

LAW AND ECONOMICS SEMINAR
Winter 2026

Professor A. Mitchell Polinsky

Thursday, February 26, 2026
4:15 - 5:45 p.m.
Stanford Law School Room 272

“The Economics of Healthcare Fraud”

Anup Malani
University of Chicago Law School

Note: It is expected that you will have reviewed this paper before the seminar. The body of the paper is just 23 double-spaced pages. The more technical appendix is optional.

The Economics of Healthcare Fraud

Jetson Leder-Luis* and Anup Malani†

October 2025

Abstract

Healthcare fraud imposes a sizable cost on U.S. public healthcare budgets and distorts health care provision. We examine the economics of health care fraud and enforcement using theory and data and connect to a growing literature on the topic. We first offer a new economic definition of health care fraud that captures and connects the wide range of activities prosecuted as fraud and allows us to taxonomize the literature. We define fraud as any divergence between the care an insurer says a patient qualifies for, the care a provider provides, and the care a provider bills for. Our definition clarifies the economic consequences of different categories of fraud and provides a framework for understanding the slate of existing studies. Next, we examine the incentives for committing and prosecuting fraud. Finally, we examine the prevalence of health care fraud prosecutions across types of fraud and types of care, and across the U.S., by machine learning on text data from Department of Justice press releases.

1 Introduction

Fraud involves the exploitation of an information asymmetry, and information asymmetries abound in the provision of healthcare (Arrow, 1963). When billing insurers, healthcare providers have opportunities to distort their care or the *reports* of their care to increase their profits. With the large-and-growing size of the healthcare sector, the frequency of fraud, and the value of anti-fraud measures, healthcare fraud is becoming a growing research topic within health economics and public finance.

While the exact magnitude of healthcare fraud is unknown, it is potentially quite large. The United States government spends nearly \$2 trillion per year on health insurance, and private insurers separately spend nearly as much. Improper payments in federal health insurance, a loose measure of funds the government knows were spent incorrectly, though not all necessarily fraud, amounted to more than \$100 billion in 2023.

*Boston University and NBER

†University of Chicago and NBER

The authors thank Mohammed Alshubaibi for excellent research assistance and Omar Khatab (developer of DSPy). We thank Zarek Brot-Goldberg, David Hyman, Aaron Kesselheim, Riley League, Ryan McDevitt, Jim Poterba, Jimmy Roberts, Maggie Shi, Matthew Stephenson, and participants of the PORTAL Seminar Series at Harvard Medical School, the ASHEcon conference, the Empirical Health Law Conference at Duke Law School, the Work-in-Progress seminar at the University of Chicago Law School, and the Harvard Law School seminar for their feedback.

The federal government recovers more than \$5 billion per year from anti-fraud enforcement (United States Government Accountability Office, 2024; U.S. Department of Justice, Office of Public Affairs, 2021), but likely recovers only a small portion of all fraud committed. These numbers reflect only the roughly 50% governmental share of insurance, so the bounds on healthcare fraud are perhaps double these levels, ranging from \$10 to \$200 billion annually. The midpoint of this range would imply that healthcare fraud in the US alone is comparable in magnitude to the annual budget of the SNAP program (~\$120 billion) or the GDP of Kenya (~\$110 billion).

While the financial value of healthcare fraud is high, the broader economic consequences are less clear. The laws defining healthcare fraud have been applied to extremely heterogeneous settings, from billing visits for patients never treated, to ordering lab tests without sufficient medical documentation, to the off-label promotion of drugs (U.S. Department of Justice, Office of Public Affairs, 2012, 2023, 2004). Similarly, the economics literature on health care fraud has examined widely varying settings—from hospitals to nursing homes and hospices, to ambulance companies, and to outpatient doctors office visits—that we highlight throughout this paper. These disparate frauds may have very different economic effects: some frauds may harm patients, some may waste scarce resources, some may impose mainly financial costs, and some may have neither allocative nor distributive effects. This heterogeneity of behaviors and potential effects suggests there is value to having an economic framework that maps disparate types of fraudulent behaviors to measures of efficiency and distribution. Importantly, such a framework would enable both a positive analysis of the drivers of fraud and its regulation, as well as a welfare analysis of both fraud and its regulation.

