

UNACCOUNTABLE

How Qualified Immunity Shields a Wide Range of Government Abuses,
Arbitrarily Thwarts Civil Rights, and Fails to Fulfill Its Promises



 INSTITUTE
for JUSTICE

By Jason Tiezzi, Robert McNamara, and Elyse Smith Pohl
February 2024

UNACCOUNTABLE

How Qualified Immunity Shields a Wide Range of Government Abuses,
Arbitrarily Thwarts Civil Rights, and Fails to Fulfill Its Promises



By Jason Tiezzi, Robert McNamara, and Elyse Smith Pohl
February 2024

TABLE OF CONTENTS

Executive Summary	4
Introduction	6
The Controversy Over Qualified Immunity	8
Methods.....	10
Results: Describing Qualified Immunity in Federal Circuit Courts	16
Discussion: A Broad and Unnecessary Shield.....	23
Conclusion	30
Appendix A: Developing the Predictive Algorithms	32
Appendix B: Predictive Algorithms' Performance on Holdout Test Data	37
Appendix C: Key Data by Circuit Court.....	40
Appendix D: Main Codebook.....	42
Appendix E: First Amendment Codebook	47
Endnotes.....	48
About the Authors.....	55
Acknowledgments.....	56

EXECUTIVE SUMMARY

Qualified immunity is perhaps America's most controversial legal doctrine, erupting into the national consciousness during debate over police misconduct in 2020. Created by the U.S. Supreme Court four decades ago, the doctrine protects government officials from being sued for violating constitutional rights—unless victims can show those rights are “clearly established.” In practice, this often means pinpointing a published opinion from the Supreme Court or the federal appellate court in their jurisdiction finding the same conduct in the same circumstances unconstitutional.

The Supreme Court intends for qualified immunity to give government officials leeway to make reasonable mistakes—especially in tense or dangerous situations requiring quick thinking—without facing lengthy litigation, onerous discovery, or financial ruin. By insisting rights be clearly established to receive protection, the Court aims to put officials on notice of conduct to avoid before they face such consequences. Critics counter, however, that qualified immunity sets too high a bar for victims of abuse to seek justice and winds up protecting officials who intentionally, maliciously, or unreasonably violate the Constitution. They also argue the doctrine does not work as the Court intends.

This study adds new evidence to the record using the largest ever collection of federal appellate cases, covering the 11-year period from 2010 through 2020. It is the first to use cutting-edge automated techniques to parse thousands of federal circuit court opinions and answer key questions about cases where government defendants claim qualified immunity—what kinds of officials and conduct it protects, its impact on civil rights cases, and whether the doctrine is achieving its aims.

Key findings include:

In the federal circuit courts, qualified immunity appeals have become more common.

- From 2010 through 2020, at least 5,526 cases before federal circuit courts raised qualified immunity on appeal, an average of about 500 cases a year.
- And from the first half of our study period to the second, the annual average of qualified immunity appeals grew by 20%, even as civil appeals of all types fell.

Contrary to popular belief, qualified immunity is not just about police accused of excessive force. It shields a wide array of government officials and conduct.

- While police were the most common defendants, fully half of appeals featured other types of government officials, either alongside or instead of police. Prison officials made up the next largest share, but in more than one in five of all appeals, or 21%, defendants were neither police nor prison officials. These other officials included mayors and city managers, university and school officials, prosecutors and judges, and child protective services workers.
- Excessive force was alleged in just 27% of appeals, followed by false arrest at 25%; some alleged both. But the third largest category, alleged in 18% of appeals, encompassed violations of First Amendment rights, including speech, association, and religious liberty.
- Altogether, only 23% of appeals fit the popular conception of police accused of excessive force.

In most First Amendment appeals, plaintiffs alleged government officials engaged in premeditated retaliation for protected speech or activity.

- In a representative sample of 125 First Amendment appeals, 59% involved plaintiffs alleging premeditated abuse by government officials in retaliation for protected First Amendment activity.
- In nearly half of such cases, government workers alleged retaliation from their superiors, while in nearly a third, private citizens claimed they were targeted for retaliation by government officials.

Qualified immunity favors government defendants and makes it harder for plaintiffs to win—regardless of the merits of their claims.

- In all, 59% of qualified immunity appeals were resolved solely in favor of government defendants, while 24% were resolved solely in favor of plaintiffs.
- Qualified immunity disadvantages plaintiffs for arbitrary reasons, such as their circuit’s population or publication rate. These vary widely and influence the volume of clearly established law in a circuit—and therefore, the volume and variety of prior cases that plaintiffs can rely on to vindicate their rights.
- Qualified immunity rulings often lack precision and clarity, again making it hard for plaintiffs to pinpoint the clearly established law required to win. In common with other legal experts, ours often could not untangle courts’ reasons for granting qualified immunity—if reasons were even offered.
- When denied qualified immunity, government defendants have the right to file special immediate appeals—a right unavailable to plaintiffs. And they can do this multiple times in the same lawsuit. Such “interlocutory appeals” accounted for 96% of all defendant appeals.
- These special appeals risk wearing down worthy plaintiffs with extended litigation. Their prevalence likely helps explain why the median duration of a qualified immunity lawsuit was three years and two months, 23% longer than the typical federal civil suit up on appeal.

Our findings provide more evidence that qualified immunity is a poor fit for achieving its goals.

- Qualified immunity confuses instead of clarifies the rules government workers must follow to avoid burdensome litigation. If legal experts struggle to make sense of qualified immunity, the average government official—let alone one facing a life-or-death situation—cannot be expected to do so.
- Qualified immunity fails to protect officials from the burdens of litigation, most notably potentially intrusive discovery. Nearly 70% of appeals came at the summary judgment stage of litigation, when courts typically have already allowed discovery.
- Qualified immunity clogs up the courts with extra, often lengthy, appeals—some 2,000 interlocutory appeals that would not have existed without the special appeal rights given to government defendants.

These results suggest qualified immunity shields a much wider array of government officials and conduct than commonly thought. They also add to a growing body of research finding qualified immunity protects officials too much and our rights too little, all while failing to achieve its goals. This strengthens the argument for the Supreme Court or Congress to temper or—better yet—abolish the doctrine.

Whether through outright abolition or significant reform, courts and lawmakers can and should act to eliminate the unbounded impunity allowed by qualified immunity.

INTRODUCTION

Sylvia Gonzalez spent a day in jail because her political opponents wanted to teach the then-72-year-old grandmother and city councilwoman a lesson. The mayor and police chief of Castle Hills, Texas, used trumped-up charges to jail Sylvia in retaliation for speech critical of another local government official. The charges were later dropped.¹

Allan Minnerath saw his company's trucks and drivers detained for hours by an overzealous county road engineer in Mahnomon County, Minnesota. The official, who opposed a state contract awarded to Allan's firm, decided to do something about it: He changed the weight limits on roads he knew the trucks would travel, then played traffic cop by personally stopping the now-overweight vehicles.²

Shaniz West had her Caldwell, Idaho, home destroyed by police after she gave them permission to search the residence for her fugitive ex-boyfriend. Rather than use the keys she gave them, police shot tear gas grenades into the home, rendering it uninhabitable and ruining all of Shaniz's and her children's belongings.³

Sylvia, Allan, and Shaniz are Institute for Justice clients whose constitutional rights were violated by government officials, and all three sought justice by suing those government officials. But Sylvia, Allan, and Shaniz have something else in common: The government officials who violated their rights invoked a contentious court-made doctrine known as qualified immunity to try to escape the lawsuits.⁴ Bringing a lawsuit is the only way for Americans to directly hold government officials accountable and enforce our rights, but qualified immunity, along with similar doctrines, has made it harder and harder to do.⁵

Created by the U.S. Supreme Court in 1982, qualified immunity is a special protection for government officials.⁶ Under the doctrine, officials cannot be sued or held civilly liable for violating a person's constitutional rights unless the person can identify a previous decision from the Supreme Court or the federal appellate court in the same jurisdiction "clearly



Sylvia Gonzalez

establishing” that such conduct is unconstitutional. And this protection applies even if officials intentionally, maliciously, or unreasonably violated the Constitution. Indeed, courts have granted qualified immunity to officials accused of all sorts of egregious conduct. This includes Denver police officers who threatened to arrest a man for recording them, then illegally searched his device for the video.⁷ It also includes Texas Medical Board investigators who conducted a warrantless search of a doctor’s office for patients’ medical records.⁸

Qualified immunity applies even if officials intentionally, maliciously, or unreasonably violated the Constitution.

This immunity does not exist in civil cases between private parties.⁹ If you sue someone for injuring you in a car crash, they must go to court to argue that they are not responsible. With qualified immunity, government officials go to court to argue that they cannot be sued—even if they did, in fact, violate your constitutional rights.

This study describes how qualified immunity works in federal appellate courts using the largest ever collection of cases, covering the 11-year period from 2010 through 2020. It is the first to use cutting-edge automated techniques to parse thousands of federal appellate opinions and answer key questions about cases where government defendants claim qualified immunity: how common they are, who wins them, how long they take, and what kinds of government officials and alleged constitutional violations they involve. The answers offer new evidence about the kinds of officials and conduct qualified immunity protects, its impact on civil rights cases, and whether the doctrine is achieving its aims.



Allan Minnerath



Shaniz West

THE CONTROVERSY OVER QUALIFIED IMMUNITY

Qualified immunity is perhaps America's most controversial legal doctrine. Particularly in the wake of the national debate over police misconduct that erupted in 2020, commentators on all sides have focused on qualified immunity as it relates to police accountability and to the ability of law enforcement officers to navigate stressful, rapidly evolving situations.¹⁰

Despite this recent focus on law enforcement, the doctrine stems from a 1982 U.S. Supreme Court case that had nothing to do with policing. Instead, the case, *Harlow v. Fitzgerald*, was brought by a former executive branch employee who alleged that two Nixon administration aides conspired to get him fired in retaliation for his testimony to a congressional committee. The Supreme Court ruled that the aides were entitled to “qualified immunity,” which meant they could not be held liable unless they had violated a “clearly established” constitutional right.¹¹

In the four decades since, qualified immunity has barred lawsuits against federal, state, and local government officials unless plaintiffs can show that officials violated a clearly established constitutional right. And in practice the bar for being clearly established can be high: Courts often interpret clearly established strictly, meaning plaintiffs must show that either the Supreme Court or the federal appellate court in the same jurisdiction has already held that the same conduct in the same circumstances is unconstitutional.¹²

Moreover, it is not enough to convince a district court judge that the right was clearly established. Qualified immunity gives government defendants a special right to immediately appeal to higher courts. In other types of civil cases (and in criminal ones), defendants can typically appeal only a final

judgment against them. But government officials invoking qualified immunity can file what is known as an “interlocutory appeal” to immediately challenge the denial of qualified immunity at any stage of the litigation.¹³

An official can ask a court to dismiss a case on qualified immunity grounds right after it is filed. If the court declines, the official can immediately file an interlocutory appeal instead of waiting for the case to be heard and decided on its merits. Later,

the official can ask the court to throw the case out before trial because the evidence unearthed in discovery is insufficient to overcome qualified immunity. If the court declines, the official can file another of these special immediate appeals. And, of course, if the case goes to trial and the plaintiff wins, the official can file a final appeal after a judgment is entered. Qualified immunity therefore triples the appellate opportunities available to government officials accused of violating people's constitutional rights—and, when used, multiplies the accompanying costs and delays.¹⁴

In granting these substantial protections to government officials, the Supreme Court acknowledges the need to hold officials responsible for misconduct—but it also aims to shield them from “harassment, distraction, and liability when they perform their duties reasonably.”¹⁵ The idea is to let some cases through the courthouse doors—those where rights were clearly established in prior cases, thus putting government officials on notice of conduct to avoid—while blocking others. Agents of the government, the argument goes, are often called upon to make difficult decisions, so they must have leeway to make reasonable mistakes without facing time-consuming litigation or financial ruin.¹⁶ Making it too easy for such claims to proceed could chill officials from performing their duties or make it harder to recruit officials in the first place.¹⁷

Unsurprisingly, the doctrine has drawn widespread criticism, much of it focused on the clearly established test.¹⁸ Critics argue that the test unfairly disadvantages victims of alleged government misconduct both now and in the future. First, when applied to stop a lawsuit because the right at issue was not clearly established by a prior case, qualified immunity can protect even intentional, malicious, or unreasonable conduct that happens to have not come up in litigation before—or that has only arisen in a different federal circuit.¹⁹ The Supreme Court itself has occasionally noted that sufficiently egregious misbehavior might warrant liability even in the absence of a ruling clearly establishing the existence of the right that was violated. Indeed, in 2002, it created an “obviousness” exception for such cases.²⁰ Yet the Court has rarely used this exception—just twice in the 24 qualified immunity cases that came before it between 2002 and 2020.²¹ Moreover, the Court has repeatedly warned lower courts “not to define clearly established law at a high level of generality,” which would, in effect, make it easier for plaintiffs to find cases to meet the test.²²

Second, critics argue the clearly established test not only burdens current plaintiffs but also disadvantages future

In the four decades since its creation, qualified immunity has barred lawsuits against federal, state, and local government officials unless plaintiffs can show that officials violated a clearly established constitutional right. And in practice the bar for being clearly established can be high.

plaintiffs, particularly in the wake of the Supreme Court's 2009 decision in *Pearson v. Callahan*.²³ That ruling gave courts facing a claim for qualified immunity a choice: They can, as they were required to do pre-*Pearson*, decide whether the government official violated a constitutional right and, if so, determine whether that right was clearly established. Or—if they would rather—they can simply determine whether any case clearly establishes the right at issue without ever deciding whether the official violated it at all.²⁴ And much of the available evidence suggests courts are more likely to avoid addressing the underlying constitutional question in qualified immunity cases when they are not required to do so.²⁵ The likely result is fewer cases clearly establishing what the Constitution requires—and thus fewer cases future plaintiffs can rely on.²⁶

Critics also decry qualified immunity's practical justifications. For example, to the extent the doctrine is meant to protect government officials from ruinous financial liability, it does not appear necessary. Indeed, evidence suggests that officials who are ordered to pay damages—including those who are denied qualified immunity—hardly ever pay out of their own pockets. Instead, damages awards are almost uniformly paid by the government.²⁷ If qualified immunity is meant to protect individual government officials from financial ruin, it appears to be protecting them from a danger that does not exist.

So too with another common justification for qualified immunity: the need to ensure that officials know what conduct is and is not constitutionally permissible before being sued or

held liable, lest they be chilled from performing their duties.²⁸ As a rule, the only sources of clearly established constitutional rights that count for qualified immunity purposes are the published opinions of federal appellate courts, and there is little evidence that officials at any level of government assiduously follow the doings of their regional circuit court.²⁹ Quite the contrary: Research on police training suggests that officers receive little or no instruction on the clearly established law in their jurisdictions.³⁰ There is little reason to assume that other government officials like tax assessors or code enforcement officers receive any more.

Amid the ongoing debate over qualified immunity, this study adds new evidence to the record by creating the largest and most comprehensive collection of federal qualified immunity appeals yet compiled.³¹ Importantly, unlike most earlier studies, ours is not limited to claims against law enforcement defendants or those alleging excessive force.³² Such cases are arguably more likely to involve the kind of difficult, often split-second, decisions that concern the Supreme Court, yet there is nothing in the doctrine to limit it to such cases.³³ This means qualified immunity applies far more broadly—and our data help to show just how far. Our extensive and detailed dataset also provides new evidence about how qualified immunity plays out in the day-to-day work of the federal courts of appeals. Our findings suggest the doctrine arbitrarily thwarts valid civil rights claims and fails to achieve the Court's goals—all while adding years to litigation and clogging circuit court dockets.

METHODS

This study aims to describe the landscape of qualified immunity appeals in federal appellate courts. Specifically, we sought to answer these questions for the study period 2010 through 2020:

1. How many federal appeals involve qualified immunity?
2. What government officials are sued in qualified immunity appeals?
3. What rights violations are alleged in qualified immunity appeals?
4. What are the key characteristics of qualified immunity appeals?
 - a. How long does litigation involving qualified immunity appeals last?
 - b. How many qualified immunity appeals are interlocutory appeals?
 - c. At what stage of litigation do qualified immunity appeals occur?
 - d. How often are plaintiffs in qualified immunity appeals represented by counsel?
 - e. How often are opinions in qualified immunity appeals published?

5. What are the overall outcomes of qualified immunity appeals, and how frequently is qualified immunity granted or denied?

To obtain the potential universe of qualified immunity opinions, we searched Thomson Reuters' Westlaw service for any federal appellate court opinion issued between January 1, 2010, and December 31, 2020, containing the phrase "qualified immunity."³⁴ This yielded 7,173 opinions. A central part of our task was determining whether opinions were relevant—that is, whether qualified immunity was raised in the appeal as opposed to merely being mentioned in the opinion.

Given the time it would take human coders to analyze thousands of opinions, we instead used algorithms—computerized instructions, rules, and models—to identify relevant opinions and label them across 33 additional variables. We collected two more variables through a separate process.³⁵ To develop the algorithms and test their reliability, we first coded a random sample of opinions by hand. This section describes our variables, how we developed and tested the algorithms, and the final dataset.



Study Variables

For each opinion, the most important information we recorded was relevance: Was qualified immunity raised on appeal, making the opinion relevant to our study?

Then, for all relevant opinions, we recorded 35 additional fields corresponding to our research questions. These fields, with coding options, are summarized in Table 1 and further defined in our main codebook, available in Appendix D. The codebook also covers exceptions and special cases.

We coded all fields at the level of the opinion rather than claim or alleged violation found within an opinion. We used this approach because our initial research questions focused on the landscape of qualified immunity appeals, not individual claims.³⁶ This approach is consistent with many other studies on qualified immunity.³⁷ However, it means that if an opinion involved multiple claims, we cannot directly link factors like the defendants, violations, or outcomes to a particular claim.

Moreover, we coded only what was before the court on appeal. For example, if a lawsuit originally involved both law enforcement and prison defendants, but only the law enforcement defendants were involved in the appeal, we coded only the law enforcement defendants.

Developing the Algorithms

To develop the algorithms, we needed a sizable sample of reliably hand-coded opinions. This sample would allow the algorithms to find patterns in the text of opinions, resulting in reliable prediction models. To test the completed algorithms, we needed a similarly reliable, but smaller, sample of hand-coded opinions.

To create these samples, we randomly selected 791 (roughly 11%) of the 7,173 opinions for hand coding. We randomly assigned 604 opinions to the training sample and 187 to the testing sample.³⁸

To ensure accuracy, our human coders were either attorneys or others with substantial knowledge of legal matters generally, if not qualified immunity specifically.³⁹ We also conducted trainings on our codebook and tested coders' accuracy by requiring them to complete a sample of practice opinions before starting the project. Finally, we employed a multistep quality-control process involving a panel of attorneys with experience in qualified immunity to resolve the thorniest coding decisions.⁴⁰

Once our human coders completed their work, we used the training sample to build our algorithms. (Appendix A details our process for developing and implementing the algorithms.)



Table 1: Fields and Variables Included in Study

	Field	Description	Type	Response Options
Basic Information	Relevance	Was qualified immunity raised on appeal in the opinion?	Binary (Y/N)	--
	Circuit Court	The circuit court for the appeal	Text	--
	Circuit Court Case Number	The circuit court case number for the appeal	Text	--
	Opinion Date	The date the opinion was filed/decided	Text (Date)	--
	Plaintiffs	The plaintiffs in the opinion	Text	--
	Defendants	The defendants in the opinion	Text	--
	Judges	The judges who heard the appeal	Text	--
	District Court of Origin	The district court where the appeal originated	Text	--
	District Court Case Number	The case number of the lawsuit in district court	Text	--
	Case Origination Date	The date the lawsuit was initiated in district court	Text (Date)	--
Procedural Details	Appellants	Which party was appealing the district court's decision?	Categorical	P – Plaintiffs D – Defendants B – Both parties (cross-appellants)
	Published	Was the opinion published?	Binary (Y/N)	--
	En Banc	Did the opinion involve an en banc hearing?	Binary (Y/N)	--
	Interlocutory Appeal	Was the appeal an interlocutory appeal?	Binary (Y/N)	--
	Pro Se Plaintiffs	Did the lawsuit include self-represented plaintiffs?	Categorical	1 – All plaintiffs were pro se for the appeal 0 – No plaintiffs were pro se at any point in the lawsuit ES – Plaintiffs were pro se at an earlier stage in the lawsuit
Case Stage	What was the procedural stage of the lawsuit at the time of the appeal?	Categorical	D – Dismissal SJ – Summary Judgment B – Both Dismissal and Summary Judgment PT – Post-trial Other – Anything else	
Government Defendant Type	Government Level of Defendants	Were the government officials being sued federal or state/local officials?	Categorical	Federal – Only federal State – Only state/local Both – Both federal and state/local
	State Law Enforcement Defendants	Was a state/local law enforcement officer listed as a defendant?	Binary (Y/N)	--
	Federal Law Enforcement Defendants	Was a federal law enforcement officer listed as a defendant?	Binary (Y/N)	--
	State Prison Defendants	Was a state/local prison official listed as a defendant?	Binary (Y/N)	--
	Federal Prison Defendants	Was a federal prison official listed as a defendant?	Binary (Y/N)	--
	Other Defendants	Was a non-law enforcement, non-prison official listed as a defendant in the appeal?	Binary (Y/N)	--
	Task Force Defendants	Were the defendants part of a state/federal law enforcement task force?	Binary (Y/N)	--

	Field	Description	Type	Response Options
Constitutional Violation Type	First Amendment	Did the plaintiffs allege violations related to their First Amendment rights?	Binary (Y/N)	--
	Religious Liberty	Did the plaintiffs allege violations of their right to freely practice their religion? (Note: This field is a sub-field of the "First Amendment" field.)	Binary (Y/N)	--
	Excessive force	Did the plaintiffs allege that the defendants committed a violation related to excessive force?	Binary (Y/N)	--
	False Arrest	Did the plaintiffs allege that the defendants committed violations related to a false arrest, malicious prosecution, or illegal seizure of a person?	Binary (Y/N)	--
	Illegal Search	Did the plaintiffs allege that the defendants committed violations related to an illegal search?	Binary (Y/N)	--
	Procedural Due Process	Did the plaintiffs allege they were deprived of fair process under the due process requirements of the Constitution?	Binary (Y/N)	--
	Care in Custody	Did the alleged violations relate to the (lack of) care provided for the plaintiffs when they were in some form of custody?	Binary (Y/N)	--
	Parental Rights	Did the plaintiffs allege that the defendants interfered with their rights as parents?	Binary (Y/N)	--
	Employment	Were at least some of the alleged violations of constitutional rights in this opinion related to an adverse employment action, a hostile work environment, or unsafe workplace conditions?	Binary (Y/N)	--
Outcomes	Overall Prevailing Party	Who was the prevailing party in the opinion?	Categorical	P – Plaintiffs D – Defendants M – Both the defendants and plaintiffs prevailed in part (mixed)
	Qualified Immunity Granted	Was qualified immunity granted to one or more defendants in this opinion?	Binary (Y/N)	--
	Qualified Immunity Denied	Was qualified immunity denied to one or more defendants in this opinion?	Binary (Y/N)	--
	Lack of Jurisdiction – Factual Dispute	Did the court decline to rule on qualified immunity as it determined it lacked jurisdiction due to a factual dispute?	Binary (Y/N)	--

Evaluating the Algorithms' Reliability

After our algorithms were finished, we needed to evaluate their reliability. To do this, we used the testing sample to compare the datapoints generated by the algorithms to those recorded by our human coders.⁴¹

Overall, our algorithms performed very well. Nearly all fields achieved performance statistics above—often well above—those in comparable legal studies.⁴² Table 2 shows the most relevant performance statistic for each field, providing a general impression of the algorithms' performance. Appendix B presents all the statistics necessary to gauge the algorithms' performance in greater detail.⁴³

For different types of fields, we used different statistics as our primary performance metric:

- For **text fields**, “accuracy” was our primary metric. In this context, accuracy simply means how often the algorithm recorded the right text (excluding minor typos and other trivial differences, such as punctuation or articles). As Table 2 shows, these fields performed extremely well, with accuracy rates of 99% to 100%.
- For **categorical fields** (i.e., fields with multiple response options), accuracy was again our primary metric. For these fields, accuracy represents the percentage of opinions labeled with the correct option. For example, the prevailing party can be either the plaintiff, the defendant, or, in the case of a mixed decision, both. The accuracy for the prevailing party field was 96.3%, meaning the algorithm applied the right label to 96.3% of opinions. Comparable legal studies report accuracies between 73% and 93%.⁴⁴ However, we generally aimed for performance accuracies at or above 95%.⁴⁵ As detailed in Table 2, four out of five categorical fields

exceeded this threshold, with the fifth just missing it at 94.4%.

