegal scheme, notwithstanding the fact that, as Elihu Root famously asserted, sometimes the proper role of the lawyer is to tell clients “that they are damned fools and should stop.”²⁰⁸
- A robust understanding of law. Individuals planning their affairs pursuant to the software’s prediction will come to experience the law purely in terms of what conduct will and will not be sanctioned—“what the courts will do in fact, and nothing more.”²⁰⁹ This impoverished view of the law will have detrimental consequences not only for compliance,²¹⁰ but also for perceptions of the legal system’s legitimacy and democratic participation in lawmaking.²¹¹
- Respect for clients’ interests. The software objectifies a user²¹² by assuming that the objective of all users is to use any asset-sheltering trust arrangement for which the projected savings outweigh the risk of detection. Some individuals engaged in estate planning seek excessively aggressive strategies, but others simply want to ensure that they are not needlessly sacrificing assets that could be shielded under well-settled law. The software simply ignores this, projecting one set of interests onto all clients.²¹³

207. Deborah L. Rhode, *The Profession and the Public Interest*, 54 STAN. L. REV. 1501 (2002) (“One of the most crucial functions of legal counsel is to help individuals evaluate short-term economic objectives in light of long-term reputational concerns and to live up to their best, not worst instincts.”); Harold Williams, *Professionalism and the Corporate Bar*, 36 BUS. LAW. 165–66 (1980).

208. Deborah L. Rhode, *Moral Counseling*, 75 FORDHAM L. REV. 1317, 1321 (2006) (quoting PHILIP C. JESSUP, 1 ELIHU ROOT 133 (1938)); see also RHODE, INTERESTS OF JUSTICE, *supra* note 100; Wendel, *Professionalism as Interpretation*, *supra* note 100; Gordon, *New Role for Lawyers?*, *supra* note 100; Rostain, *Ethics Lost*, *supra* note 100, at 1274–75.

209. Holmes, *supra* note 193, at 460–61.

210. Robert W. Gordon, *A Collective Failure of Nerve: The Bar’s Response to Kaye Scholer*, 23 LAW & SOC. INQUIRY 315 (1998) (“[T]he order of rules and norms, policies and procedures, and institutional actors and roles that make up the legal system . . . is only as effective as voluntary compliance can make it; for if people routinely start running red lights when they think no cop is watching . . . the regime will collapse.”); Stephen Pepper, *Counseling at the Limits of the Law: An Exercise in the Jurisprudence and Ethics of Lawyering*, 104 YALE L.J. 1545, 1547–48 (1995) (“In a complex legal environment much law cannot be known and acted upon, cannot function as law, without lawyers to make it accessible to those for whom it is relevant.”); W. Bradley Wendel, *Legal Ethics As “Political Moralism” or the Morality of Politics*, 93 CORNELL L. REV. 1413, 1417–18 (2008).

211. Wendel, *Professionalism as Interpretation*, *supra* note 100, at 1173 (“[T]he law cannot operate as a device to settle normative conflict and coordinate activity without a commitment on the part of law-interpreters to respect the substantive meaning standing behind the formal expression of legal norms.”).

212. Simon, *supra* note 99, at 53–54 (warning that lawyers who adhere to the dominant ideology of professionalism “impute certain basic aims to the client,” which tend to be legalistic and to “emphasize extreme selfishness”); Kruse, *Client-Centered Norms*, *supra* note 99, at 346.

213. And yet, as Kate Kruse and others have persuasively argued, lawyers can and should work to advance and represent their clients’ interests, understood holistically; not the interests that they or the legal system

- Access to reasons. The tax software, like most Big Data applications, offers no reasons for its predictions. And yet, reasons, from lawyers as much as from judges,²¹⁴ are a critical source of both stability and change in the law²¹⁵ and a critical expression of respect for participants in the legal system.²¹⁶ Without reasons, neither the taxpayer nor the financial planner could understand the law so as to follow it or extrapolate the result to similar cases.²¹⁷ Nor could they critique the result, or argue for change.²¹⁸
- Interaction with the legal system. Finally, widespread displacement of estate and tax lawyers by prediction software would eliminate a critical mechanism through which the state and society interact.²¹⁹ Lawyers translate their clients' interests into terms the legal system can understand and act upon, and the law into terms that their clients can understand and act upon.²²⁰ Here, a lawyer could educate a taxpayer regarding the IRS's regulatory goals, and suggest an arrangement that would still minimize taxes without thwarting those goals. Or the lawyer could represent the taxpayer's interests in challenging the IRS's treatment of a particular arrangement or interpretation of a particular Code provision.