This paper endeavors to provide the economic framework to understand healthcare fraud and frame the existing literature. We begin, in Section 2.1, by summarizing the legal regulations that define fraud; these laws do not reflect our framework, but they do describe the scope of the types of conduct we study. In Section 2.2, we present our basic economic framework for understanding fraud in a healthcare system. This framework hinges on (a) the care that a payor (e.g., an insurance company) deems appropriate for a patients, (b) the care that a provider (e.g., hospital, doctor, etc.) delivers, and (c) the care that the provider bills the payor for. We define fraud as any situation where the payor’s choice, the provider’s care, and the provider’s bill differ. Our framework captures three of the most common healthcare frauds prescribed by fraud regulations: upcoding, substandard care, and medical necessity fraud. Upcoding is when a provider delivers the service that the payor expects, but bills for another service with a higher reimbursement. Substandard care exists when the provider delivers care that is of lower intensity or quality than what the payor deems appropriate, but bills for the appropriate care. Medical necessity fraud occurs when a provider delivers and bills for care that is more intensive or costly than what a payor believes necessary. We also discuss kickbacks, which are prosecuted under fraud-adjacent statutes but are not about the care provided per se.

Having set forth a framework that connects the law on fraud with the economics of fraud, we connect our framework to the economic literature on healthcare fraud in Section 2.3. We show how various empirical case studies of healthcare fraud fit into our framework. While a number of studies have discussed the policies used to combat healthcare fraud, no work has presented a cohesive, economic definition of fraud or considered systematically the welfare consequences of different types of healthcare fraud.

In Section 2.4, we use our framework to characterize the distinction between fraud, medical malpractice, and defensive medicine. Unlike fraud, medical malpractice does not depend on the provider's choice of what to bill for. Instead, malpractice focuses on whether the doctor's choice of care diverges from the appropriate care, i.e., non-negligent care. This highlights the different problems fraud and malpractice target: fraud concerns agency issues between the provider and the insurer, while malpractice concerns agency issues between the doctor and the patient.

In Section 2.5, we employ our framework to address welfare. We contrast the care a payor deems appropriate with the care that a social planner would prefer. This comparison highlights cases where frauds have efficiency effects and when frauds are merely transfers. For example, upcoding – the type of fraud that occurs when the right care is performed but the bill is wrong – operates merely as a transfer to the provider. Our framework also provides the distinction between health care fraud and waste, two related but often conflated phenomena. In short, fraud describes conduct that violates the legal statutes, while waste describes spending that the social planner would deem inefficient.

In Section 3, we consider two economic questions raised by our framework: what incentives do providers have to engage in fraud, and what incentives do payors and regulators have to combat fraud? Given the large gap between the upper bound on fraud and the level of fraud prosecutions, it is reasonable to suppose that only a small percentage of fraud is prosecuted. Given that, a more pointed way to ask our incentive question is: why is fraud so infrequently prosecuted, and why don't providers commit more fraud? In our discussion of provider incentives, we connect reimbursement policy to costs: there is incentive to commit medical necessity fraud when prices are far above cost (i.e., care is profitable regardless of need, and there is incentive to commit upcoding or medical necessity fraud when prices paid for care are lower than cost (so that fraud enables cost recovery)). In our discussion of regulator incentives, we discuss a growing literature on anti-fraud enforcement, and highlight the misalignment between the incentives of payors and the incentives of those tasked with regulation of fraud.

Finally, in Section 4, we provide new empirical evidence on government fraud prosecutions. Specifically, we analyze a new dataset of roughly 5,000 press releases scraped from the U.S. Department of Justice to identify fraud prosecutions associated with Medicare and Medicaid reimbursements. We use this scraped text as data to determine the prevalence of different, economically canonical types of fraud in federal fraud cases.

We highlight how certain types of fraud have received greater attention, and trends in fraud enforcement over time, across geography, and between medical specialties. This empirical overview provides a landscape of existing frauds that we believe will prove useful for guiding further economic research.

2 Economic Framework

2.1 Legal Definition of Healthcare Fraud

Fraud is fundamentally a legal concept: fraud exists because a form of lying is not permitted under the law. Therefore, before we can discuss an economic framework for healthcare fraud, we must discuss the legal definition or standard for such fraud.