- For **binary fields** (i.e., fields with only yes/no response options), a measure called the “F1 score” was our primary performance metric. F1 scores range from 0 to 1, with 1 being perfect. While comparable legal studies report F1 scores ranging from 0.57 to 0.91, we aimed for F1 scores above 0.9, although we were generally willing to accept scores above 0.8.⁴⁶ As detailed in Table 2, we mostly succeeded, including achieving near-perfect F1 scores for several critical fields.

We focused most of our analyses on fields with strong performance—i.e., accuracy above 95% or F1 scores above 0.9. We generally avoided making detailed analyses of low-performing fields and fields with minimal data.⁴⁷

Finalizing the Dataset

After evaluating their performance, we ran the algorithms on all 7,173 opinions. In all, the algorithms generated roughly 190,000 datapoints.⁴⁸

Our full final dataset can be found on our website at <https://ij.org/report/unaccountable/data-downloads/>.

This final dataset is both comprehensive and broad: It encompasses 11 years of qualified immunity appeals and covers a range of seldom-studied attributes, including the types of government officials who were sued and the alleged rights violations at issue. Because of the algorithms' strong performance, the dataset's reliability is, for numerous critical fields, comparable to what hand coding could achieve. The scope, breadth, and reliability of this dataset allowed us to explore the landscape of qualified immunity appeals in the circuit courts.

Table 2: Summary of Algorithm Performance

Type	Field	Primary Performance Statistic	Performance
Text Fields	Circuit Court	Accuracy	100%
	Circuit Court Case Number	Accuracy	99.4%
	Opinion Date	Accuracy	100%
	Plaintiffs	Accuracy	99.4%
	Defendants	Accuracy	100%
	Judges	Accuracy	100%
	District Court of Origin	Accuracy	99.4%
Categorical Fields	Appellants	Accuracy	99.4%
	Pro Se Plaintiffs (self-represented plaintiffs)	Accuracy	99.4%
	Case Stage (at time of appeal)	Accuracy	94.4%
	Government Level of Defendants	Accuracy	96.9%
	Overall Prevailing Party	Accuracy	96.3%
Binary (Y/N) Fields	Relevance (qualified immunity raised on appeal)	F1 Score	0.95
	Published	F1 Score	1.00
	En Banc	F1 Score	1.00
	Interlocutory Appeal	F1 Score	0.99
	State Law Enforcement Defendants	F1 Score	0.96
	Federal Law Enforcement Defendants	F1 Score	--*
	State Prison Defendants	F1 Score	0.93
	Federal Prison Defendants	F1 Score	0.86
	Other Defendants	F1 Score	0.84
	Task Force Defendants	F1 Score	--*
	First Amendment Violations	F1 Score	0.98
	Religious Liberty Violations	F1 Score	1.00
	Excessive Force Violations	F1 Score	0.97
	False Arrest Violations	F1 Score	0.85
	Illegal Search Violations	F1 Score	0.71
	Procedural Due Process Violations	F1 Score	0.85
	Care in Custody Violations	F1 Score	0.82
	Parental Rights Violations	F1 Score	0.89
	Employment Violations	F1 Score	0.81
	Qualified Immunity Granted	F1 Score	0.91
Qualified Immunity Denied	F1 Score	0.86	
Lack of Jurisdiction – Factual Dispute	F1 Score	0.80	

*This field did not appear in our testing sample, meaning an F1 score could not be calculated.

Note: High-performing fields (95%+ accuracy, 0.9+ F1 score) are shaded dark green. Fields with satisfactory performance (90%+ accuracy, 0.8+ F1 score) are shaded light green. Fields with unsatisfactory performance (<90% accuracy, <0.8 F1 score) are shaded yellow. These ranges are based on the goals of our study and comparable legal studies; nevertheless, they are inherently subjective. For a full range of performance statistics and detailed data distributions for each field, see Appendix B.

RESULTS: DESCRIBING QUALIFIED IMMUNITY IN FEDERAL CIRCUIT COURTS

Volume and Growth

From 2010 through 2020, at least 5,526 cases before federal appellate courts raised qualified immunity on appeal, an average of about 500 cases a year. This does not include cases decided by state appellate courts, nor does it include non-appellate cases involving qualified immunity decided by federal district courts or state trial courts. In addition, 5,526 is likely an undercount because our algorithms for predicting whether an appeal involved qualified immunity, while highly accurate, are slightly biased in favor of underreporting.⁴⁹

The federal appellate courts varied in the number of qualified immunity appeals they decided over the study period. The 9th Circuit decided the most—and nearly nine times as many as the 1st Circuit, which decided the fewest (see Figure 1). We also saw variation in the district courts where appeals originated. Most notably, the Eastern District of Michigan accounted for more than 300 of the 6th Circuit appeals we studied, or about 1 in 20 of all appeals. Appendix C breaks out data about qualified immunity appeals for the circuit courts.⁵⁰

Qualified immunity appeals in the circuit courts have increased over time, as shown in Figure 2. In the first half of our study period (2010 through 2015), the courts decided an average of 458 qualified immunity appeals a year. But in the second half (2016 through 2020), that figure jumped to 555,

an increase of roughly 20%. Over that same timeframe, the total number of civil appeals in the circuit courts dropped by 10%, indicating that qualified immunity appeals are growing in both absolute and comparative terms.⁵¹

Driving this growth was an increase in appeals with allegations of excessive force against government defendants, typically police. From the first half of our study period to the second, excessive force appeals increased by over 50%.⁵² Meanwhile, appeals with other alleged constitutional violations also increased, although by a more modest 12%.⁵³

Government Defendants

In the appeals we analyzed, law enforcement officials were the most common defendants—but by no means the only ones (see Figure 3).⁵⁴ In fact, law enforcement officials were the sole defendants in just 50% of appeals.⁵⁵ In another 20%, only prison officials, such as corrections officers, were defendants.⁵⁶

Notably, in more than a fifth of appeals (21%), defendants were neither law enforcement nor prison officials. Our algorithms simply identified them as “other,” but our hand-coded sample suggests they tended to be mayors or city managers; university or school officials; prosecutors or judges; and child protective services workers.

Figure 1: Qualified Immunity Appeals by Circuit, 2010–2020




Figure 2: Qualified Immunity Appeals Are Growing More Common

Circuit courts decided 20% more appeals after 2015


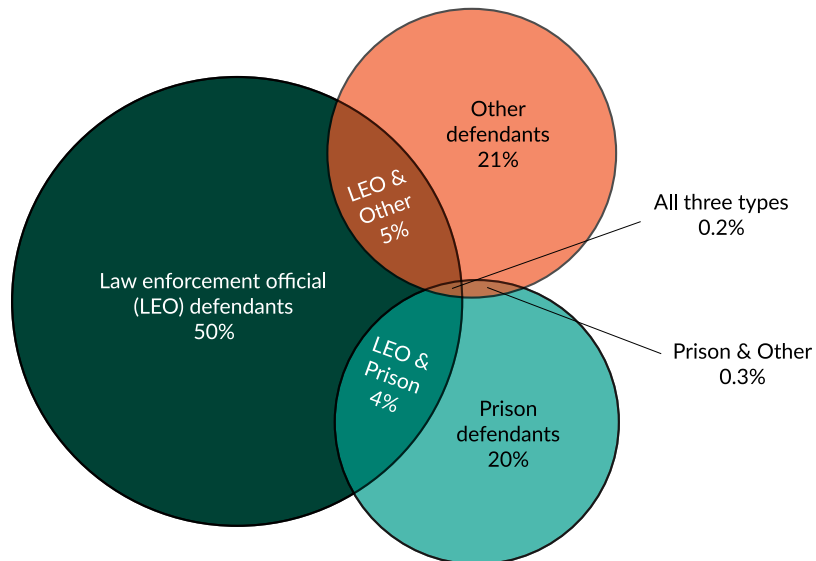


Figure 3: Government Defendants Are Diverse

While law enforcement and prison officials are most common, defendants were neither in 21% of appeals



Alleged Constitutional Violations

Excessive force and false arrest were the most common constitutional violations alleged against government defendants, raised in 27% and 25% of appeals.⁵⁷ Perhaps surprisingly, alleged First Amendment violations made up the next largest category at 18%. This category includes alleged violations of freedom of speech, assembly, association, and religious exercise.

Also common were alleged violations relating to the care of prisoners and others in custody, illegal searches, and procedural due process.⁵⁸ Each of these constitutional claims was raised in at least 10% of appeals. Figure 4 shows the prevalence of all the violation types we coded.⁵⁹

Procedural Details

We measured several procedural details about qualified immunity appeals:

- Duration of lawsuit at time of appeal decision:** The median duration of a qualified immunity lawsuit—from the initial complaint filing to the appeal decision—was three years and two months.⁶⁰ This is 23% longer than the typical civil suit up on federal appeal.⁶¹ And as shown in Figure 5, many lawsuits had been open even longer when the appeal was decided: Nearly 29% had been open for more than four years, while 8% had been open for more than six.
- Interlocutory appeals:** As detailed above, government

defendants invoking qualified immunity have a special right to immediately appeal the denial of qualified immunity. These special

Interlocutory appeals represented roughly a third of qualified immunity appeals overall—nearly 2,000 total.

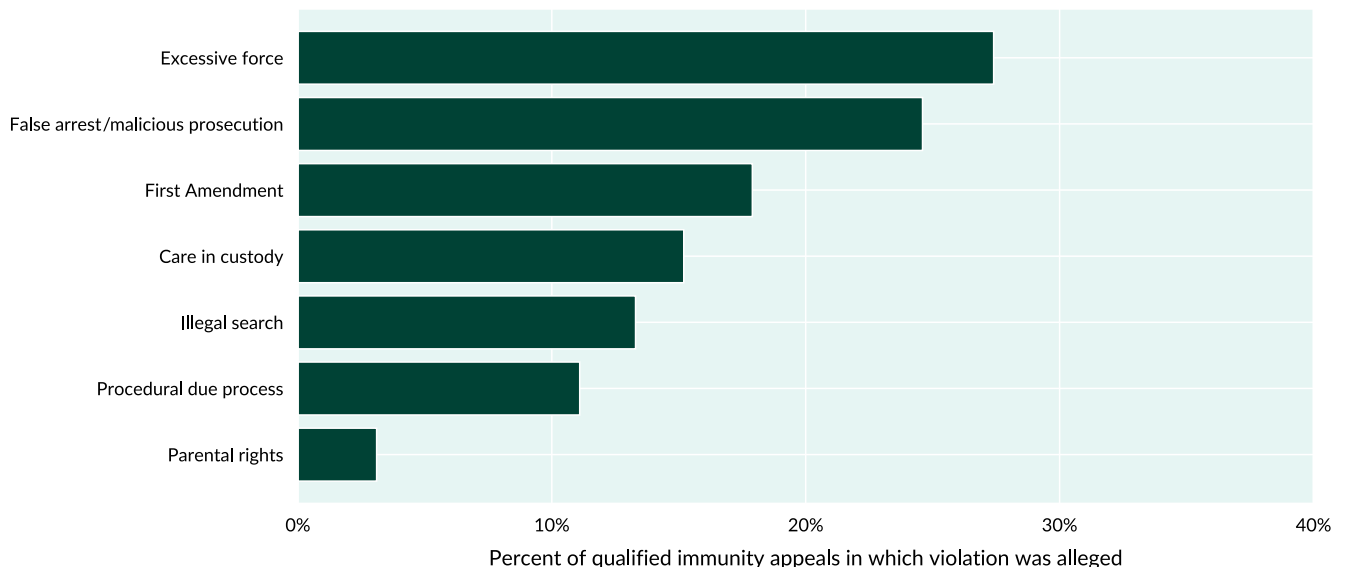
appeals represented some 96% of all qualified immunity appeals filed by defendants (see Figure 6) and roughly a third of qualified immunity appeals overall—nearly 2,000 total.

- Stage of litigation at time of appeal:** Appeals involving qualified immunity tend to occur early in lawsuits (see Figure 7). Twenty percent of the appeals we studied followed a ruling on a motion to dismiss filed by government defendants. These motions occur before the two sides in a lawsuit have exchanged information about their witnesses and evidence (i.e., before “discovery”). The point of these motions is to weed out any fatally flawed lawsuits (or individual claims). With qualified immunity, the fatal flaw that defendants argue at the motion to dismiss stage is that they simply cannot be sued because they are entitled to immunity.

Another 69% followed a ruling on a motion for summary judgment, which is a motion where one side argues that the facts are undisputed and the judge can rule without a trial to determine them; instead, this type of motion argues, the judge can simply apply the law to the known facts. Often, before deciding a summary judgment motion, judges permit some fact-finding through discovery if the two

Figure 4: Wide Array of Alleged Constitutional Violations

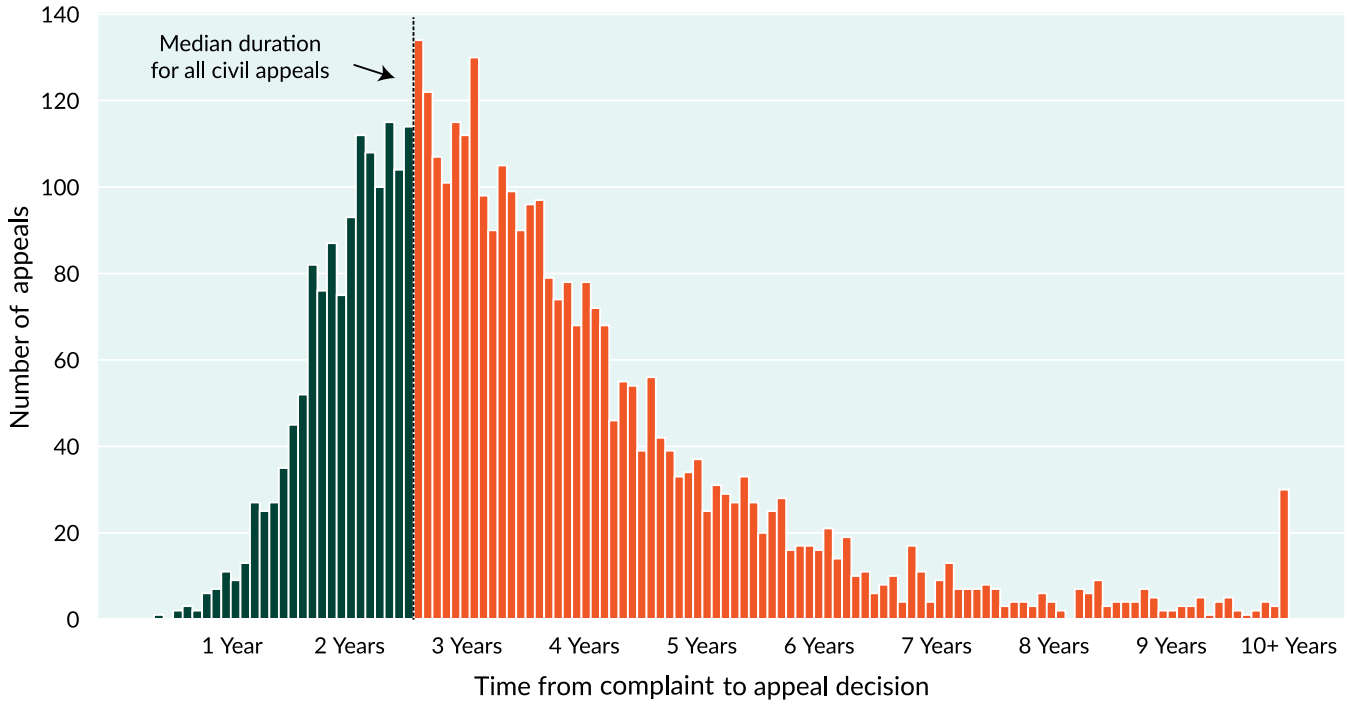
Excessive force violations are most commonly alleged, but First Amendment claims appear in 18% of appeals



Note: A single appeal may involve multiple alleged violations. Percentages therefore do not sum to 100%. In addition, we display only violation types we coded for. However, opinions sometimes involved other violation types (e.g., substantive due process).

Figure 5: Qualified Immunity Litigation Often Takes Years

Seventy percent of qualified immunity lawsuits on appeal had been open longer than the median civil appeal




Note: Lawsuits can have multiple appeals and thus appear multiple times in this chart.

Figure 6: Nearly All Defendant Appeals Are Interlocutory



Figure 7: Stage of Litigation at Time of Appeal

Most qualified immunity appeals occur prior to trial



sides disagree about key facts, though this is typically less than in a full-blown trial.

Just 4% of qualified immunity appeals occurred following a trial. In other words, only 4% happened after plaintiffs were allowed to make their full case before the court or a jury.

- **Self-representation rate:** Plaintiffs represented themselves without aid of an attorney in roughly 20% of qualified immunity appeals. However, as Figure 8 shows, the rate of self-represented plaintiffs varied considerably by circuit: Only 6% of plaintiffs represented themselves in the 1st Circuit, compared to 33% in the 5th.

- **Publication rate:** Circuit courts have discretion to publish, or not publish, their opinions. This matters because, in most circuits, only published opinions create binding precedent—that is, principles or rules, including clearly established law, that the court and the federal district courts under it are expected to follow in future cases.⁶² Overall, 35% of the qualified immunity opinions we studied were published, significantly more than the 13% publication rate across all civil appeals.⁶³ Three circuits went even further: The 1st, 7th, and 8th all published more than 75% of their qualified immunity opinions.⁶⁴ On the other hand, the 11th Circuit published fewer than 20% of its qualified immunity opinions. (See Figure 9.)

Figure 8: Plaintiffs Act as Their Own Attorneys at Different Rates Across Circuits





Figure 9: Circuits Publish Qualified Immunity Opinions at Vastly Different Rates

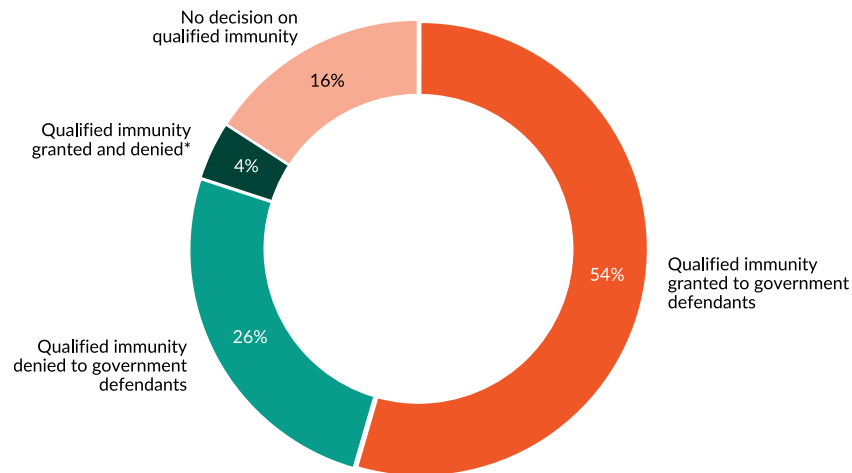


Outcomes

In qualified immunity appeals, government defendants win more than they lose. The circuit courts granted qualified immunity in 54% of appeals and denied it in just 26%. (For ease of expression, we refer to circuit courts “granting” and “denying” qualified immunity even when they are technically affirming or reversing the lower court’s grant or denial.) In the remaining appeals, the courts handed down mixed opinions (i.e., opinions with grants and denials of qualified immunity for different defendants or on different claims) or did not rule on qualified immunity at all. (See Figure 10.)


The federal appellate courts varied substantially in the rates at which they denied qualified immunity to government defendants, as Figure 11 shows. At opposite ends of the spectrum were the 5th Circuit, which denied qualified immunity in only 16% of appeals, and the 4th and 6th Circuits, which denied qualified immunity in 41%. Circuit courts similarly varied in their rates of granting qualified immunity.⁶⁵

Figure 10: Qualified Immunity Grants Outpace Denials



*Granted or denied for different defendants or claims

Figure 11: Circuits Deny Qualified Immunity at Different Rates

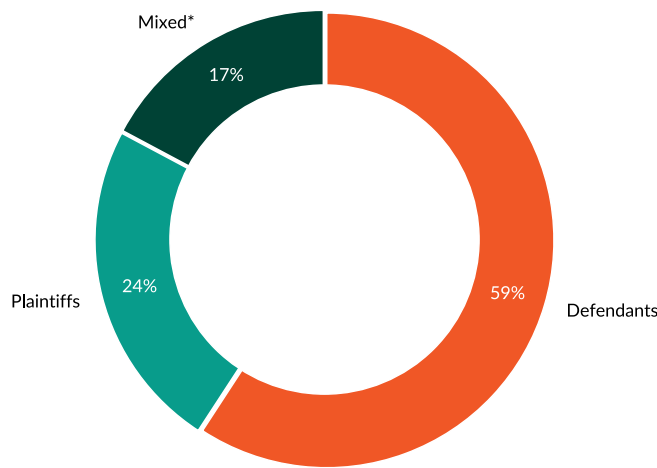


Note: Includes any denial of qualified immunity to a defendant or on an individual claim.

In addition to qualified immunity grants and denials, we looked at who prevailed on appeal—plaintiffs or defendants. We did this because a grant or denial may not always indicate who ultimately prevailed. For example, a plaintiff could have won because the court denied qualified immunity to the government defendant or because the court ruled in their favor on other grounds, such as denying a defendant’s appeal for lack of jurisdiction. Here again, government defendants saw more success than plaintiffs: 59% of appeals were resolved solely in favor of defendants, while 24% were resolved solely in favor of plaintiffs. The remaining appeals ended in mixed decisions. (See Figure 12.)

Appellate outcomes are even more tilted toward government defendants when we consider who was appealing from the district court, as Figure 13 illustrates.⁶⁶ When defendants won in district court, they nearly always won or achieved a mixed ruling on appeal; plaintiffs fully prevailed just 8% of the time. Plaintiffs had more success defending district court wins, fully prevailing 51% of the time, but government defendants still achieved a full or partial victory in nearly half of appeals after losing in the lower courts.

Figure 12: Government Defendants Usually Prevail in Qualified Immunity Appeals



*Both defendants and plaintiffs prevailed on aspects of the appeal

Figure 13: How District Court Decisions Translate Into Appellate Outcomes

When defendants win in district court, they rarely lose on appeal; when plaintiffs win in district court, the results are mixed



DISCUSSION: A BROAD AND UNNECESSARY SHIELD

In recent years, the once-obscure legal doctrine of qualified immunity has captured national attention. At the same time, as our data reveal, qualified immunity appeals have taken up a larger share of federal circuit court dockets. Because public discussion of the doctrine frequently focuses on notable cases of alleged police abuse, the debate often pits criminal justice reformers against law enforcement champions. Yet, digging more deeply into our data, we find that these conversations capture only a partial picture of qualified immunity litigation in federal circuit courts. Not only does the doctrine apply to a far wider array of government officials and conduct than is commonly understood; it also favors those officials and scuttles valid claims. Beyond these substantial downsides, our data, alongside research by others, suggest qualified immunity is a poor fit for achieving its proponents' goals.

Qualified Immunity Shields a Wider Array of Government Officials and Conduct Than Commonly Thought

A police officer who makes a split-second decision to use force against a suspect. A social worker who seizes a child without a warrant. A public university dean who restricts a student group's disfavored speech. A mayor who weaponizes local government to intimidate a political opponent. All of these government officials are eligible for protection under the one-size-fits-all doctrine of qualified immunity.