In some contexts, and with regards to some technologies, the benefits of decreased expense and increased certainty and determinacy in the law may outweigh, or may be achievable without, the costs. Moreover, the costs are those of eliminating lawyers entirely, and not a necessary consequence of the technologies themselves. The import of our thought experiment is not, therefore, to condemn legal technologies. Rather, it is to show the importance of ensuring that their development, adoption, and use are governed by norms and regulations that align with the underlying values of our legal system.

project onto clients. Kruse, *Beyond Cardboard Clients*, *supra* note 99, at 127–28 (describing client-centered lawyering, which entails “hearing clients’ stories and understanding their values, cares, and commitments,” as an answer to the problem of legal objectification); DAVID A. BINDER ET AL., *LAWYERS AS COUNSELORS: A CLIENT-CENTERED APPROACH* 2–15 (1991).

214. Wendel, *Professionalism as Interpretation*, *supra* note 100, at 1169–70 (discussing professionalism as “demand[ing] that lawyers provide a public, reasoned justification for an interpretation of legal texts—one which is plausible in light of the interpretive understandings of a professional community.”).

215. David Luban, *Natural Law as Professional Ethics: A Reading of Fuller*, in *NATURAL LAW AND MODERN MORAL PHILOSOPHY* 176, 204 (Ellen Frankel Paul et al. eds., 2001).

216. See Schauer, *supra* note 195, at 656; Luban, *supra* note 215, at 110–11 (discussing Fuller’s distinction between law and managerial direction, and view that the former implies “a certain built-in respect for [the] human dignity” of those subject to the law).

217. Schauer, *supra* note 195, at 641 (“When we provide a reason for a particular decision, we typically provide a rule, principle, standard, norm, or maxim broader than the decision itself, and this is so even if the form of articulation is not exactly what we normally think of as a principle.”).

218. *Id.* at 658.

219. See DANIEL MARKOVITS, *A MODERN LEGAL ETHICS: ADVERSARY ADVOCACY IN A DEMOCRATIC AGE* 171–200 (2008).

220. Remus, *Reconstructing Professionalism*, *supra* note 10, at 37.

5. A MORE MEANINGFUL APPROACH TO REGULATION

A roadmap for regulatory reform is beyond the scope of this Article and is a task for future work. For now, we propose proportionality as the guiding principle. Not all existing and emerging software will perform a task as well as—much less better than—a human lawyer. That is not a necessity for adoption, however, given that most software is likely to perform tasks more cheaply than a human lawyer. The issue is one of proportionality: is the less-than-human performance adequate for the task at hand, particularly given the lower cost?

The principle of proportionality recognizes that in certain contexts, lowered quality may be an acceptable and desirable tradeoff in service of increased access; in other contexts, it will not be. Many potential clients may feel that an expert system to address routine compliance or an online service provider to draft a basic will provide the level of service they want and need even if the system is not able to analyze problems as completely as a skilled human lawyer. These clients may feel that the more detailed analysis by a human lawyer does not justify its cost. Many potential clients may receive services through an expert system, such as the DoNotPay program that contests parking tickets for free, that they never could or would have received from a lawyer with little or no offsetting risk.²²¹ But clients may feel differently in other contexts, such as in the courtroom or in a child custody battle. More generally, task automation may impose a degree of standardization that loses a degree of human nuance, but the nuance may not be important for many potential users.

This, in turn, raises the key questions of client identity and autonomy. Who should make the decision of where and when these tradeoffs are acceptable? As a predictive matter, the bar has been much more willing to acquiesce to acceptance of risk by sophisticated and corporate users of legal services than by first-time individuals, but it is generally the latter who need more affordable legal services and may be the most eager for the tradeoff of proportionality. It is also the latter who are the targeted clients of new technologies that are offered without lawyer supervision, such as LegalZoom, and for whom the consumer protection rationale of regulation may be essential. Accordingly, we think these decisions cannot be entirely left to clients. We think there must be a role for regulatory bodies, populated largely though not exclusively by lawyers.