Legally, healthcare fraud is any conduct that violates a healthcare anti-fraud statute. There are two main statutes that define fraud against federal health insurance programs. First, the federal criminal Health Care Fraud Statute (18 U.S.C. § 1347) makes it a crime to knowingly obtain money from a federal healthcare program under false pretenses. Individuals convicted under the criminal statute face both jail time and monetary penalties. Second, the False Claims Act (31 U.S.C. § 3729) provides civil penalties for individuals who cause money to be paid from a healthcare program that should not have been. Because this statute has only monetary penalties, it has lower standards both for burden of proof and for how intentional the fraud must have been. There is also a criminal False Claims Act (18 U.S.C. § 287), which imposes criminal monetary penalties with a higher evidentiary standard. Administratively, providers convicted of criminal offenses are generally prohibited from any future billing. There are state analogues to these statutes that prohibit unwarranted payments from state health insurance programs.

A number of state statutes proscribe fraud against private insurance companies. Most are similar to the civil False Claims Act in that they authorize only civil penalties (e.g., Mass. Gen. Laws ch. 12, §§ 5A-5O), though a few authorize jail time (e.g., Cal. Penal Code § 550). Moreover, civil penalties are typically limited to the return of fraudulently obtained funds and forward-looking exclusion from coverage under an insurance plan.

In addition to the statutes that prohibit fraud, there are a set of federal and state laws that prohibit kickbacks or referral fees. These statutes prohibit providers from paying either patients or other providers for receiving or referring care, respectively. The goal of these laws is to eliminate incentives to provide care where care is medically inappropriate. For example, the Stark Law prohibits physicians from referring patients to other medical providers in which they have a financial interest, and the criminal Anti-Kickback statute (42 U.S.C. §1320a-7b(b)) prohibits remuneration to induce patient referrals or use of care. Various state laws

mimic the federal statutes to protect state-run Medicaid programs (305 Ill. Comp. Stat. § 5/8A-3(b)-(c)) and private insurers (e.g., Tex. Hum. Res. Code Ch. 32).

The economic framework for fraud in Section 2.2 does not include kickbacks. That taxonomy captures the different types of medical care and billing that constitute fraud, while anti-kickback laws prohibit how rents from that fraud are *divided* among providers. The latter behavior is economically distinct from directly committing fraud. We will discuss the relationship between our fraud framework and anti-kickback statutes in Section 3.

The scope of liability for fraud is quite broad. For some billing requirements, a provider may be liable for fraud even if she was unaware of those requirements or did not intend to fail those requirements, because the legal standard is recklessness, or that a provider “should have known”. Appendix A discusses federal health care billing requirements and the legal elements of a fraud claim.

2.2 An Economic Framework for Fraud

In this section, we provide an *economic* definition of healthcare fraud that encompasses different healthcare provider behaviors that all qualify as fraud, but distinguishes those frauds based on their economic consequences. To construct such a definition, we develop a framework that focuses on the difference between (a) the care that a payor, e.g., Medicare, determines is appropriate for a patient, (b) the care a provider, e.g., a doctor, delivers, and (c) the care that the provider bills the payor for. Our economic definition of fraud is any discrepancy between (a), (b) and (c).

To construct our framework, consider the relationship between the care the payor deems appropriate and the provider’s actions. The first actor in our framework is the payor, which could be the government (e.g., Medicare) or a private insurance company. This payor determines the appropriate level of care that the patient qualifies for under insurance rules. To simplify our framework, we allow the payor to determine whether the patient should receive a low level L of care or a high level H of care.¹ The second actor, the medical provider (e.g., a doctor, lab or hospital) engages in two types of actions: providing care to a patient and billing for care provided to the patients. For simplicity, we suppose that the provider can provide a low level or high level of care for a patient, and can bill for a low level or a high level of care.

Table 1 lays out the relationships between the two actors and their three actions. The rows show the payor’s action (a), whether the payor determines that the patient’s condition warrants low care or high care, L or H. The columns track the provider and indicate (b) whether the provider delivered low care or high care. Each cell tracks the billing decision of the provider; these cells have two rows, corresponding to (c)

¹We can generalize to a continuum of levels of care, but for expository purposes two levels suffice. In our framework it is best to think of low care as low-reimbursement care and high care as high-reimbursement care, rather than assume that the levels of care correspond to the medical intensity of intervention, e.g., whether to provide an arterial graft rather than a stent.

whether the provider billed for a low level of care or high level of care.