The Supreme Court intended for qualified immunity to protect government officials when they make reasonable mistakes, especially in tense or dangerous situations requiring quick thinking. And indeed, friends and foes of qualified immunity alike tend to assume it primarily protects police accused of excessive force, the most plausible type of case in which a government official might make a difficult decision under pressure. It is perhaps for this reason that, as detailed above, most prior research on qualified immunity focuses on cases with law enforcement defendants or excessive force claims.

A key advantage of our study is that we took a broader view, looking at all defendants and claims in qualified immunity appeals. We found that while law enforcement defendants and excessive force claims are well represented, so are other

defendants and claims. In fact, only 23% of the appeals we studied fit the classic mold of police accused of excessive force, showing that qualified immunity shields a far broader range of government defendants and conduct than many people think.

All told, more than a quarter of the appeals we studied involved non-police, non-prison defendants—and more than a fifth involved *only* such defendants. As noted in the results, mayors and city managers, university and school officials, prosecutors and judges, and child protective services workers were most common. Nevertheless, any government worker can invoke qualified immunity, and we saw state ethics commissioners, office of historic preservation officials, state legislators, zoning board members, horse racing commissioners, and septic system regulators listed as defendants in our hand-coded appeals. The job responsibilities of these government officials could hardly be more diverse. Yet qualified immunity applies to them all in the same way, whether they violated a person's rights during a chaotic confrontation in a dark alley or through calm deliberations in a comfortable office.

The constitutional violations claimed in qualified immunity appeals are similarly diverse. These include excessive force, false arrests, and illegal searches—policing-related violations that are perhaps more likely to involve heat-of-the-moment decisions. However, they also include alleged violations that seem less likely to involve such decisions, like procedural due process and First Amendment violations—the third largest category after excessive force and false arrests. This is not to say that First Amendment violations, for instance, never involve split-second decisions. But then, excessive force and other policing-related violations are sometimes premeditated. Our data are not granular enough to tell for every appeal whether split-second decisions were involved, but the broad categories are suggestive.

Only 23% of the appeals we studied fit the classic mold of police accused of excessive force, showing that qualified immunity shields a far broader range of government defendants and conduct than many people think.


Indeed, among a representative sample of 125 First Amendment appeals, we can tell that a sizable share likely do not involve a split-second decision.⁶⁷ As shown in Figure 14, plaintiffs in these appeals most often alleged they were victims of premeditated abuse by government officials in retaliation for protected First Amendment activity. Premeditated retaliation was the most common First Amendment violation alleged, appearing in roughly 59% of the appeals in our random sample.⁶⁸ In the remaining appeals, plaintiffs most often alleged that government officials directly restricted their speech or religious activity—for example, by arresting plaintiffs during a protest.⁶⁹

Digging deeper into premeditated retaliation claims, we found little reason to believe these allegations were less egregious than they appear at first blush. The largest share—nearly half—related to employment (see Figure 15). In these appeals,

government workers alleged retaliation from their superiors for many protected First Amendment activities, such as political association and speech, writing a book as a private citizen, union activity, speaking on matters of public concern, or even declining to speak on matters of public concern. Interestingly, police officers were the *plaintiffs* in over a third of these appeals, illustrating how qualified immunity can harm those it is often thought to protect, especially the rank and file.


The second largest share of premeditated retaliation claims, accounting for nearly a third, involved the direct targeting of private citizens over their protected First Amendment activity. In roughly three-quarters of these appeals, citizens alleged retaliation for speech that the Supreme Court has said is “at the very center of the constitutionally protected area of free discussion”—criticism of a government official, usually in a city council meeting or other public forum.⁷⁰ Police chiefs and police

Figure 14: Most First Amendment Appeals Allege Premeditated Retaliation for Protected Activity



Note: Data are based on a random sample of hand-coded appeals with First Amendment claims. For a claim to be categorized as premeditated retaliation, there had to be a delay between the protected activity and the alleged retaliatory action.

Figure 15: Premeditated Retaliation Appeals Are Mostly Alleged by Government Employees and Private Citizens



Our data show that, in practice, qualified immunity often shields government officials accused of weaponizing the power of the state to silence critics or shut down other speech they dislike.

especially in the heat of the moment—often shields government officials accused of weaponizing the power of the state to silence critics or shut down other speech they dislike. This perverse reality prompted Justice Clarence Thomas to ask, “Why should [government administrators], who have time to make calculated choices about enacting or enforcing unconstitutional policies, receive the same protection as a police officer who makes a split-second decision to use force in a dangerous setting?”⁷¹

He might have asked why government officials with time to consider their decisions sometimes receive *greater* protection than police officers tasked with making split-second decisions. This is because our data suggest alleged First Amendment violations get more protection than alleged excessive force violations. Government defendants challenging a district court loss fully prevailed in 34% of appeals with First Amendment claims but in only 23% of appeals with excessive force claims.⁷² While we cannot definitively attribute the difference to qualified immunity, it is suggestive—and concerning.⁷³

A possible explanation for the difference is that First Amendment claims are more factually diverse than excessive force claims. If true, this would make it harder for plaintiffs bringing First Amendment claims to pinpoint a prior case with sufficiently similar facts and thus to overcome qualified immunity.

Take the case of Institute for Justice client Anthony Novak. Anthony created a Facebook page poking fun at the police in his hometown of Parma, Ohio. Modeled after the police department’s real page, it had the same name, cover photo,

officers were the most common defendants, though by no means the only ones. University administrators and mayors were also frequent defendants.

Thus, our data show that, in practice, qualified immunity—a doctrine created to protect government officials when they make reasonable mistakes,

and profile picture but displayed the slogan “We no crime,” a parody of the department’s real slogan, “We know crime.” During the 12 hours the page was online, Anthony published six obviously parodic posts.

The Parma Police Department was not amused. Nearly a month after Anthony took the page down, police officers obtained a warrant for his arrest, searched his apartment, seized his electronics, and charged him with a felony under a state law that criminalizes using a computer to disrupt police operations. Anthony spent four days in jail before making bail. His case went to trial, and a jury found him not guilty.⁷⁴

What happened to Anthony was obviously retaliation for his parody and therefore unconstitutional. Nevertheless, the 6th Circuit, after first denying the police qualified immunity, later in the case reversed itself and granted it.⁷⁵ Why? Because, although his parodic posts were protected speech, Anthony had also taken steps to prevent others from spoiling the joke, deleting comments calling the page a fake and copying a notice from the police about the page. And, the second time it heard the case, the 6th Circuit held that Anthony had “not identified a case that clearly establishes deleting comments or copying the official warning is protected speech.”⁷⁶ Not only did qualified immunity give the police officers who retaliated against Anthony two bites at the apple, but the second bite succeeded because the court defined the right so narrowly as to make it impossible to find a prior case on point.

None of this is to say that alleged First Amendment violations, premeditated or otherwise, are the only kinds of suits that qualified immunity wrongfully excludes, much less that allegations of excessive force deserve less consideration by courts. It is merely to point out how frequently qualified immunity protects conduct far beyond the kinds of cases used to justify the doctrine.



Anthony Novak

Qualified Immunity Favors Government Defendants and Scuttles Valid Claims

In addition to shielding a wider range of officials and conduct than is commonly understood, qualified immunity may give government defendants the upper hand in civil rights litigation. In line with previous research, we found that qualified immunity grants outnumber denials at the appellate level, and government defendants prevail more often than plaintiffs.⁷⁷ Moreover, if plaintiffs lose in district court, they are especially unlikely to win on appeal. And even if plaintiffs win, this merely means they get one step closer to a trial—not necessarily a victory. Only 73 of the appeals we studied, or just 1.3%, gave plaintiffs a final, post-trial judgment in their favor.⁷⁸

One explanation could be that litigating against the government is hard, and this is certainly true. But our findings point to at least three ways in which qualified immunity may uniquely disadvantage plaintiffs—and do so in ways that prevent valid claims from being heard. First, qualified immunity makes the protection of rights depend not only on variation among federal circuits but also on their size and publication rates. Second, qualified immunity renders rights unclear, as even courts struggle to identify clearly established law on which successful claims can rest. Finally, qualified immunity awards special rights of appeal to government defendants, and they use these rights frequently. This gives them a huge advantage in litigation and significant leverage over plaintiffs. Combined, such challenges likely result in plaintiffs being turned away for reasons entirely unrelated to the merits of their claims or the culpability of defendants.

Arbitrarily Protecting Rights

A first way qualified immunity disadvantages plaintiffs is by making the protection of their rights depend on arbitrary factors. Instead of being the same nationwide, Americans' ability to vindicate our constitutional rights depends on the clearly established law in the federal circuit where we live. And the clearly established law in a circuit will depend on the types of appeals that happen to be brought in each circuit, as well as each circuit's differing approaches to those appeals. However, as our data show, the ability to vindicate our rights also depends on how many people live in a circuit and how often a circuit publishes its opinions. While these factors have little or nothing to do with legal reasoning or constitutional jurisprudence, they nonetheless make it harder for some plaintiffs to identify clearly established law to advance their claims.

Across the opinions we studied, there is substantial—and statistically significant—regional variation in the rates at which circuits grant and deny qualified immunity.⁷⁹ Similar variation has been found by previous research, including a 2020 study of how local factors affect plaintiff success in civil rights cases, a 2020 study of qualified immunity appeals with police excessive

force claims, and a 2021 study of more than 4,000 qualified immunity appeals.⁸⁰

Beyond the obvious factors driving circuit variation, such as types of cases appealed and differing responses to them, our data suggest two structural factors can affect how much clearly established law is found in a circuit.

First, the larger the population a circuit represents, the more appeals are likely to come before it—and the more appeals before a circuit, the more opportunities it has to create clearly established law that plaintiffs can rely on. And, indeed, the largest circuit by population represented, the 9th, accounts for nearly nine times as many qualified immunity appeals as the smallest, the 1st.⁸¹ In line with expectations, another Institute for Justice study found the 9th Circuit has substantially more clearly established law than the 1st Circuit.⁸²

Second, in most circuits, only published opinions create clearly established law. For this reason, differences in publication rates can translate into differences in the amount of clearly established law across circuits. Unsurprisingly then, higher-publishing circuits tend to have more clearly established law than lower-publishing ones, even after accounting for size. For example, the 7th and 8th Circuits, which publish more than three-quarters of their opinions, have over twice as many statements of clearly established law (216 and 236) as the similarly sized 3rd Circuit (101), which publishes only about a quarter of its opinions.⁸³

Because clearly established law is a circuit-by-circuit proposition, qualified immunity means the ability to vindicate our constitutional rights depends, to some extent, on where we live. This would be concerning enough if it were a function only of differences in the individual appeals that circuits consider and in the decisions reached in those appeals. When trivial factors such as a circuit's size and publication practices appear to affect our rights, it shows how qualified immunity can lead to the arbitrary administration of justice.⁸⁴

Making Rights Unclear

Not only do circuits differ in how they implement qualified immunity, leading to variation in the amount of clearly established law across the country, but their rulings often lack precision and clarity. This makes it hard for plaintiffs to identify relevant clearly established law and successfully argue that these precedents apply to their case. That, in turn, fuels the second challenge facing plaintiffs: Qualified immunity makes rights unclear.

Early in this project, we considered looking at how courts ruled on the two prongs of qualified immunity: (1) whether government defendants violated a constitutional right and (2) whether that right was clearly established.⁸⁵ Courts can rule on both prongs—or they can skip the first and go straight to the second. If they determine the right was not clearly established, they do not have to consider whether defendants violated the right at all. However, the experienced qualified immunity attorneys who coded the opinions found courts' rationales for granting qualified immunity were frequently ambiguous.

All told, our coders flagged as unclear roughly 1 in 4 opinions with qualified immunity grants, usually because courts had blurred their discussion of the two prongs.⁸⁶ Authors of other qualified immunity studies have reported similar problems:

- A 2009 study acknowledged that, “in many cases regarding the Fourth Amendment’s restrictions on searches and seizures, it was often unclear whether a court was resolving the substantive issue of whether a search was ‘reasonable’ or whether it was addressing whether the right was ‘clearly established.’”⁸⁷
- A 2011 study examining qualified immunity outcomes similarly noted that “some courts essentially merged the two prongs, making it difficult or impossible to classify which prong was outcome-determinative.”⁸⁸
- And a 2015 study concluded that “it is sometimes difficult to determine whether a court is resolving a claim on the constitutional merits or on qualified immunity grounds, especially for unpublished decisions.”⁸⁹

In some instances, our coders flagged opinions as unclear because courts did not state a rationale for granting qualified immunity at all—they simply affirmed a district court’s decision without discussion. Although courts were silent as to why they granted qualified immunity, these cases could be substantial.

For example, in *Young v. Borders*, police searching for a suspect mistakenly went to the wrong apartment in the middle of the night, banged on the door allegedly without identifying themselves as police, and shot to death an innocent occupant when he opened the door. The district court in the case granted the police defendants qualified immunity, and the man’s estate appealed.⁹⁰ The original 11th Circuit panel disposed of the appeal in a one-paragraph ruling, finding the defendants were entitled to qualified immunity.⁹¹ Yet the facts of the case were sufficiently troubling that, when the appeal came up for en banc review, it triggered a 6,000-word dissent and a separate 1,000-word dissent, forcing the original panelists to write a 7,000-word concurrence better explaining their original decision.⁹²

That so many qualified immunity rulings appear so sloppy or sparse in their reasoning speaks to confusion among the courts about how to even implement the doctrine. Worse, the imprecise, unclear, or absent reasoning in these rulings directly impedes plaintiffs’ ability to press their claims and win, even when their claims are strong. It also makes it difficult to say what our constitutional rights even are.

Special Appeals, Special Advantages

A third way qualified immunity gives government defendants the upper hand is by granting them special rights to file interlocutory appeals immediately after a district court denies qualified immunity. These interim appeals can be filed at any stage of litigation—or even at multiple stages—and they effectively give defendants an extra chance (or two or three) to press a claim of qualified immunity before different judges. Our data show these appeals are common and suggest they are

a factor in drawing out litigation, often by years. Thus, these appeals run the risk of wearing down worthy plaintiffs with extended litigation.

Interlocutory appeals make up a substantial portion of qualified immunity litigation: more than a third—nearly 2,000 appeals—from 2010 through 2020.⁹³ And they appear to be growing: Our data show that interlocutory appeals, along with qualified immunity appeals generally, increased over our study period.⁹⁴ Clearly, government defendants do not hesitate to take advantage of these special rights. Indeed, interlocutory appeals represented 96% of all appeals filed by defendants.

The prevalence of interlocutory appeals may explain why the median duration of a qualified immunity lawsuit in our dataset was three years and two months, 23% longer than the typical civil suit up on federal appeal. Prior research also suggests that interlocutory appeals contribute to lawsuit length, averaging more than a year—441 days—from filing to resolution.⁹⁵

To illustrate how this special right of appeal can be used to draw out litigation, consider appeals filed at the motion to dismiss stage, which make up a fifth of all interlocutory appeals. Motions to dismiss are supposed to be hard for the government to win, as they are intended to weed out only the weakest cases. Normally, if a district court judges a case strong enough to survive such a motion, it moves forward to summary judgment or trial. But if a district court judges a case strong enough to continue by denying a request for qualified immunity, government defendants can immediately file an interlocutory appeal and hope a three-judge appellate panel will decide the question differently. If they do not, this special appeal did nothing more than extend the litigation.

Not only can government defendants immediately appeal denials of qualified immunity, but they can also do so multiple times in the same lawsuit. This allows them to relitigate qualified immunity over and over. The plight of one Nevada family shows how this can play out. After a government administrator stole from their deceased father, two men filed a federal lawsuit in 2007 to hold the administrator accountable.⁹⁶ They did not get a final judgment entered in their favor until 2019—12 years later. It took more than a decade for the family to get justice because they had to run a gauntlet of procedural delays due to qualified immunity:

- After the family filed suit, the government administrator moved to dismiss the lawsuit based on qualified immunity.⁹⁷ Although the district court denied the motion, the defendant’s subsequent interlocutory appeal put the lawsuit on hold for two and a half years until the 9th Circuit affirmed the denial in 2011.⁹⁸
- In late 2012, the defendant filed a motion for summary judgment, another pretrial motion,

Our data show interlocutory appeals are common and suggest they are a factor in drawing out litigation, often by years. Thus, they run the risk of wearing down worthy plaintiffs with extended litigation.

invoking qualified immunity. This, too, was denied.⁹⁹ Once again, the defendant filed an interlocutory appeal, and once again the 9th Circuit affirmed the denial, though not until 2015.¹⁰⁰

- The lawsuit finally made it to trial in 2015, eight years after it started, and a jury awarded the plaintiffs \$2.1 million.¹⁰¹ But they *still* were not done with qualified immunity. The defendant appealed the final verdict—raising qualified immunity a third time—and received yet another denial in 2018.¹⁰² Final judgment in favor of the plaintiffs was finally satisfied in May 2019, almost exactly 12 years after the lawsuit began.¹⁰³

Although the number was excessive, the defendant in that case at least raised permissible issues in his interlocutory appeals. This is not always the case. In fact, we found evidence that government defendants may use interlocutory appeals strategically, filing meritless appeals simply to drag out litigation: 11% of interlocutory appeals were dismissed because they failed to raise an issue within the scope of the appellate court's authority. Others have raised concerns that defendants might abuse the advantages provided by interlocutory appeals. For example, civil rights attorneys interviewed for a 2020 study on qualified immunity described how interlocutory appeals were used to “wear . . . out” and “beat down the plaintiffs’ counsel.”¹⁰⁴

But even when defendants’ failure to raise an appealable issue is without malice, the impact on plaintiffs is the same: higher costs and delayed litigation, which may force them to settle or even drop their case. Yet when plaintiffs stick it out and win these appeals, all this means is that they can continue their lawsuit—and their fight against qualified immunity may not even be over.

Combined, these factors—the protection of rights turning on arbitrary factors, rights made unclear, and special advantages for government defendants—likely result in plaintiffs being turned away for reasons entirely unrelated to the merits of their claims or the culpability of defendants. Although the Supreme Court insisted that, in creating the doctrine, “we provide no license to lawless conduct,” in practice qualified immunity can arbitrarily shield officials from accountability while forcing victorious plaintiffs to endure years of litigation and vexatious appeals before receiving justice.¹⁰⁵

Qualified Immunity Is a Poor Fit for Achieving Its Goals

The Supreme Court created qualified immunity to achieve two goals: (1) preserve government officials’ ability to perform their duties without fear, especially in split-second situations, and (2) protect them from harassment, distraction, and the threat of financial harm from a flood of specious lawsuits. Unfortunately, it is not at all clear—from our results and others’—that qualified immunity achieves these aims, or that

the doctrine is necessary at all. And in addition to being a poor fit for achieving its goals, qualified immunity clogs up the courts with time-wasting appeals.

On the first goal, the problem is not just that qualified immunity is overly broad, protecting conduct far outside tough decisions in the heat of the moment. It is also that the clearly established standard fails at its core purpose: ensuring government officials have clear notice about what constitutes unconstitutional conduct before they are held accountable for it. As described above, the experienced qualified immunity attorneys who hand coded opinions for this study frequently struggled to untangle the reasoning behind qualified immunity grants, a problem encountered by other legal researchers. Even judges have acknowledged the difficulty of applying qualified immunity: Judge Charles R. Wilson of the 11th Circuit went so far as to say that “wading through the doctrine of qualified immunity is one of the most morally and conceptually challenging tasks federal appellate court judges routinely face.”¹⁰⁶

If legal experts struggle to make sense of qualified immunity, it is unreasonable to expect the average government official to comprehend all the clearly established law in their circuit—let alone apply it in tense, rapidly evolving situations.¹⁰⁷ As one district judge put it, “It strains credulity to believe that a reasonable officer, as he is approaching a suspect to arrest, is thinking to himself: ‘Are the facts here anything like the facts in *York v. City of Las Cruces*?’”¹⁰⁸ As noted earlier, previous research has shown that police are not trained on the specifics of clearly established law, but even if they were, the doctrine is so complex that such training is unlikely to be effective.¹⁰⁹

As for the second goal, protecting government officials from the burdens of possibly meritless litigation, it is again unclear that qualified immunity is necessary, let alone a good fit, for the purpose. First, as noted earlier, prior research has found that government officials are nearly always indemnified by their employers, so even those found to have violated citizens’ rights rarely face financial penalties.

Second, as qualified immunity plays out in the courts, it often does not protect officials from the practical burdens of litigation, most notably discovery, nor is it well designed to do so. In fact, the appeals in our dataset more commonly came at the summary judgment stage of litigation—when courts typically have already allowed discovery—than at the motion to dismiss stage. The gap is significant, with nearly 70% of appeals following a ruling on summary judgment compared to just 20% following a motion to dismiss, a finding similar to prior research on district court dockets.¹¹⁰ To be sure, our data are not granular enough to tell how often discovery took place prior to the appeals in our dataset, but there is likely a sizable share of such cases.¹¹¹ In cases like these, qualified immunity is not protecting officials from the burdens of discovery, let alone

If legal experts struggle to make sense of qualified immunity, it is unreasonable to expect the average government official to comprehend all the clearly established law in their circuit—let alone apply it in tense, rapidly evolving situations.

quickly disposing of meritless claims.

Despite the Supreme Court's hope of cutting off litigation early, it makes sense that qualified immunity would be more of a factor later in litigation. This is because qualified immunity often arises in complicated situations—including those involving split-second decisions—when factual development is especially important for deciding a case, making it difficult to resolve before discovery or trial. Anecdotally, our coders observed appellate courts concluding exactly that when denying qualified immunity at the summary judgment stage—disputed facts required a trial. The fact-intensive nature of a qualified immunity inquiry may also help explain why we found that defendants have less success facing excessive force than First Amendment allegations, as excessive force claims may be more likely to involve disputed facts and situations in which the only witnesses are the plaintiffs and the defendants, making these cases ill suited for quick resolution.¹¹²

Third, courts have other tools for weeding out non-viable claims.¹¹³ One is basic pleading standards. For a civil rights lawsuit, a plaintiff must allege concrete facts that, if true, could plausibly give rise to a constitutional violation. If they fail to do so, then the case is dismissed—no qualified immunity necessary.

Another tool is protection afforded by the Constitution to government officials who make reasonable mistakes in difficult situations. For example, under the Fourth Amendment, only *unreasonable* searches, seizures, and force are prohibited.¹¹⁴ This reasonableness standard has been interpreted to give government officials—especially those facing difficult and dangerous situations—leeway in their decisions. As stated by the Supreme Court in a prominent excessive force case: “The calculus of reasonableness must embody allowance for the fact that police officers are often forced to make split-second

judgments—in circumstances that are tense, uncertain, and rapidly evolving.”¹¹⁵ If there is no constitutional violation, then the lawsuit can be disposed of on this ground without invoking qualified immunity.