To make informed regulatory decisions, lawyers generally and bar committees in particular will have to become more informed and more skilled with new legal technologies. Both groups will also need to struggle with the bounds of the “practice of law” and with the increasingly mixed nature of legal expertise and

221. “In the 21 months since the free service was launched . . . DoNotPay has taken on 250,000 cases and won 160,000, giving it a success rate of 64% appealing over \$4m of parking tickets.” Gibbs, *supra* note 95. The creator plans to expand to address flight delay compensation, rights of HIV positive individuals, and refugees navigating foreign legal systems.

other forms of expertise. Only by doing so will the bar be able to adjust to current realities and fulfill its obligation to society.

CONCLUSION

At the risk of oversimplifying, we think it is fair to characterize much of the current debate regarding legal technologies as existing at the extremes. With respect to employment effects, headlines proclaim the end of the legal profession. Traditionalists respond with unauthorized practice of law rules, arguing that new technologies threaten client interests and undermine the core values of the profession. Many scholars and commentators push back, arguing that we should automate as many legal services as possible in an attempt to reduce prices and increase access.

This Article has sought to add detail and nuance to the discussion. First, we showed that while technology is undoubtedly advancing and changing the nature of legal practice, it is displacing lawyers at a modest pace. Second, we argued that while current approaches to the professional regulation of legal technologies are ineffective and undesirable, the answer cannot be to abandon professional regulation. Instead, we must begin the difficult but important task of designing more effective regulatory structures that draw upon both legal and technical expertise, while protecting both clients and the values of our legal system.

APPENDIX

APPROXIMATE ESTIMATES OF MODERATE AND LIGHT
EMPLOYMENT EFFECTS

Moderate Employment Effects: As noted above, moderate employment effects arise when a largely unstructured legal task has a significant structured component that can be computerized. To calibrate the employment impact of this level of innovation, we refer to a case study of search-related innovation in exceptions processing at a large bank.²²² Exceptions processing requires determining the proper disposition of

checks written on accounts that have been closed, checks written for amounts greater than the balances in the accounts on which they are drawn, checks that customers request stop payments on, checks written for large amounts that require signature verification, and fraudulent checks.²²³

Each department employee reconciled a single type of exception. The work was made more complex because a single check could involve multiple exceptions. For example, individuals short of cash might buy time by writing multiple checks to creditors and by then submitting multiple stop-check orders. The result was substantial time spent both searching boxes of checks for particular items and coordinating work among employees addressing different exceptions for the same account.

When digital check images were substituted for paper checks in the workflow, employees who handled exceptions gained rapid access to a particular check, resulting in reduced search time and, therefore, increased productivity. Simultaneously, the exceptions departments reorganized their workflow such that employees no longer focused on a particular type of exception but instead handled all exceptions for a particular set of accounts.

This new technology and reorganization, which could have taken place with paper checks (though it did not), increased productivity. The combined effect was to reduce the number of employees required to handle a constant volume of exceptions from 650 to 470—a reduction of twenty-eight percent. A computer-friendly estimate attributes two-thirds of this reduction to the digitized images and one-third to the reorganization. Correspondingly, we assume that lawyering tasks in which computers have a Moderate Employment Effect reduce lawyer time devoted to those tasks by nineteen percent.

Light Employment Effects: This category includes tasks that entail largely unstructured work with limited room for automation—e.g., Fact Investigation or

222. See Autor et al., *supra* note 22, at 437.

223. *Id.*

Advising Clients.²²⁴

To calibrate Light Employment Effects, we use a case study of a limited computer innovation in healthcare: Adler-Milstein and Huckman's study of the impact of electronic medical record (EMR) use on clinician productivity.²²⁵ Productivity in the study is measured by "Relative Value Units" billed per clinician workday, which is the standard medical accounting measure of the volume and intensity of services provided. The study's sample consists of forty-two medical practices, which were observed over three years during which they implemented EMRs at various rates. Findings indicate that one standard deviation in the use of an EMR increases clinician productivity by five percent. Services per patient visit did not increase, but physicians could see more patients per workday by using the EMR to delegate some services to physicians' assistants. Relying on this example, we posit that adopting a computer innovation with Light Employment Effects would decrease required lawyer employment for a given task by five percent.

224. Light employment effects also arise in tasks relating to document management. Because these tasks are usually performed by clerical staff, automation does not affect lawyer employment per se.

225. Adler-Milstein & Huckman, *supra* note 125.