And in fact, among the appeals in our dataset where government defendants did not receive qualified immunity but nevertheless prevailed, our coders observed that appellate courts often relied on such grounds to dispose of meritless lawsuits.¹¹⁶ Instead of ruling on qualified immunity, they often ruled that no constitutional violation had occurred or that the plaintiffs' initial lawsuit filings were deficient for reasons unrelated to qualified immunity. This is consistent with earlier research showing that when defendants raised qualified immunity in pretrial motions in district courts, those courts were more likely to grant the motions on grounds *other than* qualified immunity.¹¹⁷

Not only does qualified immunity fail to achieve its goals, but it also brings a substantial downside: The special right of interlocutory appeal clogs the circuit courts with extra appeals. Above, we described how interlocutory appeals can lengthen litigation and wear down plaintiffs, even if they have strong cases. Here, we note how it can create a burden for appellate courts—one that is becoming more common, according to our data. Without the special treatment given to qualified immunity defendants, the 2,000 interlocutory appeals in our dataset would not exist. For good reason, these special immediate appeals are supposed to be a “narrow exception” to the normal order of business.¹¹⁸ Yet, as we found, government defendants in qualified immunity lawsuits use them frequently, adding an estimated cumulative total of over 2,300 *years* to qualified immunity litigation over our study period.¹¹⁹

CONCLUSION

Judge Don Willett of the 5th Circuit summed up the problem with qualified immunity well:

Plaintiffs must produce precedent even as fewer courts are producing precedent. Important constitutional questions go unanswered precisely because no one's answered them before. Courts then rely on that judicial silence to conclude there's no equivalent case on the books. No precedent = no clearly established law = no liability. An Escherian Stairwell. Heads defendants win, tails plaintiffs lose.¹²⁰

And lose plaintiffs do—even those who have had their constitutional rights violated in the most egregious ways, as these examples of policing abuses from our dataset illustrate:

- *Jessop v. City of Fresno*: Officers allegedly stole over \$200,000 in cash and rare coins from the plaintiffs while executing a search warrant. The 9th Circuit granted qualified immunity because there was no previous case holding it unconstitutional to steal property during a search.¹²¹
- *Collie v. Barron*: A police officer shot and paralyzed David Collie, an innocent, unarmed man who matched the broad description (shirtless, black) of suspects in a nearby robbery. The officer shot him in the back as he pulled his hand out of his pocket to indicate where he was headed. Although the 5th Circuit noted that “this tragic case exemplifies an individual’s being in the wrong place at the wrong time,” it nonetheless stated that “under current governing law, we must affirm [the grant of qualified immunity].”¹²²
- *Cass v. City of Abilene*: Police obtained a warrant to search a gold exchange store for a possible misdemeanor reporting violation, allegedly in retaliation for speech by one of the owners, Marcus Cass, criticizing a policy they favored. Instead of simply serving the warrant, the defendants used Cass and his co-owner’s supposed “anti-police”

attitude as an excuse to conduct a tactical raid. Allegedly without announcing himself as a police officer, an officer entered Cass’ office with his gun drawn and his badge out of view. Seeing only an armed man dressed in street clothes and body armor, Cass did not know he was police and drew his own firearm. He was shot and ultimately died from his wounds. The 5th Circuit granted the defendants qualified immunity.¹²³

These and many other cases show how courts can grant qualified immunity even when police officers blatantly violate the Constitution. However, as we found, qualified immunity’s protections do not stop with police. The doctrine in fact shields a much wider array of government officials than commonly thought, from social workers to university deans to mayors and beyond. It also shields a wide array of conduct, from allegations of punishment for protected speech—as in the ordeals faced by Sylvia Gonzalez and Anthony Novak—to alleged due process violations, including cases with extreme consequences:

- *Cope v. Cogdill*: A suicidal inmate was placed in a cell with a 30-inch telephone cord, which he used to strangle himself as the sole jailer on duty watched. The jailer failed to call emergency services after the inmate became unresponsive. The 5th Circuit noted that “watching an inmate attempt suicide and failing to call for emergency medical assistance is not a reasonable response.” Nevertheless, given the lack of sufficiently similar case law, it granted the jailer qualified immunity.¹²⁴
- *Sampson v. County of Los Angeles*: A woman who had recently taken custody of her niece was repeatedly sexually harassed by a social worker assigned to the case. The 9th Circuit noted that the right to be free from sexual harassment by public officials was clearly established in workplaces and schools. However, because the sexual harassment occurred in the context of providing social services, the court granted the social worker qualified immunity. In doing so, the court lamented that “the Supreme Court’s

exceedingly narrow interpretation of what constitutes a ‘clearly established’ right precludes us from holding what is otherwise obvious to us.”¹²⁵

Qualified immunity adds disadvantages for plaintiffs even beyond the challenge of identifying clearly established precedents—making the ability to vindicate rights turn on arbitrary factors, making rights unclear, and affording government defendants special advantages. And it is not clear that qualified immunity even serves the goals motivating the Supreme Court’s creation of, and continued support for, the doctrine.

Our results add to a growing body of research finding qualified immunity unacceptably burdens plaintiffs and fails at its goals. This strengthens the argument for the Supreme Court to overturn *Harlow*, the case in which it created the doctrine four decades ago. If it will not, Congress can and should act to end qualified immunity.

In addition, state legislatures and city councils can create their own causes of action to hold government officials accountable and ban qualified immunity as a potential defense in state or local civil rights lawsuits.¹²⁶ To date, two states have taken such action: Colorado banned the defense for police officers, while New Mexico banned it for all government workers.¹²⁷ In a similar vein, New York City created a cause of action allowing victims of unreasonable search and seizure or excessive force to

sue New York Police Department employees; importantly, the legislation disallows the qualified immunity defense.¹²⁸

Short of ending qualified immunity, the Supreme Court could temper the doctrine by broadening the clearly established standard to encompass not just precedent from the Supreme Court and the same federal circuit but also precedent from the other circuits and state courts, as well as government-issued guidance. This would make the ability to vindicate rights less arbitrary, and it would also prevent absurdities like government officials receiving qualified immunity for conduct their own agency explicitly disallows simply because the relevant circuit court has not yet encountered similar conduct.¹²⁹

Better yet, the Supreme Court could eliminate the clearly established standard and instead allow courts to evaluate claims based on whether a government official was acting in good faith. Ending qualified immunity altogether would better ensure that government officials who violate the Constitution are held accountable. However, under such a good faith standard, at least intentional and obvious constitutional violations would not receive protection.

Whether through outright abolition or significant reform, courts and lawmakers can and should act to end the unbounded impunity allowed by the current doctrine of qualified immunity.



Sylvia Gonzalez

APPENDIX A: DEVELOPING THE PREDICTIVE ALGORITHMS

In this appendix, we describe our process for developing the predictive algorithms, beginning with how we turned the opinions into data. We then describe how we built the algorithms. And we conclude with how we ran the final algorithms on all 7,173 opinions in our dataset.

Turning Opinions Into Data

Before we could build our predictive algorithms, we needed to turn the opinion text into data the algorithms could manipulate. To do this, we imported, indexed, and annotated the text of the opinions as described below.

Importing the Opinions

After downloading the opinions (in Rich Text Format) from Westlaw into a local directory, we read them into a Python 3.7 development environment using the `striptrf` Python library.¹³⁰ One small complication was that Westlaw frequently stores important information as hyperlinks, which the `striptrf` reader was unable to parse. To fix this, we wrote a short function that converted the hyperlinks into plain text.

At the end of the importing process, we were left with an unorganized block of text like the example shown in Figure A1.

Figure A1: Initial Format of an Imported Opinion

```
'\n\n\n567 F.3d 672\nUnited States Court of Appeals,\nDistrict of Columbia Circuit.\nPeter James ATHERTON, Appellant\nv.\nDISTRICT OF COLUMBIA OFFICE OF the MAYOR, et al., Appellees.\nNo. 07-5195.\nArgued March 6, 2009.\nDecided June 2, 2009.\nRehearing En Banc Denied Aug. 25, 2009.\nSynopsis\nBackground: Dismissed grand juror brought civil rights action against District of Columbia, juror officer, Assistant United States Attorney, and others, alleging he was unlawfully removed from grand jury service because of his deliberative judgments and ethnicity. Defendants moved to dismiss. The United States District Court for the District of Columbia, Henry H. Kennedy, Jr., J., 2007 WL 1041659, granted motions. Juror appealed.\nHoldings: The Court of Appeals, Edwards, Senior Circuit Judge, held that:\njuror officer was not absolutely immune from suit;\nassistant United States attorney was not absolutely immune from suit;\njuror failed to state a claim under § 1983 that his dismissal violated equal protection; and\njuror failed to state a claim under § 1985 that defendants conspired to remove juror from grand jury for ethnic purposes.\nAffirmed in part, reversed in part, and remanded.\nProcedural Posture(s): On Appeal; Motion to Dismiss; Motion to Dismiss for Failure to State a Claim.\n*676 Appeal from the United States District Court for the District of Columbia (No. 04cv00680).\nAttorneys and Law Firms\nSarah Campbell, Student Counsel, argued the cause as amicus curiae in support of appellant. With her on the briefs were James E. Coleman, Jr., appointed by the court, Sean E. Andru
```

© 2021 Thomson Reuters.

Hierarchical Indexing

Because these text blocks were difficult to use and manipulate, our next step was to organize them using hierarchical indexing. Once we successfully read an opinion into our development environment, we structured it into an ordered hierarchy that consisted of sections, paragraphs, sentences, and words.¹³¹ Through this structuring, we assigned each word in the opinion a unique index or location—for example, a word might be the fifth word in the first sentence of the third paragraph of the introduction section.

For the paragraph, sentence, and word indexing, we used standard methods (e.g., splitting on newline characters) or pre-programmed functions from the Natural Language Toolkit (NLTK) library in Python.¹³² However, we created the section indexing specific to the structure of the opinions as initially downloaded from Westlaw. It included (among other sections) the header, the Westlaw synopsis/background (if applicable), the main text of the opinion, the footnotes, and any concurrences and dissents (which Westlaw often includes within the same document). In addition, we broke out both the introduction and conclusion of the opinion section using separate subfunctions. We stored each word and its accompanying index in tabular form, as illustrated in Table A1.

Table A1: Example of Indexed Data Stored in Tabular Form

Word	Section	Para Num	Sent Num	Word Num
Qualified	Opinion	1	1	1
immunity	Opinion	1	1	2
frequently	Opinion	1	1	3
involves	Opinion	1	1	4
alleged	Opinion	1	1	5
1 st	Opinion	1	1	6
Amendment	Opinion	1	1	7
violations.	Opinion	1	1	8
These	Opinion	1	2	1
violations	Opinion	1	2	2
are	Opinion	1	2	3
often	Opinion	1	2	4
alleged	Opinion	1	2	5
against	Opinion	1	2	6
non-police,	Opinion	1	2	7
non-prison	Opinion	1	2	8
defendants.	Opinion	1	2	9

This indexing allowed us to group the text in different ways. For example, in certain cases, it was helpful to keep each word stored separately. In others, it was preferable for each row in the table to list an entire sentence. Using the indexing, we could easily change the table above into groups of sentences, paragraphs, or even entire sections. Table A2 shows how the data from Table A1 could be grouped into sentences instead of broken out by word.

Table A2: Example of Indexed Data Grouped by Sentence

Sentence	Section	Para Num	Sent Num
Qualified immunity frequently involves alleged 1 st Amendment violations.	Opinion	1	1
These violations are often alleged against non-police, non-prison defendants.	Opinion	1	2

Model-based approaches for prediction (described in more detail below) were more likely to rely on tables with each word listed separately, while rules-based approaches were more likely to rely on tables with sentences or entire paragraphs grouped together.

Annotating the Opinions

Aside from hierarchical indexing, we added several annotations to each word in a table:

- The word with all capitalization and punctuation removed.
- The part of speech. We added these annotations using the NLTK library’s part-of-speech tagger.¹³³
- A binary indicator for whether the word included (or was) a number. For example, if “1st” appeared in the text, we tagged the word as including a number.
- The word stem or root. For example, the stem of the word “drafting” is “draft.” We used the Porter word stemmer from the NTLK library to perform the stemming.
- A binary indicator for whether the word was a “stop” word—that is, a word that has little inherent meaning or performs a mostly grammatical function (e.g., “a,” “the,” “which,” etc.). We used the NLTK library’s built-in list of stop words supplemented with our own manually developed list of words that appear frequently in legal opinions.¹³⁴

The result was a table that included each word, its hierarchical index location, and all additional annotations.¹³⁵ For a brief sample of what this looked like, see Table A3, where blue represents the hierarchical index and green represents the additional annotations.¹³⁶

Table A3: Example of Indexed Data With Annotations

Word	Section	Para Num	Sent Num	Word Num	Term	Part of Speech	Number?	Word Stem	Stop Word?
Qualified	Opinion	1	1	1	qualified	adj.	0	qualif	0
immunity	Opinion	1	1	2	immunity	noun	0	immun	0
frequently	Opinion	1	1	3	frequently	adv.	0	frequent	0
involves	Opinion	1	1	4	involves	verb	0	involv	0
alleged	Opinion	1	1	5	alleged	adj.	0	alleg	0
1 st	Opinion	1	1	6	1st	adj.	1	1st	0
Amendment	Opinion	1	1	7	amendment	noun	0	amend	0
violations.	Opinion	1	1	8	violations	noun	0	violat	0
These	Opinion	1	2	1	these	pronoun	0	these	1
violations	Opinion	1	2	2	violations	noun	0	violat	0
are	Opinion	1	2	3	are	verb	0	are	1
often	Opinion	1	2	4	often	adv.	0	often	0
alleged	Opinion	1	2	5	alleged	verb	0	alleg	0
against	Opinion	1	2	6	against	prep.	0	against	0
non-police,	Opinion	1	2	7	non-police	noun	0	non-police	0
non-prison	Opinion	1	2	8	non-prison	noun	0	non-prison	0
defendants.	Opinion	1	2	9	defendants	noun	0	defend	1

Ultimately, we did not use all the annotations in our predictive algorithms. For example, we generally found that word stems hurt performance during the development process, so we avoided including them in our final algorithms. Nevertheless, having these annotations available allowed for more flexibility in the algorithm-development process.

Building the Predictive Algorithms

Once we had turned the opinions into data and indexed and annotated them, we could build our predictive algorithms. We randomly sampled roughly 11% of our 7,173-opinion dataset for model development and evaluation, using stratified sampling to ensure each year was proportionally represented.¹³⁷ We hand coded these 791 opinions across 34 fields and assigned them to one of two samples using random stratified sampling by year:

- Training sample: This consisted of 604 opinions we used to develop the predictive algorithms. We further divided these opinions into two sub-groups, again using random stratified sampling by year:
 - Primary training sample: We used 529 opinions to build the predictive algorithms.
 - Validation sample: We used 75 opinions to assess performance during the development process and fine-tune the algorithms.
- Testing sample: This consisted of 187 opinions we used solely for the final evaluation of the algorithms’ performance. Because we used these only once the algorithms were final, these holdout test data were not part of the development process.

We used the primary training sample to build our predictive algorithms, constructing the algorithms for each field separately. To do so, we used four different approaches:

Rules-based extraction: We extracted, rather than predicted, seven of the 34 fields. This method of obtaining information began with narrowing the search area. For example, most basic qualitative information was present at the top of an opinion, so we searched only headers for these fields.

Next, we used “regular expressions,” a special programming tool that enables powerful and focused text searches, to find and extract the desired information.¹³⁸

For example, one of our simplest rules-based extraction algorithms targeted the circuit court that issued the opinion. To find this information, we searched an opinion’s header for a paragraph that started with “United States Court.” Using regular expressions, we then extracted the text that came immediately after that phrase in the paragraph (e.g., “United States Court, *Sixth Circuit*”). If no further text was present in the paragraph, we took the first words of the line immediately after, as those typically referred to the circuit court in question.

Rules-based prediction: In addition to using rules to extract information, we used them to predict 11 of the 34 fields. These rules most often took the form of “if . . . then” statements: If certain phrases were found in a particular part of an opinion, then we would classify the opinion a certain way. For example, one of our most effective rules for predicting whether an opinion involved qualified immunity was a search for the court’s recitation of the two-pronged standard of review—that is, something like “(1) . . . constitutional violation . . . (2) . . . clearly established.”¹³⁹ If this text pattern appeared, we classified the opinion as involving qualified immunity.

Rules-based approaches relied heavily on hierarchical indexing (and especially the section indexing) to narrow the search area, as specific pieces of information often showed up in certain parts of an opinion document. For example, defendants frequently appeared in the first paragraph of the main section of an opinion. To predict whether an opinion dealt with state law enforcement defendants, we therefore searched sentences in the opinion’s introduction for phrases such as “against . . . police.”

We again used regular expressions to create flexible searches. Sometimes, this flexibility meant searching for many synonyms or variations. In the state law enforcement defendant example above (“against . . . police”), we searched for several similar law enforcement officials in addition to police: “highway patrol officers,” “state troopers,” “sheriffs,” and so forth. We also used regular expressions to avoid false positives. Continuing with the same example, if the word “police” was preceded by “capitol,” we did not count it as a match since Capitol Police are federal, not state, officials.

Given the variety of language used by courts, individual rules like these rarely found all or even most opinions with a characteristic we sought. Instead, we often strung separate rules together, with each rule tuned to avoid false positives. For example, our algorithm for predicting whether an appeal dealt with state law enforcement defendants used eight separate rules, all tuned to avoid false positives.

Model-based prediction: Beyond rules, we used statistical models to predict three of the 34 fields.

To build these models, we used algorithms such as naïve bayes, penalized logistic regression (generally, ridge logistic regression), and support vector machines. (We used different models for different fields.) Specifically, we fed into the models small collections of highly relevant words and phrases identified using feedback from our experienced qualified immunity attorneys and data exploration methods.¹⁴⁰ For example, our model for predicting whether an opinion dealt with excessive force violations included only five highly specific inputs: three related to common “use of force” phrases such as “excessive force” and “deadly force,” one related to the standard of review for prison excessive force appeals, and one related to overly tight handcuffs.

Since words and phrases will usually appear more often in longer opinions than in shorter ones, we normalized the

frequency counts by dividing by the number of words in the opinion.¹⁴¹ For example, if a word appeared three times in a 500-word opinion, we would divide the raw frequency (3) by the number of words (500).

In addition to dividing by the word count, we divided by the word’s standard deviation in our primary training sample.¹⁴² This step was necessary given how we compiled our model inputs.¹⁴³

To achieve satisfactory performance, we tuned our models using five-fold cross-validation.¹⁴⁴ We used cross-validation to pick the model type, engineer features (i.e., choose model inputs), and calibrate various “knobs and dials” on certain models.¹⁴⁵ We also conducted internal assessments of performance using the validation sample as an additional assurance that the models worked well on unseen opinions. Although we did use naïve bayes and support vector machine models for certain fields, we generally found, through cross-validation, that ridge logistic regression models performed best.¹⁴⁶

Rules-model hybrids: For the remaining 13 of the 34 fields, we used a hybrid method involving both rules and models. Typically, we integrated the two approaches by tuning both the rules and the models to avoid false positives. If the model predicted that a field was present, that became the final prediction. For example, if the model predicted the opinion dealt with a First Amendment violation, that is how the algorithm classified it.

But if the model predicted that the field was not present—that the opinion did not deal with a First Amendment violation—then the algorithm turned to a rules-based method to classify the opinion.

This approach proved effective because the rules and models tended to target slightly different types of opinions: Models worked well when distinctive vocabularies were present throughout an opinion, while rules worked well when highly specific phrases were present.

For example, the model that predicted whether an opinion involved qualified immunity was excellent when an opinion extensively discussed qualified immunity and clearly established rights. But not all relevant opinions featured substantial discussion of qualified immunity. In some, the court focused on other issues but then granted qualified immunity in the alternative in a footnote using a highly specific language pattern such as “in the alternative . . . we hold defendants are entitled to qualified immunity.” This is not enough text for a model to make reliable predictions, but we could write a rule to find patterns like it.

Running the Predictive Algorithms

Once the algorithms were complete, we put them into a broader Python function that created the dataset. This function first read in the opinion, organizing and indexing it as described above. Next, it sequentially predicted the 34 fields. The function repeated this process for each opinion in the file directory, resulting in a final dataset that included predictions for every opinion and field.

Also once the algorithms were complete, we formally evaluated their effectiveness using the 187-opinion holdout test sample. Comprehensive statistics for how each field performed on this holdout test set are available in Appendix B.

Finally, after we assessed the algorithms' performance, we performed minor manual cleanup on our final dataset.¹⁴⁷ Readers interested in the programming code we used to generate the final dataset can contact us through our website at <https://ij.org/report/unaccountable/data-downloads/>.¹⁴⁸

APPENDIX B: PREDICTIVE ALGORITHMS' PERFORMANCE ON HOLDOUT TEST DATA

This appendix provides a comprehensive overview of our predictive algorithms' performance on every field in this study. For the relevance field (i.e., whether qualified immunity was raised on appeal), we calculated performance statistics based on the entire 187-opinion testing sample. For all other fields, we calculated performance statistics based on the 162 relevant opinions from the testing sample. We did not use these holdout test data at any point during the algorithm-development process. This ensured that the algorithms' performance on these opinions was representative of their performance overall.

To measure performance, we used the following statistics:

- **Accuracy:** Accuracy measures the number of correct predictions out of the total number of predictions for a field. For example, if an algorithm predicts a field 100 times and 91 predictions are correct, then the accuracy is 91%. Accuracy is our primary metric for evaluating text and categorical fields. Although we report accuracy for all predicted fields, other metrics—precision, recall, and F1 score—are better indicators of performance for binary (yes/no) fields.
- **Precision:** In a nutshell, precision measures how good an algorithm is at avoiding false positives. For example, if an algorithm predicts 20 opinions involve state law enforcement officials, but only 18 actually do (meaning two were false positives), then the precision is 0.9 (18/20). The maximum value is 1. Precision is applicable only for binary fields.
- **Recall:** Put simply, recall measures how good an algorithm is at picking out the field in question. For example, if there are 25 true interlocutory appeals and the algorithm correctly identifies 23 of them, then recall is 0.92 (23/25). The maximum value is 1. Recall is applicable only for binary fields.
- **F1 Score:** The F1 score is a widely used performance metric that combines precision and recall into a single statistic by taking their harmonic mean. A harmonic mean is a mean that penalizes divergence between the values being averaged. For example, the harmonic mean of 0.8 and 0.8 is 0.8, but the harmonic mean of 0.6 and 1 is 0.75, because 0.6 and 1 have greater divergence than 0.8 and 0.8. A good F1 score indicates that both precision and recall performed well. The maximum value is 1.
- **Confusion Matrix:** We also provide confusion matrices, which give additional context for an algorithm's performance. Although not technically a statistic, a confusion matrix compares an algorithm's predictions to the true values by putting them into a table. In our matrices, values shaded in green are correct predictions—that is, predictions that matched the true value. Blue cells are false positives (predicted positives but true negatives). Yellow cells are false negatives (predicted negatives but true positives), although for categorical variables, all incorrect cells are shaded yellow. Qualitative fields (e.g., plaintiffs, defendants) do not have confusion matrices.

	Field	Type	Algorithm Prediction Method	Accuracy	Precision	Recall	F1 Score	Confusion Matrix Truth: y-axis; Prediction: x-axis																								
Basic Information	Relevance (qualified immunity raised on appeal)	Binary	Rules-Model Hybrid	90.9%	0.980	0.914	0.95	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>22</td><td>3</td></tr> <tr><td>1</td><td>14</td><td>148</td></tr> </table>		0	1	0	22	3	1	14	148															
		0	1																													
	0	22	3																													
	1	14	148																													
	Circuit Court	Text	Rules-Based Extraction	100%	--	--	--	--																								
	Circuit Court Case Number	Text	Rules-Based Extraction	99.4%	--	--	--	--																								
	Opinion Date	Text (date)	Rules-Based Extraction	100%	--	--	--	--																								
	Plaintiffs	Text	Rules-Based Extraction	99.4%	--	--	--	--																								
	Defendants	Text	Rules-Based Extraction	100%	--	--	--	--																								
Judges	Text	Rules-Based Extraction	100%	--	--	--	--																									
District Court of Origin	Text	Rules-Based Extraction	99.4%	--	--	--	--																									
District Court Case Number	Text	Manually Coded	--	--	--	--	--																									
Case Origination Date	Text (date)	Manually Coded	--	--	--	--	--																									
Procedural Details	Appellants	Categorical P – Plaintiffs B – Both (cross-appellants) D – Defendants	Rules-Based Prediction	99.4%	--	--	--	<table border="1"> <tr><td></td><td>P</td><td>B</td><td>D</td></tr> <tr><td>P</td><td>107</td><td>1</td><td>0</td></tr> <tr><td>B</td><td>0</td><td>4</td><td>0</td></tr> <tr><td>D</td><td>0</td><td>0</td><td>50</td></tr> </table>		P	B	D	P	107	1	0	B	0	4	0	D	0	0	50								
		P	B	D																												
	P	107	1	0																												
	B	0	4	0																												
	D	0	0	50																												
	Published	Binary	Rules-Based Prediction	100%	1.000	1.000	1.00	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>99</td><td>0</td></tr> <tr><td>1</td><td>0</td><td>63</td></tr> </table>		0	1	0	99	0	1	0	63															
		0	1																													
0	99	0																														
1	0	63																														
En Banc	Binary	Rules-Based Prediction	100%	1.000	1.000	1.00	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>160</td><td>0</td></tr> <tr><td>1</td><td>0</td><td>2</td></tr> </table>		0	1	0	160	0	1	0	2																
	0	1																														
0	160	0																														
1	0	2																														
Interlocutory Appeal	Binary	Rules-Based Prediction	99.4%	1.000	0.979	0.99	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>114</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>47</td></tr> </table>		0	1	0	114	0	1	1	47																
	0	1																														
0	114	0																														
1	1	47																														
Pro Se Plaintiffs (self-represented plaintiffs)	Categorical 0 – No pro se in lawsuit ES – Pro se at earlier stage 1 – All pro se in appeal	Rules-Based Prediction	99.4%	--	--	--	<table border="1"> <tr><td></td><td>0</td><td>ES</td><td>1</td></tr> <tr><td>0</td><td>122</td><td>0</td><td>0</td></tr> <tr><td>ES</td><td>0</td><td>15</td><td>1</td></tr> <tr><td>1</td><td>0</td><td>0</td><td>24</td></tr> </table>		0	ES	1	0	122	0	0	ES	0	15	1	1	0	0	24									
	0	ES	1																													
0	122	0	0																													
ES	0	15	1																													
1	0	0	24																													
Case Stage (at time of appeal)**	Categorical SJ – Summary Judgment D – Dismissal B – Both PT – Post-trial	Rules-Model Hybrid	94.4%	--	--	--	<table border="1"> <tr><td></td><td>SJ</td><td>D</td><td>B</td><td>PT</td></tr> <tr><td>SJ</td><td>108</td><td>0</td><td>5</td><td>1</td></tr> <tr><td>D</td><td>0</td><td>30</td><td>0</td><td>0</td></tr> <tr><td>B</td><td>2</td><td>1</td><td>8</td><td>0</td></tr> <tr><td>PT</td><td>0</td><td>0</td><td>0</td><td>7</td></tr> </table>		SJ	D	B	PT	SJ	108	0	5	1	D	0	30	0	0	B	2	1	8	0	PT	0	0	0	7
	SJ	D	B	PT																												
SJ	108	0	5	1																												
D	0	30	0	0																												
B	2	1	8	0																												
PT	0	0	0	7																												
Government Level of Defendants	Categorical S – State B – Both F – Federal	Rules-Based Prediction	96.9%	--	--	--	<table border="1"> <tr><td></td><td>S</td><td>B</td><td>F</td></tr> <tr><td>S</td><td>153</td><td>3</td><td>1</td></tr> <tr><td>B</td><td>0</td><td>0</td><td>0</td></tr> <tr><td>F</td><td>1</td><td>0</td><td>4</td></tr> </table>		S	B	F	S	153	3	1	B	0	0	0	F	1	0	4									
	S	B	F																													
S	153	3	1																													
B	0	0	0																													
F	1	0	4																													
Government Defendant Type	State Law Enforcement Defendants	Binary	Rules-Model Hybrid	95.1%	0.990	0.932	0.96	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>58</td><td>1</td></tr> <tr><td>1</td><td>7</td><td>96</td></tr> </table>		0	1	0	58	1	1	7	96															
		0	1																													
	0	58	1																													
	1	7	96																													
	Federal Law Enforcement Defendants	Binary	Rules-Model Hybrid	99.4%	0.000	--*	--*	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>161</td><td>1</td></tr> <tr><td>1</td><td>0</td><td>0</td></tr> </table>		0	1	0	161	1	1	0	0															
		0	1																													
	0	161	1																													
1	0	0																														
State Prison Defendants	Binary	Rules-Model Hybrid	97.5%	1.000	0.875	0.93	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>130</td><td>0</td></tr> <tr><td>1</td><td>4</td><td>28</td></tr> </table>		0	1	0	130	0	1	4	28																
	0	1																														
0	130	0																														
1	4	28																														
Federal Prison Defendants	Binary	Rules-Model Hybrid	99.4%	1.000	0.750	0.86	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>158</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>3</td></tr> </table>		0	1	0	158	0	1	1	3																
	0	1																														
0	158	0																														
1	1	3																														
Other Defendants	Binary	Rules-Based Prediction	90.7%	0.826	0.844	0.84	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>109</td><td>8</td></tr> <tr><td>1</td><td>7</td><td>38</td></tr> </table>		0	1	0	109	8	1	7	38																
	0	1																														
0	109	8																														
1	7	38																														
Task Force Defendants	Binary	Rules-Based Prediction	100%	--*	--*	--*	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>162</td><td>0</td></tr> <tr><td>1</td><td>0</td><td>0</td></tr> </table>		0	1	0	162	0	1	0	0																
	0	1																														
0	162	0																														
1	0	0																														

	Field	Type	Algorithm Prediction Method	Accuracy	Precision	Recall	F1 Score	Confusion Matrix Truth: y-axis; Prediction: x-axis																
Constitutional Violation Type	First Amendment Violations	Binary	Rules-Model Hybrid	99.4%	1.000	0.969	0.98	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>130</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>31</td></tr> </table>		0	1	0	130	0	1	1	31							
		0	1																					
	0	130	0																					
	1	1	31																					
	Religious Liberty Violations	Binary	Rules-Model Hybrid	100%	1.000	1.000	1.00	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>158</td><td>0</td></tr> <tr><td>1</td><td>0</td><td>4</td></tr> </table>		0	1	0	158	0	1	0	4							
		0	1																					
	0	158	0																					
	1	0	4																					
	Excessive Force Violations	Binary	Rules-Model Hybrid	98.1%	0.980	0.961	0.97	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>110</td><td>1</td></tr> <tr><td>1</td><td>2</td><td>49</td></tr> </table>		0	1	0	110	1	1	2	49							
	0	1																						
0	110	1																						
1	2	49																						
False Arrest Violations	Binary	Rules-Model Hybrid	91.4%	0.975	0.750	0.85	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>109</td><td>1</td></tr> <tr><td>1</td><td>13</td><td>39</td></tr> </table>		0	1	0	109	1	1	13	39								
	0	1																						
0	109	1																						
1	13	39																						
Illegal Search Violations	Binary	Rules-Model Hybrid	92.6%	0.714	0.714	0.71	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>135</td><td>6</td></tr> <tr><td>1</td><td>6</td><td>15</td></tr> </table>		0	1	0	135	6	1	6	15								
	0	1																						
0	135	6																						
1	6	15																						
Procedural Due Process Violations	Binary	Rules-Model Hybrid	96.9%	0.737	1.000	0.85	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>143</td><td>5</td></tr> <tr><td>1</td><td>0</td><td>14</td></tr> </table>		0	1	0	143	5	1	0	14								
	0	1																						
0	143	5																						
1	0	14																						
Care in Custody Violations	Binary	Model-Based Prediction	95.7%	0.941	0.727	0.82	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>139</td><td>1</td></tr> <tr><td>1</td><td>6</td><td>16</td></tr> </table>		0	1	0	139	1	1	6	16								
	0	1																						
0	139	1																						
1	6	16																						
Parental Rights Violations	Binary	Model-Based Prediction	99.4%	1.000	0.800	0.89	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>157</td><td>0</td></tr> <tr><td>1</td><td>1</td><td>4</td></tr> </table>		0	1	0	157	0	1	1	4								
	0	1																						
0	157	0																						
1	1	4																						
Employment Violations	Binary	Model-Based Prediction	95.7%	1.000	0.682	0.81	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>140</td><td>0</td></tr> <tr><td>1</td><td>7</td><td>15</td></tr> </table>		0	1	0	140	0	1	7	15								
	0	1																						
0	140	0																						
1	7	15																						
Outcomes	Overall Prevailing Party	Categorical P – Plaintiffs M – Mixed D – Defendants	Rules-Model Hybrid	96.3%	--	--	--	<table border="1"> <tr><td></td><td>P</td><td>M</td><td>D</td></tr> <tr><td>P</td><td>30</td><td>3</td><td>0</td></tr> <tr><td>M</td><td>1</td><td>26</td><td>0</td></tr> <tr><td>D</td><td>1</td><td>1</td><td>100</td></tr> </table>		P	M	D	P	30	3	0	M	1	26	0	D	1	1	100
		P	M	D																				
	P	30	3	0																				
	M	1	26	0																				
D	1	1	100																					
Qualified Immunity Granted	Binary	Rules-Based Prediction	90.1%	0.914	0.914	0.91	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>61</td><td>8</td></tr> <tr><td>1</td><td>8</td><td>85</td></tr> </table>		0	1	0	61	8	1	8	85								
	0	1																						
0	61	8																						
1	8	85																						
Qualified Immunity Denied	Binary	Rules-Based Prediction	92.6%	0.902	0.822	0.86	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>113</td><td>4</td></tr> <tr><td>1</td><td>8</td><td>37</td></tr> </table>		0	1	0	113	4	1	8	37								
	0	1																						
0	113	4																						
1	8	37																						
Lack of Jurisdiction – Factual Dispute	Binary	Rules-Based Prediction	98.1%	1.000	0.667	0.80	<table border="1"> <tr><td></td><td>0</td><td>1</td></tr> <tr><td>0</td><td>153</td><td>0</td></tr> <tr><td>1</td><td>3</td><td>6</td></tr> </table>		0	1	0	153	0	1	3	6								
	0	1																						
0	153	0																						
1	3	6																						

*Statistic cannot be calculated due to insufficient data.

**For the “case stage” field, the confusion matrix does not reflect the “other” category as it did not appear as either an actual or a predicted case stage.

APPENDIX C: KEY DATA BY CIRCUIT COURT

This appendix breaks out data about qualified immunity appeals for each federal court of appeals except the Federal Circuit, which generally does not hear qualified immunity cases, and the District of Columbia Circuit, which had few relevant appeals. Disaggregating such a small sample would not yield reliable estimates. (The 41 qualified immunity appeals from the District of Columbia Circuit are included in the “all circuits” row of this appendix’s tables, however.) Because dividing our main dataset into 11 subsets magnifies the potential impact of incorrect algorithm predictions, we have largely included only fields with excellent algorithmic performance.¹⁴⁹

Table C1: Basic Facts

Circuit	Number of Appeals	Number of Appeals per Million Pop.	Opinion Publication Rate	Percent With Plaintiffs Representing Themselves	Percent Where Plaintiffs Appealed a District Court Loss	Percent Where Defendants Appealed a District Court Loss
1st	105	7.5	86%	6%	64%	34%
2nd	382	16.0	35%	18%	61%	35%
3rd	335	14.8	25%	22%	68%	30%
4th	244	7.8	41%	13%	58%	36%
5th	708	20.2	29%	33%	70%	28%
6th	819	25.1	32%	9%	49%	46%
7th	273	10.8	77%	22%	71%	28%
8th	407	19.2	81%	11%	53%	44%
9th	936	14.4	22%	24%	63%	33%
10th	540	30.1	33%	25%	64%	34%
11th	736	20.9	17%	18%	60%	38%
All Circuits*	5,526	17.0	35%	20%	61%	36%

*The “all circuits” row reflects data for all qualified immunity appeals in our dataset, including those from the District of Columbia Circuit.

Note: The percentage of appeals that involved cross-appellants is not shown. Accordingly, the percentages of appeals where plaintiffs and defendants appealed a district court loss do not sum to 100%.

Table C2: Defendants and Allegations

Circuit	Percent With Law Enforcement Defendants	Percent With Prison Defendants	Percent With Non-Police, Non-Prison Defendants	Percent With Excessive Force Allegations	Percent With First Amendment Allegations
1st	61%	10%	31%	18%	19%
2nd	57%	23%	32%	14%	21%
3rd	57%	20%	33%	18%	26%
4th	62%	24%	24%	29%	18%
5th	55%	30%	25%	29%	16%
6th	63%	19%	26%	34%	18%
7th	54%	29%	24%	16%	20%
8th	65%	25%	20%	33%	16%
9th	52%	25%	29%	28%	18%
10th	56%	31%	27%	23%	19%
11th	70%	21%	18%	36%	13%
All Circuits*	59%	24%	26%	27%	18%

*The “all circuits” row reflects data for all qualified immunity appeals in our dataset, including those from the District of Columbia Circuit.

Note: For simplicity, we have combined state and federal law enforcement defendants into a single law enforcement defendants column, and state and federal prison defendants into a single prison defendants column. These results should be considered reasonable approximations as opposed to precise values.

Table C3: Qualified Immunity Grants and Denials

Circuit	Appeal With QI Grant	Appeal With QI Denial
1st	56%	26%
2nd	56%	27%
3rd	56%	20%
4th	46%	41%
5th	67%	16%
6th	48%	41%
7th	52%	32%
8th	63%	30%
9th	61%	31%
10th	66%	26%
11th	60%	34%
All Circuits*	59%	30%

*The “all circuits” row reflects data for all qualified immunity appeals in our dataset, including those from the District of Columbia Circuit.

Table C4: Overall Prevailing Party

Circuit	Plaintiff Success Rates			Defendant Success Rates		
	All Appeals	When Appealing a District Court Loss	When Defending a District Court Win	All Appeals	When Appealing a District Court Loss	When Defending a District Court Win
1st	21%	5%	53%	61%	33%	76%
2nd	26%	11%	52%	57%	36%	69%
3rd	18%	8%	43%	64%	42%	74%
4th	34%	13%	67%	42%	18%	58%
5th	16%	4%	46%	72%	40%	86%
6th	34%	12%	59%	47%	23%	68%
7th	25%	14%	51%	59%	42%	66%
8th	22%	7%	40%	63%	43%	80%
9th	22%	10%	44%	57%	36%	67%
10th	19%	4%	49%	64%	34%	79%
11th	25%	6%	54%	61%	29%	82%
All Circuits*	24%	8%	51%	59%	33%	74%

*The “all circuits” row reflects data for all qualified immunity appeals in our dataset, including those from the District of Columbia Circuit.

Note: This table uses the “prevailing party” field, which records the party who won an appeal overall. Corresponding columns (i.e., the two “all appeals” columns; the plaintiff loss and defendant win columns; and the plaintiff win and defendant loss columns) will not sum to 100%, as we do not display mixed outcomes (where both plaintiffs and defendants won on different aspects of an appeal). Mixed outcomes can be inferred based on the data shown.

APPENDIX D: MAIN CODEBOOK

This appendix contains definitions for all the fields we hand coded and subsequently predicted with our algorithms. We provided it to all hand coders to use as a guide during the coding process. We have made minor edits for clarity.

Basic Relevance

Relevance: Was qualified immunity raised on appeal in this opinion? Phrased another way, was it clear the court *could have* ruled on qualified immunity even if it did not? Relevant opinions included those where the court (1) directly discussed qualified immunity as a legal issue; (2) noted qualified immunity as an alternate justification; and (3) noted qualified immunity as moot based on its other rulings in the opinion.

- 1 – Qualified immunity was raised on appeal in this opinion, either by the defendants or the court itself. As noted above, this included appeals where qualified immunity was before the court but was not ruled upon. (Note: If qualified immunity was granted for a claim in district court, including in the alternative, and then that claim was appealed, we considered the opinion relevant.)
- 0 – Qualified immunity was not before the court on appeal. Examples include:
 - Opinions that were not addressing claims under 42 U.S. Code Sections 1983 or 1985 (which covers conspiracies to violate a person’s civil rights) or claims under *Bivens v. Six Unknown Named Agents* (i.e., opinions that had nothing to do with qualified immunity or that raised it solely with respect to state law claims).
 - Opinions that were part of Section 1983, 1985, or *Bivens* lawsuits but where qualified immunity was not raised as a defense.
 - Opinions where there was *previously* an interlocutory appeal of qualified immunity, but the *instant* appeal related to issues at trial that were not relevant to qualified immunity.
 - Opinions where qualified immunity was raised at the district court stage, but qualified immunity was not among the areas being appealed (e.g., the appeal related to a sanction for attorney fees in a case that previously involved qualified immunity).
 - Opinions where it was unclear if qualified immunity was being appealed (the default

code was “0” unless it was reasonably clear qualified immunity was raised on appeal).

- Opinions in a post-trial appeal that was related to qualified immunity but was not an appeal of qualified immunity itself (e.g., an appeal of jury instructions provided regarding qualified immunity).
- Opinions where qualified immunity was raised but was waived because the defendants failed to adequately brief the issue for the court.

If qualified immunity was not before the court in the opinion (i.e., if relevance was entered as “0”), we did not complete the following sections.

Basic Information

Circuit_court: The circuit court for the appeal.

Circuit_case_no: The circuit court case number for the appeal. We captured this exactly as it appeared in the opinion.

Date: The date the opinion was filed/decided. If only one date was listed in the header of the opinion, we assumed it was the date filed/decided. If the opinion was amended or modified, we used the amended date.

Plaintiffs: The plaintiffs in the opinion. We recorded this exactly as it appeared in the header of the opinion, with the exception that we removed the trailing words “plaintiff” and “appellant” or “appellee.” (Though we use the plural throughout, there could be only one plaintiff.)

Defendants: The defendants in the opinion. We recorded this exactly as it appeared in the header of the opinion. (Though we use the plural throughout, there could be only one defendant.)

Judges: The judges who heard the appeal. We recorded this exactly as it appeared in the header of the opinion, barring cleanup of extraneous text and punctuation.

District_court: The district court where the appeal originated.

District_court_case_no: The case number of the lawsuit in district court. We recorded this as “NP” (not present) if the district court case number did not appear in the opinion.

Procedural Details

Appellants: Which party was appealing the district court's decision?

- P – The plaintiffs were appealing the district court's decision.
- D – The defendants were appealing the district court's decision.
- B – Both parties were appealing the district court's decision (i.e., there were cross-appellants).

Published: Was the opinion published or unpublished?

- 1 – The opinion was published.
- 0 – The opinion was unpublished.

En_banc: Did the opinion involve an en banc hearing?

- 1 – The opinion involved an en banc hearing.
- 0 – The opinion did *not* involve an en banc hearing. This included instances where an en banc hearing was denied. We also used this category if it was not immediately clear whether an en banc hearing occurred.

Interloc: Was the appeal an interlocutory appeal?

- 1 – The appeal was interlocutory. Unless the opinion indicated otherwise, we considered any immediate appeal of a *denial* of summary judgment or dismissal to be an interlocutory appeal, even if the phrase “interlocutory appeal” did not explicitly appear.
- 0 – The appeal was *not* interlocutory.

Pro_se: Did the lawsuit include pro se plaintiffs?

- 1 – All plaintiffs were listed as pro se for the appeal (this included attorneys who represented themselves).
- 0 – No plaintiffs were and, *as best we could determine from the instant opinion*, had never been pro se at any point in the lawsuit. (We used this option unless it was clear the plaintiffs *were* pro se at some point in the lawsuit.)
- ES (earlier stage) – The plaintiffs were previously (at some point in the lawsuit) listed as pro se but were being represented by an attorney on appeal. Such opinions occasionally listed the plaintiffs as pro se. However, Westlaw editors listed separate attorneys for the plaintiffs, suggesting that the plaintiffs were not pro se for the full duration of the appeal. (Additionally, we used this option if one plaintiff in the appeal represented themselves, but another plaintiff in the appeal was represented by an attorney. This rarely occurred in our data.)

Case_stage: What was the procedural stage of the lawsuit at the time of the appeal? We did not consider motions to alter or amend judgment, motions for reconsideration, motions to amend the complaint, any appellate motions (e.g., motions

to dismiss the appeal for a lack of jurisdiction), or any other motions not specified below.

- D – Dismissal. This also included motions on the pleadings and failure to state a claim.
- SJ – Summary Judgment.
- B (stands for both) – Dismissal and Summary Judgment concurrently (some cases featured appeals of both at the same time).
- PT (stands for post-trial) – Any opinion where at least some of the claims proceeded to trial and were being appealed.
- Other – Anything that did not fit neatly into one of the categories above.

Government Defendant Type

The fields below pertain to the types of government officials sued. A few notes:

- We coded only for defendants who were part of the appeal. For example, if a lawsuit previously included both state prison and state law enforcement defendants, but the appeal only included the state prison defendants, then we coded the government defendants solely as “state_prison.”
- We did *not* code official-capacity defendants or government entities such as municipalities, state boards, and so forth.
- We included all types of defendants who were part of the appeal, regardless of their connection to qualified immunity. (That being said, it was extremely rare for qualified immunity to be raised for some types of defendants but not others.)
- We always considered police defendants such as sheriffs, deputies, detectives, and so forth to be law enforcement officials, regardless of the exact role they were performing when sued (e.g., we still considered a sheriff operating a detention center to be a law enforcement defendant). Similarly, we always considered corrections officers and jailers to be prison officials.
- We categorized any state officer working in a law enforcement role (even if they were not police) as a state law enforcement officer, regardless of their agency of employment. (Generally, we considered officials with the power to arrest and/or obtain warrants law enforcement officials.)
- We considered probation and parole officers to be law enforcement officials.
- We considered officials working in conjunction with a prison system (e.g., a lawyer with the department of corrections, a chaplain within a prison) prison officials.

Gov_level: Were the government officials being sued federal or state/local officials (i.e., was the lawsuit brought under *Bivens* or Section 1983)?

- Fed – The government officials being sued were federal officials.
- State – The government officials being sued were state/local government officials.
- Both – Both federal and state/local officials were being sued.

State_LEO: Was a state/local law enforcement officer listed as a defendant? (They did *not* need to be the sole defendant type.) This included sheriffs sued in conjunction with the operation of a jail or police officers assisting in pre-trial detention. However, we did *not* count corrections officers in prisons or jails as law enforcement.

- 1 – State/local law enforcement officials were included as defendants.
- 0 – State/local law enforcement officials were *not* included as defendants.

Federal_LEO: Was a federal law enforcement official (e.g., FBI agent, Park Service officer, Secret Service agent) listed as a defendant? (They did *not* need to be the sole defendant type.) Corrections officers in federal prisons were not counted as law enforcement.

- 1 – Federal law enforcement officials were included as defendants.
- 0 – Federal law enforcement officials were *not* included as defendants.

State_prison: Was a state/local prison official (e.g., a warden, corrections officer, jailer) listed as a defendant? (They did *not* need to be the sole defendant type.) This included corrections officers involved in pretrial detention.

- 1 – State/local prison officials were included as defendants.
- 0 – State/local prison officials were *not* included as defendants.

Federal_prison: Was a federal prison official (e.g., a warden, corrections officer) listed as a defendant? (They did *not* need to be the sole defendant type.)

- 1 – Federal prison officials were included as defendants.
- 0 – Federal prison officials were *not* included as defendants.

Other_ind_capacity_def: Was a non-law enforcement, non-prison official (e.g., child protective services officials, fire fighters, city government officials, state regulatory board officials, state legislators) listed as a defendant in this opinion? (They did *not* need to be the sole defendant type.)

- 1 – Non-law enforcement, non-prison officials were included as defendants.
- 0 – Non-law enforcement, non-prison officials were *not* included as defendants.

Task_force: If a defendant was law enforcement, was a state/federal task force involved or did the appeal involve a cross-deputized officer (i.e., a state/local officer designated to act as a federal officer)?

- 1 – At least one defendant was part of a federal/state task force or was cross-deputized (or both).
- 0 – A federal/state task force was not involved in this appeal, nor were any officers cross-deputized.

Constitutional Violation Type

The fields below pertain to the types of substantive constitutional violations alleged by the defendants, regardless of the constitutional amendment under which a claim was brought. A few notes:

- These fields were *not* mutually exclusive—that is, a single opinion could have contained multiple claim types below.
- Our coding was not exhaustive. That is, we coded only for the categories listed below, even if an appeal contained additional alleged violations (e.g., substantive due process, Second Amendment). And some appeals included none of the violations we coded for.
- We included constitutional violations before the court regardless of their connection to qualified immunity. (That being said, it was relatively rare for qualified immunity to be raised for only some claims and not others.)
- We coded violations alleged and raised in the appeal, regardless of whether they were proven. We ignored violations that were part of the lawsuit but not raised on appeal. In general, we coded for the specific violations plaintiffs alleged, even if the facts in the opinion suggested other or additional violations. For example, if a claim generally involved personal property that was seized, but the claim was brought solely under the due process clause (alleging the plaintiffs were deprived of a proper procedure), we classified the opinion solely as procedural due process—we did not code for “illegal search” (which covered illegal seizing of personal property), since the specific violation alleged by the plaintiffs related solely to an improper procedure. (An exception to this practice occurred for excessive force violations framed as illegal seizures, as detailed below.)
- We had to be reasonably confident to code a type of constitutional violation. We permitted strong educated inferences in the absence of explicit clarity; if substantial uncertainty existed, however, we coded the field as “0.”

First Amendment: Did the plaintiffs allege violations of their rights to free speech (e.g., retaliation for expressing their First Amendment rights), free association, religious liberty, or any other First Amendment right?

- 1 – First Amendment violations were alleged in this opinion.
- 0 – No First Amendment violations were alleged in this opinion.

Religious liberty: Did the plaintiffs allege violations of their right to freely practice their religion? This field is a sub-field of the “First Amendment” field above. (Note: Although Religious Land Use and Institutionalized Persons Act, or RLUIPA, claims are not technically constitutional claims, we captured them under “religious liberty” if they involved a prisoner’s right to religious freedom.)

- 1 – Religious liberty violations were alleged in this opinion.
- 0 – No religious liberty violations were alleged in this opinion.

Excessive force: Did the plaintiffs allege that the defendants committed a violation related to excessive force? This included cases involving non-traditional defendants such as corrections officers or government administrators. We still categorized excessive force claims framed as unauthorized seizures as “excessive force,” since technically all excessive force claims can be considered unauthorized seizures.

- 1 – Excessive force violations were alleged in this opinion.
- 0 – No excessive force violations were alleged in this opinion.

False arrest: Did the plaintiffs allege that the defendants committed violations related to a false arrest or malicious prosecution? This category also included illegal seizures of a person that did not result in an arrest (e.g., an illegal traffic stop)—with the exception that removals of a child from a guardian’s custody are covered instead by the “parental rights” category.

- 1 – False arrest or malicious prosecution violations were alleged in this opinion.
- 0 – No false arrest or malicious prosecution violations were alleged in this opinion.

Illegal search: Did the plaintiffs allege that the defendants committed violations related to an illegal search? This category also included the alleged illegal seizure of personal property.

- 1 – Illegal search violations were alleged in this opinion.
- 0 – No illegal search violations were alleged in this opinion.

Procedural due process: Did the plaintiffs allege they were deprived of fair process under the due process requirements of the Constitution? Because a large number of violations

could have potentially included procedural due process, we only coded this field if either (1) procedural due process was explicitly mentioned as an alleged constitutional violation or (2) due process was explicitly mentioned, and it was clear that the due process violation was procedural in nature. (Note: We considered *Brady* violations for a failure to disclose exculpatory evidence to be procedural due process violations.)

- 1 – Procedural due process violations were alleged in this opinion.
- 0 – No procedural due process violations were alleged in this opinion.

Care in custody: Did the alleged violations relate to the (lack of) care provided for the plaintiffs when they were in some form of custody? Deliberate indifference to medical needs in prison, or a failure to protect an inmate from harm by another inmate, fell under this category. These violations must have specifically related to a lack of care or attention, however. Excessive force cases against corrections officers, First Amendment retaliation against prisoners, or cases restricting a prisoner’s religious liberty did not necessarily fall under this category. Additionally, “custody” was not limited to jail or prison and included, for example, temporary detentions by police officers.

- 1 – Violations relating to care in custody were alleged in this opinion.
- 0 – No violations relating to care in custody were alleged in this opinion.

Parental rights: Did the plaintiffs allege that the defendants interfered with their rights as parents? Child custody cases fell under this category, but cases involving direct violations of minors’ rights (e.g., using excessive force on a minor child) did not necessarily implicate parental rights.

- 1 – Violations of parental rights were alleged in this opinion.
- 0 – No parental rights violations were alleged in this opinion.

Employment: Were at least some of the alleged violations of constitutional rights in this opinion related to an adverse employment action (e.g., a termination, demotion, salary decrease), a hostile work environment, or unsafe workplace conditions? We did not consider contractors suing a government agency over a contract dispute (e.g., an IT company contracted with a local government suing that government) to be “employment.”

- 1 – At least some alleged violations in this opinion related to an adverse employment action or a hostile/unsafe work environment.
- 0 – No alleged violations in this opinion related to an adverse employment action or a hostile/unsafe work environment.

Outcomes

Prevail: Who was the prevailing party in the opinion? (Note: We did not consider attorney fees or damages in determining the prevailing party.)

- P – The plaintiffs were the prevailing party.
- D – The defendants were the prevailing party.
- M – Both the defendants and plaintiffs prevailed in parts of the opinion. (M stands for mixed.)

QI_grant: Was qualified immunity granted to one or more of the defendants in this opinion? We considered an opinion vacating a denial of qualified immunity a qualified immunity grant.

- 1 – Qualified immunity was granted to at least one defendant in the opinion (even if other defendants had qualified immunity denied). This could have been an express grant, or it could have been a decision that reversed or vacated a qualified immunity denial and remanded the issue back to the district court.
- 0 – No grants of qualified immunity occurred in the opinion, regardless of reason (e.g., no ruling on qualified immunity, qualified immunity denied, lawsuit thrown out for procedural reasons).

QI_denied: Was qualified immunity denied for one or more defendants in this opinion? We did *not* consider denials of qualified immunity to municipalities to be denials, as municipalities can never be granted qualified immunity. We did consider an opinion vacating a grant of qualified immunity to be a qualified immunity denial.

- 1 – At least one defendant had qualified immunity denied in the opinion (even if other defendants had qualified immunity granted). This could have been an express denial, or it could have been a decision that reversed or vacated a qualified immunity grant and remanded the issue back to the district court.
- 0 – No denials of qualified immunity occurred in opinion, regardless of reason (e.g., no ruling on qualified immunity, qualified immunity granted for all defendants, lawsuit thrown out for procedural reasons).

LJFD: Did the court decline to rule on qualified immunity as it determined it lacked jurisdiction due to a factual dispute?

- 1 – The court found that for at least one claim, it lacked jurisdiction due to disputed facts (and therefore declined to rule on qualified immunity).
- 0 – There were no rulings that the court lacked jurisdiction due to disputed facts in this opinion.

APPENDIX E: FIRST AMENDMENT CODEBOOK

Separate from the majority of our coding, we hand coded supplemental information about a random sample of qualified immunity opinions with First Amendment claims. Unlike with the bulk of our coding, we did not use the resulting data for building or evaluating predictive algorithms. Instead, we directly analyzed the data and incorporated the results into our study beginning on p. 24.

First Amendment Violation Type (Categorical Response):

What type of First Amendment violation was alleged? (If multiple First Amendment violations were alleged in the appeal, we selected all of the categories below that were applicable.)

- **Retaliation:** The defendants were alleged to have engaged in premeditated retaliation against the plaintiffs in response to the plaintiffs' exercise of a protected First Amendment right. Critically, there must have been temporal separation between the protected activity and the retaliatory action, such that no "split-second" decision-making was involved.¹⁵⁰
 - **Subcategories:**
 - **Targeting of Private Citizen:** Government officials retaliated against a private citizen in response to protected First Amendment conduct.
 - **Employment Retaliation:** A government employee suffered adverse employment consequences in response to protected First Amendment activity.
 - **Prison Retaliation:** A prisoner or detainee was subjected to retaliatory action in response to protected First Amendment activity.
 - **Miscellaneous Retaliation:** Any First Amendment retaliation that did not fit into one of the subcategories above.
- **Direct Restriction on Speech/Association/Assembly:** The defendants were alleged to have acted to directly restrict the plaintiffs from exercising a First Amendment right (e.g., a plaintiff who was arrested for protesting, a plaintiff who was arrested for filming the police, a plaintiff who was arrested during a city council meeting). (Note: There were no subcategories for this category.)
- **Direct Restriction on Religious Liberty:** The defendants were alleged to have prohibited the plaintiffs from practicing their religion.
 - **Subcategories:**
 - **Prison:** The plaintiffs were prevented from freely practicing their religion in prison.
 - **Religious Discrimination:** Non-prison plaintiffs were discriminated against for religious reasons.
 - **Miscellaneous Religious Liberty:** Any direct restrictions on religious liberty that did not fit into one of the subcategories above.
- **Other:** Any First Amendment violations that did not neatly fit into one of the categories above.

Government Position of Defendant (Open Ended): What were the job titles/positions of the government officials sued?

We coded the following fields only if the appeal involved a retaliation claim.

Protected First Amendment Activity (Open Ended): What was the protected First Amendment activity that the plaintiffs alleged subjected them to retaliation? (If not a retaliation claim, we marked this as N/A.)

Retaliatory Action (Open Ended): What were the retaliatory acts allegedly undertaken by the defendants in response to the protected First Amendment activity? (If not a retaliation claim, we marked this as N/A.)

We coded the following field only if the appeal involved an employment retaliation claim.

Police Plaintiff (Y/N): Were any of the *plaintiffs* in the appeal police officers?

ENDNOTES

- ¹ Petition for Writ of Certiorari, *Gonzalez v. Trevino*, No. 22-1025 (U.S. Apr. 20, 2023) (No. 22-1025), available at <https://ij.org/wp-content/uploads/2023/04/Gonzalez-Petition-for-a-Writ-of-Certiorari.pdf>; Wilson, J. J. (2023, Oct. 13). *Supreme Court accepts First Amendment retaliation case* [Press release]. Arlington, VA: Institute for Justice. <https://ij.org/press-release/supreme-court-accepts-first-amendment-retaliation-case/>; Brief for Petitioner, *Gonzalez v. Trevino*, No. 22-1025 (U.S. Dec. 11, 2023), available at <https://ij.org/wp-content/uploads/2020/09/Brief-for-Petitioner.pdf>
- ² Petition for Writ of Certiorari, *Central Specialties, Inc. v. Large*, 143 S. Ct. 369 (2022) (No. 21-1552), available at <https://ij.org/wp-content/uploads/2022/06/42266-pdf-Purple-br.pdf>; Institute for Justice. (n.d., a). *Minnesota CSI* [Background]. Arlington, VA. <https://ij.org/case/minnesota-csi/>
- ³ Petition for Writ of Certiorari, *West v. Winfield*, 141 S. Ct. 111 (2020) (No. 19-899), available at <https://ij.org/wp-content/uploads/2020/01/39265-pdf-McNamara-br.pdf>
- ⁴ The officials in Allan's and Shaniz's cases succeeded. The district court in Allan's case granted the rogue county engineer qualified immunity, and the 8th Circuit upheld the grant after Allan appealed. *Central Specialties, Inc. v. Large*, 482 F. Supp. 3d 886 (D. Minn. 2020), *aff'd*, 18 F.4th 989 (8th Cir. 2021). In Shaniz's case, the district court denied the police officers qualified immunity, but the 9th Circuit reversed the decision on appeal. *West v. City of Caldwell*, 2018 WL 1526011 (D. Idaho Mar. 28, 2018), *rev'd*, 931 F.3d 978 (9th Cir. 2019). Both Allan and Shaniz petitioned the U.S. Supreme Court to hear their cases, but it refused. Kramer, J. (2022, Oct. 31). *Supreme Court refuses to hear case of government engineer who acted like a police officer, got qualified immunity anyway* [Press release]. Arlington, VA: Institute for Justice. <https://ij.org/press-release/supreme-court-refuses-to-hear-case-of-government-engineer-who-acted-like-a-police-officer-got-qualified-immunity-anyway/>; Kramer, J. (2020, June 15). *Supreme Court refuses to hear cases challenging qualified immunity* [Press release]. Arlington, VA: Institute for Justice. <https://ij.org/press-release/supreme-court-refuses-to-hear-cases-challenging-qualified-immunity/>. As this report goes to press, Sylvia still has a shot at justice. The mayor and police chief filed a motion to dismiss on qualified immunity grounds. The district court denied their motion, but they appealed—and the 5th Circuit granted their motion to dismiss on the ground that Sylvia failed to allege a constitutional violation. *Gonzalez v. City of Castle Hills*, 2021 WL 4046758 (W.D. Tex. Mar. 12, 2021), *rev'd in part sub nom. Gonzalez v. Trevino*, 42 F.4th 487 (5th Cir. 2022). Fortunately, the Supreme Court agreed to hear Sylvia's case, which it will do in early 2024. *Gonzalez v. Trevino*, 144 S. Ct. 325 (2023); Wilson, 2023.
- ⁵ Similar doctrines include de facto federal immunity, municipal immunity, and prosecutorial immunity. For more information, see Institute for Justice. (n.d., c). *Project on Immunity and Accountability*. Arlington, VA. <https://ij.org/issues/project-on-immunity-and-accountability>
- ⁶ *Harlow v. Fitzgerald*, 457 U.S. 800, 814 (1982).
- ⁷ *Frasier v. Evans*, 992 F.3d 1003 (10th Cir. 2021).
- ⁸ *Zadeh v. Robinson*, 928 F.3d 457 (5th Cir. 2019).
- ⁹ Qualified immunity also does not exist in criminal proceedings, which are generally supposed to provide greater protections to defendants than civil proceedings.
- ¹⁰ See, e.g., McNamara, R. (2020, Jan. 15). *Tear-gas grenades and 'qualified immunity.'* *The Wall Street Journal*. <https://www.wsj.com/articles/tear-gas-grenades-and-qualified-immunity-11579133525>; Jaicomo, P., & Bidwell, A. (2020, May 30). *Police act like laws don't apply to them because of 'qualified immunity.'* They're right. *USA Today*. <https://www.usatoday.com/story/opinion/2020/05/30/police-george-floyd-qualified-immunity-supreme-court-column/5283349002/>; Schwartz, J. (2023). *Shielded: How the police became untouchable*. New York, NY: Viking; Levinson, R. & Berens, M. (2020, June 4). *Special report: How union, Supreme Court shield*
- ¹¹ *Harlow v. Fitzgerald*, 457 U.S. 800, 814 (1982). For a history of qualified immunity and how government officials were held accountable before its creation, see Jaicomo, P., & Bidwell, A. (2022a). *Recalibrating qualified immunity: How *Tanzin v. Tanvir*, *Taylor v. Riojas*, and *McCoy v. Alamu* signal the Supreme Court's discomfort with the doctrine of qualified immunity.* *Journal of Criminal Law and Criminology*, 112(1), 105–144. And for a history of immunity for federal employees, see Jaicomo, P., & Bidwell, A. (2022b). *Unqualified immunity and the betrayal of *Butz v. Economou*: How the Supreme Court quietly granted federal officials absolute immunity for constitutional violations.* *Dickinson Law Review*, 126(3), 719–767.
- ¹² *Harlow v. Fitzgerald*, 457 U.S. 800, 818 (1982); *Ashcroft v. al-Kidd*, 563 U.S. 731, 741 (2011). Examples of cases with outrageous facts where government officials received qualified immunity due to a lack of clearly established law on point include *Frasier v. Evans*, 992 F.3d 1003 (10th Cir. 2021); *Zadeh v. Robinson*, 928 F.3d 457 (5th Cir. 2019); *Jessop v. City of Fresno*, 936 F.3d 937 (9th Cir. 2019); *Collie v. Barron*, 747 F. App'x 950 (5th Cir. 2018) (per curiam); *Cass v. City of Abilene*, 814 F.3d 721 (5th Cir. 2016); *Cope v. Cogdill*, 3 F.4th 198 (5th Cir. 2021); *Sampson v. County of Los Angeles ex rel. L.A. Cnty. Dep't of Children & Fam. Servs.*, 974 F.3d 1012 (9th Cir. 2020).
- ¹³ *Mitchell v. Forsyth*, 472 U.S. 511, 530 (1985) (“We hold that a district court's denial of a claim of qualified immunity, to the extent that it turns on an issue of law, is an appealable 'final decision' within the meaning of 28 U.S.C. § 1291 notwithstanding the absence of a final judgment.”).
- ¹⁴ One court has observed that if officials avail themselves of all opportunities to file intermediate appeals, they could get even more bites at the apple. *Joseph ex rel. Joseph v. Bartlett*, 981 F.3d 319, 330–331 & n.32 (5th Cir. 2020) (“An official can take multiple immediate appeals because the official can raise qualified immunity at any stage in the litigation—from Rule 12(b)(6) motions to dismiss, to Rule 12(c) motions for judgment on the pleadings, to Rule 56 motions for summary judgment, to Rule 50(b) post-verdict motions for judgment as a matter of law—and continue to raise it at each successive stage.”).
- ¹⁵ *Pearson v. Callahan*, 555 U.S. 223, 231 (2009).
- ¹⁶ *Pearson v. Callahan*, 555 U.S. 223, 231 (2009).
- ¹⁷ *Harlow v. Fitzgerald*, 457 U.S. 800, 814 (1982) (stating that the societal costs of civil rights claims “include the expenses of litigation, the diversion of official energy from pressing public issues, and the deterrence of able citizens from acceptance of public office. Finally, there is the danger that fear of being sued will 'dampen the ardor of all but the most resolute, or the most irresponsible [public officials], in the unflinching discharge of their duties.'”). It is important to note that when the Supreme Court created

- qualified immunity, it did so operating on numerous assumptions, including that qualified immunity would reduce litigation for government officials and protect them from personal bankruptcy while also putting them on notice of what conduct is illegal, all in the hopes of combatting the chilling effect that civil rights litigation is presumed to pose. This and other studies challenge those untested assumptions.
- ¹⁸ In addition to criticizing the practical impact of qualified immunity, critics argue that the doctrine lacks historical and legal support. They argue that the Supreme Court usurped the policymaking power of the legislative branch, creating a kind of immunity that did not exist in common law when Congress enacted 42 U.S.C. § 1983, giving people the right to sue state and local government officials who violate their constitutional rights. See, e.g., Jaicomo and Bidwell, 2022a; Baude, W. (2018). Is qualified immunity unlawful? *California Law Review*, 106, 45–90; Jeffries, J. C., Jr. (2010). What's wrong with qualified immunity? *Florida Law Review*, 62(4), 851–869; Blum, K. M. (2018). Qualified immunity: Time to change the message. *Notre Dame Law Review*, 93(5), 1887–1935; Schwartz, J. C. (2018). The case against qualified immunity. *Notre Dame Law Review*, 93(5), 1797–1852.
- ¹⁹ See Schwartz, 2018; Chen, A. K. (2018). The intractability of qualified immunity. *Notre Dame Law Review*, 93(5), 1937–1968.
- ²⁰ See *Hope v. Pelzer*, 536 U.S. 730 (2002).
- ²¹ See Jaicomo and Bidwell, 2022a. On the other hand, in *McCoy v. Alamu*, 141 S. Ct. 1364 (2021), the Supreme Court indicated an increased willingness to hold defendants accountable. There, it vacated a 5th Circuit decision granting qualified immunity to a prison guard who pepper-sprayed an inmate in the face for no reason. The 5th Circuit granted qualified immunity because it said there was no previous case establishing that precisely that conduct was unconstitutional—there were only cases establishing that *punching, tasing, or striking an inmate with a baton* for no reason was unconstitutional. *McCoy v. Alamu*, 950 F.3d 226 (5th Cir. 2020). The Supreme Court rejected the lower court's reasoning and remanded the case to be decided in light of the Court's then-recent decision in *Taylor v. Riojas*, 141 S. Ct. 52 (2020), where it ruled jailors had fair warning that their egregious acts were unconstitutional even in the absence of a prior case finding their specific conduct unconstitutional. *McCoy*, 141 S. Ct. at 1364. It remains to be seen whether the Court will continue to curb qualified immunity as it did in these cases.
- ²² See *Ashcroft v. al-Kidd*, 563 U.S. 731, 742 (2011). See also, e.g., *Kisela v. Hughes*, 138 S. Ct. 1148 (2018); *City & County of San Francisco v. Sheehan*, 575 U.S. 600 (2015); *Mullenix v. Luna*, 577 U.S. 7 (2015). For a discussion of the problems with requiring a high level of specificity, see Chen, 2018. This is not to say that federal appellate courts never deny qualified immunity absent direct precedent. In a recent Institute for Justice case, the 10th Circuit reversed the lower court's grant of qualified immunity. The case involved an incident of road rage in which an off-duty sheriff's deputy driving his personal vehicle, his child in tow, followed Mario Rosales to his home, blocked him in his driveway, shouted at him without identifying himself as law enforcement, and pointed a gun at him. Though the deputy was criminally convicted of assaulting Mario (and of endangering his own child), the district court ruled his behavior did not violate clearly established law because there was no prior case precisely on point. The 10th Circuit disagreed, finding the deputy's "conduct was obviously unconstitutional." *Rosales v. Bradshaw*, 72 F.4th 1145, 1156–1159 (10th Cir. 2023).
- ²³ See, e.g., Blum, K., Chemerinsky, E., & Schwartz, M. (2013). Qualified immunity developments: Not much *Hope* left for plaintiffs. *Touro Law Review*, 29(3), 633–658; Blum, 2018.
- ²⁴ See *Pearson v. Callahan*, 555 U.S. 223, 231 (2009). *Pearson* is the Supreme Court's most recent decision on how courts should adjudicate qualified immunity. Earlier, in *Stegert v. Gilley*, 500 U.S. 226 (1991), the Court suggested that courts should first rule on whether the government defendant violated the Constitution and then decide whether the right violated was clearly established. Because some courts continued to take the shortcut of deciding only whether a right was clearly established, in *Saucier v. Katz*, 533 U.S. 194 (2001), the Supreme Court made it mandatory for courts to rule first on the constitutional question. The Court then backtracked in *Pearson*, saying that the courts had discretion to rule on both questions or on only the clearly established one.
- ²⁵ See Rolfs, C. (2011). Qualified immunity after *Pearson v. Callahan*. *UCLA Law Review*, 59(2), 468–502 (finding that circuit courts more often avoided the constitutional question after *Pearson*, while district courts did not); Tokson, M. (2015). Judicial resistance and legal change. *The University of Chicago Law Review*, 82(2), 901–973 (finding that circuit courts more often avoided the constitutional question after *Pearson*); Nielson, A. L., & Walker, C. J. (2015). The new qualified immunity. *Southern California Law Review*, 89(1), 1–65 (finding that circuit courts more often avoided the constitutional question after *Pearson*). But see Sampsell-Jones, T., & Yauch, J. (2011). Measuring *Pearson* in the circuits. *Fordham Law Review*, 80(2), 623–641 (finding that two years after *Pearson*, circuit courts by and large had not changed their practice and continued to rule on the constitutional question).
- ²⁶ On the other hand, some research suggests that when courts are required to rule on the constitutional question, as they were between *Saucier* and *Pearson*, they are more likely to find that no violation occurred, which also does not help future plaintiffs. See Leong, N. (2009). The *Saucier* qualified immunity experiment: An empirical analysis. *Pepperdine Law Review*, 36(3), 667–714 (finding district and circuit courts were more likely to find no constitutional violation occurred following *Saucier*). See also Hughes, P. W. (2009). Not a failed experiment: *Wilson-Saucier* sequencing and the articulation of constitutional rights. *University of Colorado Law Review*, 80(2), 101–130 (finding that while circuit courts were more likely to rule on the constitutional violation following *Saucier*, their articulation was often restrictive of rights). But see Sobolski, G., & Steinberg, M. (2010). An empirical analysis of Section 1983 qualified immunity actions and implications of *Pearson v. Callahan*. *Stanford Law Review*, 62(2), 523–564 (finding no statistically significant increase in the rate at which circuit courts found no constitutional violation occurred following *Saucier*).
- ²⁷ See Pfander, J. E., Reinert, A. A., & Schwartz, J. C. (2020). The myth of personal liability: Who pays when *Bivens* claims succeed. *Stanford Law Review*, 72(3), 561–640; Schwartz, J. C. (2014). Police indemnification. *New York University Law Review*, 89(3), 885–1005. However, some observers argue policies indemnifying government officials are premised on the existence of qualified immunity. See Nielson, A. L., & Walker, C. J. (2020). Qualified immunity and federalism. *The Georgetown Law Journal*, 109(2), 229–303. This is a matter of some debate. See Schwartz, J. C. (2020b). Qualified immunity and federalism all the way down. *The Georgetown Law Journal*, 109(2), 305–351. Regardless, legislatures can choose to indemnify officials absent qualified immunity, as New Mexico did when it abolished its version of the doctrine for state-law claims in 2021. N.M. Stat. Ann. §§ 41-4A-1 *et seq.*
- ²⁸ See *Saucier v. Katz*, 533 U.S. 194, 202 (2001) ("If the law did not put the officer on notice that his conduct would be clearly unlawful, summary judgment based on qualified immunity is appropriate."). See also *Harlow v. Fitzgerald*, 457 U.S. 800, 819 (1982) ("Where an official could be expected to know that certain conduct would violate statutory or constitutional rights, he should be made to hesitate; and a person who suffers injury caused by such conduct may have a cause of action. But where an official's duties legitimately require action in which clearly established rights are not implicated, the public interest may be better served by action taken 'with independence and without fear of consequences.'").
- ²⁹ In some instances, unpublished federal appellate opinions may also factor into the clearly established question. See Cleveland, D. R. (2009). Clear as mud: How the uncertain precedential status of unpublished opinions muddles qualified immunity determinations. *University of Miami Law Review*, 65(1), 45–77.
- ³⁰ See Schwartz, J. C. (2021). Qualified immunity's boldest lie. *The University of Chicago Law Review*, 88(3), 605–684. In a similar vein, research has found that most police and sheriff's departments do not closely monitor civil rights cases against their officers or even have computerized systems for tracking such litigation, which would provide them with information about improvements needed in police training. See Schwartz, J. C. (2010). Myths and mechanics of deterrence: The role of lawsuits in law enforcement decisionmaking. *UCLA Law Review*, 57(4), 1023–1093. That said, research has also found that some departments do systematically track litigation and use it for training. See Schwartz, J. C. (2012). What police learn from lawsuits. *Cardozo Law Review*, 33(3), 841–894.
- ³¹ State appellate courts also decide federal qualified immunity appeals because they have general jurisdiction to hear any appeals that interest them. But because state courts defer to federal precedent, our study focuses on appeals decided by the federal appellate courts. Prior to our study, the largest empirical study on qualified immunity was one that looked at 4,054 federal appellate cases. See Reinert, A. A. (2021). *Qualified immunity on appeal: An empirical assessment* (Faculty Research Paper No. 634). New York, NY: Jacob Burns Institute for Advanced Legal Studies, Cardozo Law School, Yeshiva University.

- Other studies of federal cases analyzed far fewer cases, ranging from 29 cases to 979, but none more than 1,000. See, e.g., Braaten, C. N., & Vaughn, M. S. (2023). Police officers' qualified immunity in excessive force claims, warrant execution, and warrantless searches and arrests: Tracing the evolution of, and stagnation in, U.S. Supreme Court precedents. *American Journal of Criminal Justice*, 48, 65–95. <https://doi.org/10.1007/s12103-021-09617-w> (29 Supreme Court cases); Hughes, 2009 (414 circuit court cases); Lammon, B. (2022). Reforming qualified-immunity appeals. *Missouri Law Review*, 87(4), 1137–1218 (773 circuit court cases); Nielson and Walker, 2015 (844 circuit court cases); Botts, J., Dowdell, J., & Januta, A. (2020, Dec. 23). Taking the measure of qualified immunity: How Reuters analyzed the data. *Reuters*. <https://www.reuters.com/investigates/special-report/usa-police-immunity-methodology/> (529 circuit court cases and 435 district court cases); Schwartz, J. C. (2017). How qualified immunity fails. *The Yale Law Journal*, 127(1), 2–76 (979 district court cases).
- ³² Examples of studies that limit their scope to claims against law enforcement defendants or those alleging excessive force or related Fourth Amendment violations include Schwartz, 2017 (study of cases involving claims against state and local police brought in five federal districts over two years, 2011–2012); Botts et al., 2020 (study of cases involving police who raised the qualified immunity defense in circuit courts from 2005 to 2019 and in federal district courts in two states between 2014 and 2018); Braaten and Vaughn, 2023 (study of all cases heard by the Supreme Court up to June 6, 2020, involving police who raised the qualified immunity defense). Other studies, however, look beyond excessive force and related violations. Among those studies, we are aware of only two that report on multiple alleged constitutional violations, though not as many as our study codes for. See Reinert, 2021; Reinert, A. A. (2018). Qualified immunity at trial. *Notre Dame Law Review*, 93(5), 2065–2092.
- ³³ Through its recent docket and rulings, the Supreme Court has repeatedly emphasized that courts should give substantial deference to law enforcement officials who are forced to make split-second decisions in difficult situations. See, e.g., *Rivas-Villegas v. Cortesluna*, 595 U.S. 1 (2021); *City of Tablequah v. Bond*, 595 U.S. 9 (2021); *City of Escondido v. Emmons*, 139 S. Ct. 500 (2019); *Kisela v. Hughes*, 138 S. Ct. 1148 (2018); *White v. Pauly*, 580 U.S. 73 (2017); *Mullenix v. Luna*, 577 U.S. 7 (2015); *City & County of San Francisco v. Sheehan*, 575 U.S. 600 (2015); *Plumbhoff v. Rickard*, 572 U.S. 765 (2014).
- ³⁴ The exact search we used was “qualified immunity” & “DA aft 12-31-2009 & bef 01-01-2021.” We also tested slightly different search terms (e.g., using the root “qualif” as opposed to the word “qualified,” allowing one word between “qualified” and “immunity”). While these searches returned a handful of additional opinions, our review of a sample of those opinions revealed that they did not involve qualified immunity being raised on appeal.
- ³⁵ The date the lawsuit was filed in district court, or the case origination date, was never present in the opinion documents, while the district court case number was not always present. However, both the case origination date and the district court case number did generally appear in the “dockets” section of the “filing” tab on Westlaw’s user interface. Because these two fields did not require any discretion, we had a third-party contractor code the information for all 7,173 opinions. When the information was not present in the filings tab, we instructed the contractor to code the field as “NP” (not present). The contractor coded those fields as NP for about 20% of opinions, meaning the information was present in the filings tab for about 80%. Where the district court case number was not present in the filings tab, we used an algorithm to search the opinion documents for it, thereby identifying the district court case number for another 10% of opinions. The algorithm had an overall accuracy of 95.7%. Finally, we employed several quality-control measures to ensure the hand coding for both fields was accurate. In a final review of every 30th entry, we found only one mistake, an error rate of 0.5%.
- ³⁶ In addition, we suspected our algorithms would struggle to consistently and accurately attribute specific pieces of information to individual claims.
- ³⁷ See, e.g., Hughes, 2009; Sampsell-Jones and Yauch, 2011; Reinert, A. A. (2010). Measuring the success of *Bivens* litigation and its consequences for the individual liability model. *Stanford Law Review* 62(3), 809–862; Reinert, 2021; Schwartz, 2017; Botts et al., 2020.
- ³⁸ To ensure each year was proportionally represented among our hand-coded opinions, we used stratified sampling. For example, for 2010, Westlaw returned 611 relevant results. We therefore randomly selected 11% of these 611 opinions (67) for our hand coding. We also used stratified sampling when assigning opinions to the training and testing samples, so all years were proportionally represented in both samples.
- ³⁹ In addition, to ensure our final evaluation of the algorithms’ performance was as accurate as possible, we had only attorneys code the testing sample.
- ⁴⁰ As part of this process, we assigned roughly 20% of opinions to two different coders and strongly encouraged them to code a field with a question mark (“?”) if they were unsure of the correct code. All discrepancies between coders in the overlapping 20% of opinions, and all fields coded with a question mark, went through additional review. Through this review, we corrected clear errors and referred tougher cases to the panel of attorneys with experience in qualified immunity. They made the final coding choices and, if needed, proposed corrections or clarifications to the codebook. The overlapping 20% of opinions also allowed us to ensure all coders understood our codebook definitions and were coding the opinions accurately.
- ⁴¹ Critically, we did not use the testing sample at any point during the process of building the algorithms. This ensured that the testing sample was completely novel to the algorithms (and their developer), which meant that the algorithms’ performance on the sample was representative of their overall performance. For the testing sample, there were only 209 instances (out of roughly 5,500 individual codes) in which the initial hand coding and the algorithm’s prediction differed. For these discrepancies, we conducted a two-step review process. First, we screened out obvious algorithm errors. Second, a senior attorney with experience in qualified immunity conducted a fully blind review of the remaining discrepancies to determine the correct code. Through this process, we determined that 153 of the discrepancies were algorithm errors, while 56 were errors—generally administrative errors—in the original hand coding. We, of course, counted the 153 algorithm errors against the algorithms when determining the final performance statistics listed in Table 2 and Appendix B. However, we corrected those errors in our final dataset. As for the 56 hand-coding errors, the final performance statistics reflect the fact that the algorithms got those items right.
- ⁴² By comparable legal studies, we mean legal studies that used a hand-coded sample to build algorithms, and then used those algorithms to perform classification on an entire collection of documents. Still, performance is highly dependent on the nature of the classification task and data, so these are at best approximate comparators.
- ⁴³ No single statistic fully captures an algorithm’s effectiveness at predicting a field. Several factors, including how a field is used, whether a field is rare, and how many response categories are possible, dictate which performance statistics are most appropriate and how those statistics are interpreted.
- ⁴⁴ See Fagan, F. (2015). From policy confusion to doctrinal clarity: Successor liability from the perspective of big data. *Virginia Law & Business Review*, 9(3), 391–451; Macey, J. & Mitts, J. (2014). Finding order in the morass: The three real justifications for piercing the corporate veil. *Cornell Law Review*, 100(1), 99–155; Rauterberg, G., & Talley, E. (2017). Contracting out of the fiduciary duty of loyalty: An empirical analysis of corporate opportunity waivers. *Columbia Law Review*, 117(5), 1075–1151. These studies serve as only rough comparators to ours because of the divergent nature of their datasets and goals.
- ⁴⁵ This is an oversimplification. Our goals for each field varied based on our intended usage of that field and the distribution of data. (All else being equal, fields with categories that appear less frequently tend to have higher accuracy because it is easier to randomly guess the right answer.) Still, 95% serves as a reasonable approximation of our goals.
- ⁴⁶ The best-performing model from one study (a “how to” for making automated predictions on legal texts) reports 0.4484 precision at 0.80 recall, which translates to an F1 score of 0.57. See Table 4 in Keeling, R., Chhatwal, R., Huber-Fliflet, N., Zhang, J., Zhao, H. (2020). Using machine learning on legal matters: Paying attention to the data behind the curtain. *Hastings Science and Technology Law Journal*, 11(1), 9–36. Another study reported precision of 0.89 and recall of 0.94, which translates to an F1 score of 0.91. See Nyarko, J. (2019). We’ll see you in . . . court! The lack of arbitration clauses in international commercial contracts. *International Review of Law and Economics*, 58, 6–24. These studies serve as only rough comparators to ours because of the divergent nature of their datasets and goals.
- ⁴⁷ With an F1 score of 0.71, the illegal search field was the only field whose performance fell substantially below our goals. We therefore avoided detailed analysis of this field. Fields with minimal or no data included en banc, task force, government level, federal law enforcement, federal prison, parental rights, and religious liberty. We report a few general statistics for these fields (e.g., their frequencies in our dataset) but generally avoided additional analysis as we cannot be confident that the algorithms can accurately predict fields

- with minimal data.
- ⁴⁸ This includes the relevance prediction for all 7,173 opinions and, for those opinions predicted as relevant, predictions for the other 33 fields. It does not include the date the lawsuit was filed in district court or the district court case number field.
- ⁴⁹ If our algorithms predicted that an appeal involved qualified immunity, it almost always did. In the few instances where the algorithms misclassified appeals, they usually predicted that qualified immunity was *not* raised in an appeal when, in fact, it was. (See Appendix B for further details about the algorithms' performance.) In most such instances, qualified immunity was not a central part of the court's opinion: In our testing sample, the relevant opinions our algorithms missed used the phrase "qualified immunity" an average of less than one time each (not including footnotes). Meanwhile, the opinions our algorithms correctly classified as relevant used the phrase roughly nine times each.
- ⁵⁰ The two exceptions are the District of Columbia Circuit, which had few relevant appeals (41), and the Federal Circuit, which generally does not hear qualified immunity cases.
- ⁵¹ Our sources for the number of circuit appeals overall were the Cases Filed, Terminated, and Pending, by Nature of Proceeding tables (Table B-1) in the Administrative Office of the U.S. Courts' annual *Judicial Business of the United States Courts* reports: <https://www.uscourts.gov/statistics-reports/analysis-reports/judicial-business-united-states-courts>. To facilitate as close to an apples-to-apples comparison with our data as possible, we used the number of appeals terminated each year rather than the number initiated. We also averaged the number of appeals from 2010 to 2015 and compared it to the average from 2016 to 2020.
- ⁵² The average number of excessive force appeals was 112 during the first half of the study period and 168 during the second.
- ⁵³ The average number of appeals *not* involving excessive force increased from 346 during the first half of the study period to 387 during the second.
- ⁵⁴ We generally categorized any state or local official with the power to make arrests or obtain warrants as a state or local law enforcement officer, even if they were not employed by the police. See our codebook in Appendix D for further details.
- ⁵⁵ Among law enforcement defendants, state and local officials are far more common, making up 96% of defendants in qualified immunity appeals. Federal law enforcement officials were defendants in just 3% of appeals, while the other 1% had both state or local and federal officials as defendants. Even after accounting for the relative populations of sworn state and local law enforcement officers versus federal ones, state and local law enforcement officials are four times more likely than federal ones to be defendants in a qualified immunity appeal. Our sources for the populations of law enforcement officers were Gardner, A. M., & Scott, K. M. (2022). *Census of state and local law enforcement agencies, 2018 – Statistical tables*. U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics. <https://bjs.ojp.gov/sites/g/files/xyckuh236/files/media/document/cslea18st.pdf>, and Brooks, C. (2019). *Federal law enforcement officers, 2016 – Statistical tables*. U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics. <https://bjs.ojp.gov/content/pub/pdf/fleo16st.pdf>
- ⁵⁶ As with law enforcement defendants, prison defendants were typically state or local, as opposed to federal, officials. Among appeals with prison defendants, 94% solely involved state or local officials, while 6% solely involved federal officials. None of the appeals we studied involved both state or local and federal prison defendants.
- ⁵⁷ False arrest violations include false arrests as well as malicious prosecutions and other non-arrest seizures of a person. See our codebook in Appendix D for further details.
- ⁵⁸ Procedural due process violations generally relate to violations of the right to fair treatment in some procedure. For example, students expelled from a public university without clear notice of their alleged transgressions might claim a procedural due process violation against university officials. These types of violations are more commonly alleged against administrative officials than law enforcement officers.
- ⁵⁹ Opinions sometimes involved violation types we did not code, such as substantive due process. And a small number of opinions included *only* violation types we did not code for.
- ⁶⁰ Because the date the lawsuit was filed in district court was missing for roughly 20% of opinions, we were unable to calculate the duration of those lawsuits at the time of the appeal decision. Nevertheless, our estimate of three years and two months likely understates the duration of qualified immunity litigation. The longest lawsuits—by far—were those where appeals occurred post-trial, and they were overrepresented among the opinions missing the lawsuit filing date.
- ⁶¹ On average from 2010 through 2020, the median length of a civil lawsuit was two years and seven months. We calculated this figure using data on "other civil appeals" from the Median Times for Civil and Criminal Cases Terminated on the Merits tables (Table B-4A) in the Administrative Office of the U.S. Courts' annual *Judicial Business of the United States Courts* reports: <https://www.uscourts.gov/statistics-reports/analysis-reports/judicial-business-united-states-courts>. To facilitate as close to an apples-to-apples comparison with our data as possible, we used the "filing in lower court to . . . final order in appeals court" column. (Note that this is called the "from filing in lower court to final disposition in appellate court" column in reports for 2010–2012.) We averaged the median from each individual year within our reporting period to arrive at the average annual median. Importantly, the AOUSC data do not include "prisoner petitions," while our dataset includes at least a subset of them (prisoner civil rights appeals). These prisoner petitions tend to take substantially less time than other civil appeals (an average of four months less, depending on the year), meaning our data likely understate the length of qualified immunity appeals relative to all civil appeals.
- ⁶² The circuits vary substantially in how they treat unpublished opinions, which speaks to the muddled and confusing implementation of qualified immunity across the country. As of 2009, however, only the 9th Circuit definitively allowed for the use of unpublished opinions in creating clearly established law. For further details about how circuits treat unpublished opinions when it comes to qualified immunity, see Cleveland, 2009.
- ⁶³ Our sources for the publication rate in civil appeals overall were the Type of Opinion or Order Filed in Cases Terminated on the Merits tables (Table S-3 in reports for 2010–2013 and Table B-12 in reports for 2014–2020) in the Administrative Office of the U.S. Courts' annual *Judicial Business of the United States Courts* reports: <https://www.uscourts.gov/statistics-reports/analysis-reports/judicial-business-united-states-courts>. Specifically, we tabulated the number of published appeals and the total number of non-consolidated appeals. We then calculated the publication rate by dividing the number of published appeals by the total number of non-consolidated appeals. (For 2010 and 2011, we divided by the total number of appeals as the total number of non-consolidated appeals was not available.)
- ⁶⁴ Though we did not include the District of Columbia Circuit in our main analyses, it also had a high publication rate (78%).
- ⁶⁵ The 5th Circuit granted qualified immunity most often, in 67% of appeals. In contrast, the 4th Circuit granted qualified immunity least often, in 46% of appeals.
- ⁶⁶ As shown on the left-hand side of Figure 13, plaintiffs were appealing a district court loss in 61% of appeals, while defendants were appealing a district court loss in 36% of appeals. The remaining 3% were cross-appeals, where both plaintiffs and defendants were appealing aspects of the district court's decision. These data represent district court outcomes when the losing party opted to appeal; we do not have data regarding all district court outcomes.
- ⁶⁷ For this analysis, we randomly selected 125 First Amendment appeals from our training and testing samples. The 125 appeals represented 12.6% of all First Amendment appeals in our dataset. Our First Amendment codebook, available in Appendix E, provides complete definitions and response categories for the fields we coded. Though we used a different codebook, we generally coded these fields in the same manner as the main fields in this study. This included assigning 20% of opinions to two different coders to ensure accuracy.
- ⁶⁸ The margin of error (at a 95% confidence interval) for our First Amendment sample is roughly 8 percentage points. This means among all First Amendment claims, we estimate the percentage alleging premeditated retaliation to be between 51% and 67%—our estimate of 59% plus or minus 8 percentage points. Analyses within subgroups (e.g., analysis of premeditated retaliation claims specifically) will have higher margins of error. As such, these results should be interpreted as estimates and not precise figures.
- ⁶⁹ Often, claims that we considered direct restrictions of speech or religious liberty were framed by the plaintiff as retaliation. However, we only counted claims as "premeditated retaliation" if there was separation in time between the protected activity and the alleged retaliation. For example, a plaintiff might be arrested during a protest and file a First Amendment retaliation claim, alleging the government official arrested them in retaliation for their protest. But unless the plaintiff alleged that the arrest was motivated by speech or conduct that

occurred *prior* to the protest, we considered this a direct restriction on speech and not premeditated retaliation. This conservative approach ensured that our “premeditated retaliation” category did not involve any split-second decisions by government officials.

⁷⁰ *Rosenblatt v. Baer*, 383 U.S. 75, 85 (1966). See also *N.Y. Times Co. v. Sullivan*, 376 U.S. 254 (1964). Among the “direct targeting” claims that did not involve criticism of a government official, two alleged targeting for the plaintiff’s speech at a public university (expressing disfavored opinions in class, privately reporting professor misconduct), two alleged targeting for lawsuits the plaintiff had filed, one alleged targeting for the plaintiff’s testimony in court, and one alleged targeting for the plaintiff’s running for political office.

⁷¹ *Hoggard v. Rhodes*, 141 S. Ct. 2421, 2422 (2021) (Thomas, J., respecting denial of certiorari).

⁷² This finding holds when looking at all appeals regardless of the party appealing. Defendants fully prevailed in only 48% of all appeals with excessive force allegations, compared to 59% of all appeals with First Amendment violations.

⁷³ We are unable to attribute specific outcomes (i.e., a grant or denial of qualified immunity) to specific claims. As a workaround, we used the prevailing party field to measure the success rate on First Amendment and excessive force claims since it indicates the overall outcome of the appeal. If an appeal involved a First Amendment claim and government defendants prevailed overall, we can be reasonably confident that they prevailed on the First Amendment claim, too. Furthermore, we highlight defendant appeals because nearly all defendant appeals are interlocutory appeals. Thus, even though we cannot attribute the outcomes on specific claims to qualified immunity, qualified immunity likely plays a central role in such appeals. Finally, we cannot be sure which party won the relevant claims for the 18% of excessive force appeals or the 19% of First Amendment appeals with “mixed judgment” outcomes, where defendants and plaintiffs each prevailed in part. While unlikely that these mixed judgments would close the 11-percentage-point gap in the success rate, we cannot entirely rule it out.

⁷⁴ Petition for Writ of Certiorari, *Novak v. City of Parma*, 143 S. Ct. 773 (2023) (No. 22-293), available at <https://ij.org/wp-content/uploads/2022/09/Petition-for-Writ-of-Certiorari-Novak-v.-Parma-FINAL.pdf>; Harper, J. (2016, Aug. 12). Jury acquits Parma man who made fake Cleveland page mocking police department. *Cleveland.com*. <https://www.cleveland.com/court-justice/2016/08/jury-acquits-parma-man-who-mad.html>; Institute for Justice. (n.d., b). *Novak v. Parma* [Background]. Arlington, VA. <https://ij.org/case/novak-v-parma/>

⁷⁵ *Novak v. City of Parma*, 932 F.3d 421 (6th Cir. 2019); *Novak v. City of Parma*, 33 F.4th 296 (6th Cir. 2022). Unfortunately, the Supreme Court declined to hear the case in February 2023. Wimer, A. (2023, Feb. 21). *Supreme Court rejects appeal of man arrested for Facebook police parody* [Press release]. Arlington, VA: Institute for Justice. <https://ij.org/press-release/supreme-court-rejects-appeal-of-man-arrested-for-facebook-police-parody/>

⁷⁶ *Novak v. City of Parma*, 33 F.4th 296, 305 (6th Cir. 2022).

⁷⁷ Other studies finding that the federal circuits more often rule in favor of government officials include Nielson and Walker, 2015 (finding that plaintiffs prevailed only 27.7% of the time); Reinert, 2021 (finding that qualified immunity was granted in full in about 60% of opinions, denied in full in only 30%, and granted in part and denied in part in another 7%); Lammon, 2022 (finding that among interlocutory appeals, and when counting only those appeals where the lower court was reversed or affirmed in full, defendants prevailed on 34% of appeals and plaintiffs prevailed 31% of the time). In addition, studies have found that in cases appealed to the Supreme Court, police defendants succeeded more often than plaintiffs (Braaten and Vaughn, 2023) and that trial court juries also appear to favor government defendants (Reinert, 2018). On the other hand, one study of appeals across five circuits found plaintiffs were more successful than defendants when defendants filed interlocutory appeals of qualified immunity, with appeals being affirmed three times as often as they were reversed. (Looking at appeals of qualified immunity grants, however, the study found plaintiffs prevailed just 7.7% of the time compared to 65.4% for defendants.) The same study also found that district courts more often ruled in favor of plaintiffs, denying qualified immunity 31.6% of the time and granting it 12% of the time. (Another 5.9% of rulings were mixed.) Schwartz, 2017.

⁷⁸ This includes all post-trial judgments where the plaintiff was the sole prevailing party. It is likely overbroad, as some plaintiff victories post-trial may result in retrials instead of final judgments. It also does not include rare instances

of final judgments for plaintiffs at the summary judgment stage of litigation. (For technical reasons, we are unable to pinpoint the number of final plaintiff judgments at the stage of summary judgment, but they are clearly rare: In our hand-coded data, we did not have a single such opinion.) Finally, this only reflects the number of post-trial judgments for plaintiffs in our data. Some cases will settle before trial, so plaintiffs may “win” those. Further, it could be that defendants raise qualified immunity less frequently at trial, so our dataset is not capturing some final victories for plaintiffs. (If the opinion did not raise qualified immunity on appeal, it did not end up in our dataset.) Even with these caveats, it is notable that so few final judgments are in favor of plaintiffs.

⁷⁹ In this context, “statistically significant” means that we can attribute the differences in qualified immunity grants and denials across circuits to real differences between the circuits and not random variation. Using a chi-squared test, we found that differences in both grants and denials of qualified immunity across circuits were statistically significant, with p-values well under 0.001. It is worth noting that our results remain statistically significant after controlling for variation in the rate of self-represented plaintiffs, which means more plaintiff losses in a circuit cannot be chalked up to more plaintiffs proceeding without counsel.

⁸⁰ Schwartz, J. C. (2023). Civil rights without representation. *William & Mary Law Review*, 64(3), 641–706; Botts et al., 2020; Reinert, 2021.

⁸¹ Setting aside the District of Columbia Circuit, the 9th Circuit represented the largest population as of 2015 (~65 million), while the 1st Circuit represented the smallest (~14 million). We calculated circuit populations using U.S. Census Bureau. (2023). *County population totals: 2010–2019*. <https://www.census.gov/data/datasets/time-series/demo/popest/2010s-counties-total.html>

⁸² Miller, M., Cairns, M., Bidwell, A., Jaicomo, P., & Morton, K. (2022). *Constitutional GPA: Is your government preventing accountability?* Arlington, VA: Institute for Justice. <https://ij.org/report/constitutional-gpa/>

⁸³ Miller et al., 2022. The 7th and 8th Circuits also have more statements of clearly established law than the low-publishing 5th and 11th Circuits (177 and 162) despite representing considerably smaller populations. On the other hand, the 1st Circuit’s high publication rate is somewhat offset by its small population (~14 million), while the 9th Circuit’s low publication rate is somewhat offset by its large population (~65 million). We calculated circuit populations using U.S. Census Bureau. (2023). *County population totals: 2010–2019*. <https://www.census.gov/data/datasets/time-series/demo/popest/2010s-counties-total.html>

⁸⁴ In addition, we found an interesting inverse relationship between plaintiffs’ likelihood of success when represented by counsel in a circuit and the rate at which plaintiffs represent themselves in that circuit. Put simply, the harder it was for plaintiffs *with an attorney* to win, the more often plaintiffs acted as their own attorney. To be sure, a number of additional variables—including the types of cases considered in each circuit and other regional demographic factors—may influence the ability of plaintiffs to retain counsel. And indeed, preliminary exploration suggests that the percentage of prison cases in a circuit explains some, although not all, of the correlation. We therefore avoid drawing any conclusions in this primarily descriptive study. Still, given others’ research on the subject, this warrants further investigation. For additional research on the subject, see, e.g., Schwartz, J. C. (2020a). Civil rights ecosystems. *Michigan Law Review*, 118(8), 1539–1601, and Schwartz, 2023.

⁸⁵ We originally coded three fields related to the rationale for granting qualified immunity: two related to whether a constitutional violation did or did not occur and one that detailed the court’s ruling (if applicable) on the clearly established question. We hand coded these fields for our primary training sample (just over 400 relevant opinions) but not for our validation or testing samples. We abandoned these fields relatively early in the project as the algorithms were unable to predict them with satisfactory accuracy.

⁸⁶ During the coding process, we instructed our coders to use a question mark if they were unsure of the correct code for a particular field. For roughly 1 in 4 opinions with qualified immunity grants, coders indicated that they were unsure of the correct code for at least one of the three fields relating to the court’s rationale for granting qualified immunity. This far exceeded the uncertainty rate for any other fields.

⁸⁷ Hughes, 2009.

⁸⁸ Sampsell-Jones and Yauch, 2011.

⁸⁹ Nielson and Walker, 2015.

⁹⁰ *Young v. Borders*, 850 F.3d 1274 (11th Cir. 2017).

⁹¹ *Young v. Borders*, 620 F. App’x 889 (11th Cir. 2015) (per curiam).

⁹² *Young v. Borders*, 850 F.3d 1274 (11th Cir. 2017) (mem.).

- ⁹³ This translates to just under 200 interlocutory appeals per year and is consistent with a 2022 study that found federal appellate courts decided 773 interlocutory appeals over a four-year period from 2017 through 2020—or just under 200 per year. See Lammon, 2022.
- ⁹⁴ Interlocutory appeals of qualified immunity rose from an average of 165 during the first part of the study period (2010–2015) to 190 during the second half (2016–2020)—an increase of 15%.
- ⁹⁵ Schwartz, J. C. (2020c). Qualified immunity's selection effects. *Northwestern University Law Review*, 114(5), 1101–1178.
- ⁹⁶ Complaint, *Mathis v. County of Lyon*, No. 2:07-CV-00628 (D. Nev. May 14, 2007), 2007 WL 4637091.
- ⁹⁷ Motion for Partial Judgment on the Pleadings, *Mathis v. County of Lyon*, No. 2:07-CV-00628 (D. Nev. Jan. 2, 2008), 2008 WL 7279623.
- ⁹⁸ *Mathis v. County of Lyon*, 633 F.3d 877 (9th Cir. 2011). The 9th Circuit affirmed the denial of qualified immunity in a 2-1 decision. Since this decision occurred at the pleadings stage, the court was required to (for purposes of the appeal) accept the plaintiffs' allegations as true. That one member of the court believed outright theft failed to violate clearly established law shows how difficult it can be for plaintiffs to overcome qualified immunity even when violations are obvious.
- ⁹⁹ *Mathis v. County of Lyon*, No. 2:07-CV-00628, 2013 WL 3725036 (D. Nev. July 12, 2013).
- ¹⁰⁰ *Mathis v. County of Lyon*, 591 F. App'x 635 (9th Cir. 2015).
- ¹⁰¹ Perea, R. (2015, Nov. 16). Jury awards \$2.1m to heirs for charges against Glover. *Reno Gazette Journal*. <https://www.rgj.com/story/news/local/mason-valley/2015/11/16/jury-awards-21m-heirs-charges-against-glover/75895178/>
- ¹⁰² *Mathis v. County of Lyon*, 757 F. App'x 542 (9th Cir. 2018). The defendant appealed numerous additional aspects of the lawsuit in this final appeal, but we focus solely on the qualified immunity component. Nonetheless, the jury's verdict was generally upheld.
- ¹⁰³ Satisfaction of Judgment, *Mathis v. County of Lyon*, No. 2:07-CV-00628 (D. Nev. May 10, 2019), ECF No. 518. For this summary, we also relied on the timeline and case overview from the docket sheet for the case provided by Westlaw.
- ¹⁰⁴ Schwartz, 2020c.
- ¹⁰⁵ *Harlow v. Fitzgerald*, 457 U.S. 800, 819 (1982).
- ¹⁰⁶ Wilson, C. R. (2000). "Location, location, location": Recent developments in the qualified immunity defense. *N.Y.U. Annual Survey of American Law*, 57(4), 445–475, p. 447.
- ¹⁰⁷ Circuits have nearly 200 statements of clearly established law on average. See Miller et al., 2022.
- ¹⁰⁸ *Favela v. City of Las Cruces ex rel. Las Cruces Police Dep't*, 398 F. Supp. 3d 858, 894 n.11 (D.N.M. 2019).
- ¹⁰⁹ Schwartz, 2021. Instead, police are taught to apply the general principles described in a few foundational cases to a wide variety of circumstances and situations.
- ¹¹⁰ See Schwartz, 2017 (showing that 75.7% of summary judgment motions for applicable civil rights claims included an argument based on qualified immunity, compared to only 33.3% of motions to dismiss).
- ¹¹¹ Although our data cannot speak directly to this point, prior research suggests that at least some discovery occurs prior to summary judgment for nearly all lawsuits involving qualified immunity. Schwartz, 2017.
- ¹¹² In addition, government officials may still be forced to participate in litigation even if the personal liability claims against them are dismissed on qualified immunity grounds. This is because the lawsuit may include other related claims for which qualified immunity is inapplicable, such as claims against the government entity or municipality. These lawsuits would continue and the government officials, as parties with relevant information, might still be forced to provide depositions or testify. Our hand coders observed that personal liability claims against government defendants were often accompanied by claims against the relevant government entity ("official capacity" claims), claims against the municipality ("Monell" claims), or state law claims.
- ¹¹³ See *Crawford-El v. Britton*, 523 U.S. 574, 597–600 (1998). This Supreme Court decision lists a myriad of tools for protecting government officials from spurious lawsuits.
- ¹¹⁴ U.S. Const. amend. IV.
- ¹¹⁵ *Graham v. Connor*, 490 U.S. 386, 396–397 (1989).
- ¹¹⁶ Our anecdotal observation from coding a representative sample of qualified immunity appeals is that lawsuits with meritless or trivial claims seem relatively rare, at least among appeals. Of course, defining a meritless or trivial claim is a highly subjective exercise, so we do not report any concrete data on this point.
- ¹¹⁷ See Schwartz, 2017. For motions to dismiss or motions on the pleadings in which qualified immunity was raised, the motion was granted in full or in part on qualified immunity grounds 13.6% of the time. Meanwhile, the motion was granted in full or in part for other reasons 33.8% of the time. See Table 7 in Schwartz, 2017. Motions for summary judgment showed similar results: 20.5% of motions in which qualified immunity was raised were granted in full or in part based on qualified immunity, while 30.0% were granted in full or in part on other grounds. See Table 8 in Schwartz, 2017. (Motions granted on other grounds included motions where qualified immunity was granted in the alternative as well as motions where the court ruled that no constitutional violation occurred.)
- ¹¹⁸ *Firestone Tire & Rubber Co. v. Risjord*, 449 U.S. 368, 374 (1981). Justice William Brennan provided several reasons why interlocutory appeals should be rare and trial courts should generally be permitted to see cases through to final judgment "without a court of appeals peering over [their] shoulder[s] every step of the way." *Mitchell v. Forsyth*, 472 U.S. 511, 544 (1985) (Brennan, J., concurring in part and dissenting in part). This final judgment rule "preserves scarce judicial resources that would otherwise be spent in costly and time-consuming appeals. Trial court errors become moot if the aggrieved party nonetheless obtains a final judgment in his favor, and appellate courts need not waste time familiarizing themselves anew with a case each time a partial appeal is taken. Equally important, the final judgment rule removes a potent weapon of harassment and abuse from the hands of litigants. . . . In many cases in which a claim of right to immediate appeal is asserted, there is a sympathetic appellant who would undoubtedly gain from an immediate review of his individual claim. But lurking behind such cases is usually a vastly larger number of cases in which relaxation of the final judgment rule would threaten all of the salutary [sic] purposes served by the rule." *Mitchell*, 472 U.S. 511 at 543–544.
- ¹¹⁹ (1,939 interlocutory appeals * 441 days) / 365 days = 2,343 years. This is *not* the net effect of interlocutory appeals on qualified immunity lawsuit length. Where a government defendant wins an interlocutory appeal, that appeal may shorten that lawsuit. However, defendants fully prevailed in only a third of interlocutory appeals, suggesting that these appeals drew out more lawsuits than they cut short.
- ¹²⁰ *Zadeh v. Robinson*, 928 F.3d 457, 479–480 (5th Cir. 2020) (Willett, J., concurring).
- ¹²¹ *Jessop v. City of Fresno*, 936 F.3d 937 (9th Cir. 2019).
- ¹²² *Collie v. Barron*, 747 F. App'x 950, 950 (5th Cir. 2018); Chung, A., Hurley, L., Januta, A., Botts, J., & Dowdell, J. (2020, Aug. 25). Shot by cops, thwarted by judges and geography. *Reuters*. <https://www.reuters.com/investigates/special-report/usa-police-immunity-variations>
- ¹²³ *Cass v. City of Abilene*, 814 F.3d 721 (5th Cir. 2016).
- ¹²⁴ *Cope v. Cogdill*, 3 F.4th 198 (5th Cir. 2021).
- ¹²⁵ *Sampson v. County of Los Angeles ex rel. L.A. Cnty. Dep't of Children & Fam. Servs.*, 974 F.3d 1012, 1025 (9th Cir. 2020).
- ¹²⁶ The Institute for Justice has model legislation states and municipalities can use to rein in qualified immunity in their jurisdictions. See Institute for Justice. (n.d., d). *Protecting Everyone's Constitutional Rights Act – (Q.I. reform)* [Model legislation]. Arlington, VA. <https://ij.org/legislation/protecting-everyones-constitutional-rights-act/> and Institute for Justice. (n.d., e). *Protecting Everyone's Constitutional Rights Ordinance* [Model legislation]. Arlington, VA. <https://ij.org/legislation/protecting-everyones-constitutional-rights-ordinance/>. State supreme courts also have a role to play. For example, the high courts of Montana and Nevada have declined to adopt qualified immunity doctrines of their own, ruling that people can sue state and local government officials who violate their state constitutional rights. *Dorwart v. Caraway*, 2002 MT 240, 312 Mont. 1, 58 P.3d 128; *Mack v. Williams*, 522 P.3d 434 (Nev. 2022).
- ¹²⁷ S.B. 20-217, 72d Gen. Assemb., Reg. Sess. (Colo. 2020), <https://leg.colorado.gov/bills/sb20-217>; H.B. 4, 55th Leg., Reg. Sess. (N.M. 2021), <https://www.nmlegis.gov/Legislation/Legislation?chamber=H&legtype=B&legno=4&year=21>
- ¹²⁸ N.Y.C. Admin. Code §§ 8-801 et seq.
- ¹²⁹ See, e.g., *Fraser v. Evans*, 992 F.3d 1003 (10th Cir. 2021). Mentioned in our introduction, this case involved Denver police officers threatening to arrest a man for recording them and illegally searching his device for the video. Even though four other circuits had found a First Amendment right to record the police and the Denver Police Department had trained the officers to respect that right, the officers received qualified immunity because the 10th

Circuit ruled the right was not clearly established in that jurisdiction. See also Brief of Amici Curiae Institute for Justice, ACLU, and ACLU-Colo., *Frasier v. Evans*, 992 F.3d 1003 (10th Cir. 2021), available at <https://ij.org/wp-content/uploads/2020/01/Frasier-IJ-amicus-brief.pdf>

¹³⁰ We used striptrf version 0.0.11.

¹³¹ For a discussion of this approach, see Derose, S. J., Durand, D. G., Mylonas, E., & Renear, A. H. (1990). What is text, really? *Journal of Computing in Higher Education Winter*, 1(2), 3–26.

¹³² We used NLTK version 3.4.5. Bird, S., Klein, E., & Loper, E. (2009). *Natural language processing with Python*. Sebastopol, CA: O'Reilly Media.

¹³³ The NLTK library offers 45 options for the part of speech, including categories for different types of punctuation and categories that break standard parts of speech into multiple subtypes. For example, it offers four categories for nouns: singular common, singular proper, plural common, and plural proper.

¹³⁴ These words were (with punctuation removed) court, courts, v, district, districts, f3d, us, sct, law, laws, judgment, id, defendant, defendants, plaintiff, plaintiffs, case, cir, led2d, see, rules, united, states, also. We identified these words through several exploratory methods using our primary training sample. Tagging stop words did not automatically remove them from subsequent analysis. Rather, it gave us the flexibility to easily remove these words if the situation dictated.

¹³⁵ In practice, this information was stored as two separate tables that could be merged. For simplicity, we treat it as one table here.

¹³⁶ For illustration purposes, we simplify the part-of-speech tagging and word stemming here. In reality, and as noted above, the NLTK tagger offers 45 options for the part of speech, many of which are too complicated to address here. Similarly, the word stemmer can sometimes make unusual stemming decisions. Accordingly, we showcase its application in theory.

¹³⁷ Originally, we sampled 10% of our dataset, totaling 716 opinions. However, after we began developing the algorithms, we decided to sample an additional 75 opinions to assess performance during the development process and fine-tune the algorithms (the validation data).

¹³⁸ Text searches using “regular expressions” go far beyond simply looking for a given word or phrase—they can seek out patterns based on the *types* of characters. For example, circuit court case numbers often appear in the format ##-#####, with the number of digits following the dash varying. We could therefore use regular expressions to search for two numbers, followed by a dash, followed by at least two more numbers. Regular expressions also allow for more flexibility by permitting multiple search options and optional characters. For instance, courts sometimes use a dash to separate the day, month, and year in a date, while other times they use a forward slash. Regular expressions allow for a single search to include both options. Using these features in concert (along with additional advanced options not discussed here) allows for powerful and precise text searches.

¹³⁹ The exact regular expression was `(?:\b[1iI]\b)first(?:\b[a-z0-9]\b)(?#avoids stopping at ... ; regex flags end of sentence)(?:right|violate|constitutional) (?:\b[a-z0-9]\b)(?:right|violate)* (?:\b[2iI]i?I\b)second[next|then] (?:\b[a-z0-9]\b)(?:clearly established|objectively reasonable)`

¹⁴⁰ Researchers working with these algorithms typically use a “bag of words” approach, measuring the frequency of each word’s usage in each text. This typically results in thousands of word frequencies, which researchers then feed into the model. Using a large number of words allows the algorithms to find unexpected associations that are helpful for prediction. However, it can also result in an abundance of statistical noise, especially with small datasets. And, indeed, given our relatively small training sample, this approach generally proved ineffective for reaching our performance goals.

¹⁴¹ In practice, we often used hierarchical indexing to search only a given portion of an opinion document—for example, the main text or the header. In such cases, the word count was for the specific portion of the document searched.

¹⁴² We calculated the standard deviation after normalizing the frequency counts.

¹⁴³ Some model inputs sum the frequencies of multiple related phrases, while others use the frequency for a single word. All else being equal, model inputs of multiple words or phrases are likely to result in higher frequencies and more variation than model inputs of a single word. Standardizing helps account for this variation, ensuring that all inputs are scaled equally. (We used inputs with multiple words or phrases to cut down on statistical noise. For example,

in our model predicting whether an opinion involved a violation of parental rights, we searched for several child welfare agencies [e.g., department of youth services, division of family services] and summed the resulting frequencies.

By combining these phrases, we created an input representing the concept of a child welfare agency, capable of identifying such agencies regardless of their specific name.)

¹⁴⁴ Cross-validation is a method commonly used in data science to approximate a model’s performance on unseen data. It involves splitting the data into model-building and evaluation subsets, with the subsets being continually shuffled so that all data are used at least once for evaluation. For example, when using five-fold cross-validation, four-fifths of the data is used for building the model, and one-fifth is used for evaluating the model. This process is then repeated five times, each time with a completely different evaluation dataset.

¹⁴⁵ Most notably, we used cross-validation to calibrate the regularization parameter on penalized logistic regression and support vector machine models.

¹⁴⁶ We used a ridge logistic regression model as opposed to a lasso or elastic net model because we frequently used multiple correlated predictors that we wanted to keep in the model. In our experience, when applying regularization, ridge tends to shrink correlated predictors in tandem, while lasso tends to choose one correlated predictor over the other. Lasso also can shrink predictors to zero (which we did not want), while ridge does not do this. Finally, ridge consistently outperformed lasso through cross-validation.

¹⁴⁷ Most notably, we removed duplicate opinions from the final dataset. We identified duplicate opinions using the circuit court and circuit court case number fields, both of which had greater than 99% accuracy. We found 269 opinions that shared a circuit court and circuit court case number with another opinion in our dataset. We reviewed these opinions manually to determine which opinion to keep. Generally, duplicates involved an original opinion that was later superseded by another opinion, most often as a result of panel rehearings, en banc rehearings, or remands from the Supreme Court. In these instances, we used only the most recent opinion, discarding any prior opinions. Occasionally, duplicates were standalone dissents or concurrences (often for en banc rehearings). These we also discarded. In addition to the duplicate review, we performed minor manual cleanup for fields that involved the direct extraction of text. To avoid inflating our algorithms’ performance, this was done following the evaluation of their effectiveness. Next, we merged in the “case origination date” and “district court case number” fields that were hand coded by our third-party contractor. And as a final measure, we corrected the machine prediction errors in our training and testing datasets by replacing the incorrect machine predictions with the human coding for the 11% of our dataset that was hand coded.

¹⁴⁸ The legal opinions themselves are available from Westlaw.

¹⁴⁹ We generally considered excellent algorithmic performance to be accuracy above 95% or an F1 score above 0.9. The fields included in Appendix C that did not meet these thresholds are federal law enforcement defendants (combined with state law enforcement defendants, which did have excellent algorithmic performance, in Table C2), federal prison defendants (combined with state prison defendants, which did have excellent algorithmic performance, in Table C2), non-police, non-prison (“other”) defendants, and qualified immunity denials. These fields’ algorithmic performance was still solid, however. Given their importance to our study, we have opted to include them.

¹⁵⁰ It was sometimes difficult to determine whether an opinion involving university, college, or school defendants involved premeditated retaliation or a direct restriction on speech. If plaintiffs were directly punished for their conduct (e.g., sexually harassing another student) and did not allege the punishment was for unrelated reasons, we considered that a direct restriction, regardless of how the claim was framed. However, if plaintiffs were punished but alleged the punishment was *really* because of a separate reason (e.g., their political speech), we considered it retaliation.

ABOUT THE AUTHORS



Jason Tiezzi

Jason Tiezzi is a data scientist who frequently collaborates with the Institute for Justice, including on projects such as *Too Many Licenses?* and the third editions of both *License to Work* and *Policing for Profit*. He holds a bachelor's degree from the College of William and Mary and a master's degree in data science from the University of Virginia.



Robert McNamara

Robert McNamara is the deputy director of litigation at the Institute for Justice. Since joining IJ in 2006, he has defended fundamental American liberties in state and federal courts nationwide, winning precedent-setting victories on behalf of entrepreneurs, property owners, and other victims of government overreach. His writing has been published in *The New York Times*, *The Washington Post*, and *The Wall Street Journal*, along with numerous other outlets, and his work has been featured by television and radio programs ranging from CBS's *Sunday Morning* to NPR's *All Things Considered*.



Elyse Smith Pohl

Elyse Smith Pohl is a legal research and policy attorney at the Institute for Justice. She brings legal analysis to the social science and policy research developed by IJ's strategic research team. Her areas of research include forfeiture, occupational licensing, and government accountability. Prior to joining IJ, she worked as an attorney in private practice. She earned her law degree from Ave Maria School of Law, her master's degree in law from the University of Navarre in Spain, and her undergraduate degree from the University of Virginia.

ACKNOWLEDGMENTS

The authors are deeply grateful to those who helped in ways large and small (mostly large). Lisa Knepper has served as a guiding light from the initial stages of this project, contributing in more ways than are possible to list. Perhaps more than anything, her knack for solving thorny problems saved us from running aground many a time. Mindy Menjou's influence is also everywhere: She expertly edited the text, contributed key substantive ideas, and handled a truly outrageous volume of logistical tasks. This project does not happen without them.

We similarly could not have done this project without the wealth of expertise possessed by attorneys Anya Bidwell, Patrick Jaicomo, and Keith Neely. They helped us create our codebook definitions, generated ideas for the predictive algorithms, and answered our seemingly endless parade of technical questions. Likewise, Dick Carpenter has been there every step of the way, beginning with the original idea for this project. His statistical knowledge and all-around wisdom have been invaluable throughout.

The text benefited from thoughtful feedback from Scott Bullock, Dana Berliner, Caroline Grace Brothers, and Marie Miller, all of whom contributed ideas that wound up in our final report. Additionally, Trey Barrineau's line edits were immensely helpful for streamlining and refining the final text.

Laura Maurice-Apel brought the report to life with her keen eye for color and expertly designed layouts, while Rima Gerhard skillfully transitioned the print report to the web.

Keith Diggs, Brennen VanderVeen, Brandon Livengood, and the aforementioned Keith Neely performed the thankless and time-consuming task of hand coding and labeling the data we used to build the algorithms. Their generosity of time and attention to detail were pivotal to this project's success. Our data checkers Sarah Eckhardt, Harrison Weeks, Zachary Popovich, and Anthony Laudadio left no stone unturned, as did our legal cite checker and proofreader Evan Lisull.

ABOUT IJ

Founded in 1991, the Institute for Justice is a nonprofit, public interest law firm. Our goal is to end widespread abuses of government power and secure the constitutional rights that allow all Americans to pursue their dreams.