Table 1: Economic Categories of Fraud

		Provider	
		L	H
Payor	L	L: <Not fraud>	L: <Not observed>
	H: Upcoding	H: Med. unnecessary care	
H	L: <Not prosecuted>	L: <Not observed>	
	H: Substandard care	H: <Not fraud>	

The combination of these three decisions – low or high for each of the payor’s determination, the provider’s choice of care, and provider’s billing – fully taxonomizes all types of healthcare fraud, as well as a number of combinations that are not fraud.

First, we describe what behavior is not fraud, as well as the set of actions that are rarely observed or prosecuted. If the provider delivers and bills for the same level of care that the payor determines is appropriate for the patient, then there is no fraud. These combinations of action are marked as <Not Fraud> in the table. There are two such cases: the provider delivers, provider bills, and payor deems appropriate a high level of care (HHH) or all three actions coincide with low care (LLL).

Our taxonomy also identifies theoretical conduct that is rarely observed: situations where the provider delivers high care but bills low care. This is not observed because the doctors have little incentive to provide care without billing for it when the patient is insured. We mark these cases with <Not observed>. In theory, these actions could be considered fraudulent (including by legal standards) because the bill mismatches the care; but they are not observed and therefore we do not explore them any further.² Third, fraud that is rarely if ever prosecuted includes situations where the payor determines that the patient deserves high care, but the provider delivers and bills for low care. Although the patient potentially suffers, the lack of financial harm to the insurer makes this type of fraud unlikely to be enforced through anti-fraud statutes.

Our framework highlights three distinct types of healthcare fraud. The first type involves situations where the payor chooses and the provider delivers low care, but the provider bills high. This type of behavior is called “upcoding”, because coding is a synonym for billing. Upcoding is the canonical example of economic

²Providing high care and billing for low care is sometimes called “undercoding”. The AAPC (American Academy of Professional Coders) writes, “deliberate under coding is, in reality, ‘making a false statement’ about the services provided, and is ultimately a ‘misrepresentation’ of the facts. The [CMS fraud] fact sheet gives an example of fraud as ‘knowingly billing for services that were not furnished,’ which would apply if services are purposefully under coded.” See: <https://www.aapc.com/blog/26957-undercoding-is-no-better-than-overcoding/>. However, in practice this appears to rarely if ever be enforced.

(and legal) fraud. A prominent example of upcoding was the Columbia-HCA lawsuit settled in 2000, wherein the hospital billed for patients as if they had been treated for severe pneumonia when in fact they were only treated for simple pneumonia (U.S. Department of Justice, Office of Public Affairs, 2000).

A second type of fraud involves situations where the payor chooses and the provider bills for high care, but the provider delivers low care. This type of behavior is called providing “substandard care” because of the (downward) discrepancy between appropriate/billed care and the care delivered. A clear example is a case where nursing home owners provided inadequate care in poor quality facilities when patients needed more serious treatment (U.S. Department of Justice, Office of Public Affairs, 2019b).

A third type of fraud in our framework involves cases where the provider delivers and bills for high care, but the payor believes low care is appropriate. We call this behavior “medical necessity fraud”. In general, this comes with a lie, wherein the doctor falsifies the documentation that allows for them to bill high, by overstating the patient’s medical necessity. An example from litigation is the widespread use of expensive genetic cancer testing for patients who had no indication what is warranted these tests (U.S. Department of Justice, Office of Public Affairs, 2019a).

2.3 Mapping to the Literature

Our model of fraud allows us to sort the large and growing literature on health care fraud based on the category of fraud committed. This literature also evaluates anti-fraud enforcement, which we discuss in Section 3. Table 2 maps case studies of fraud from the economics literature (and from Department of Justice press releases) into the three common types of healthcare frauds in our taxonomy. Here we highlight specific papers and their relationship to our framework.

Upcoding. In early work, Silverman and Skinner (2004), Dafny (2005), and Becker et al. (2005) explored the prevalence of hospital upcoding, i.e., hospitals submitting bills for more expensive care than they provided. Many of these cases of upcoding concerned hospital complications and chronic conditions, events for which hospitals obtain higher reimbursement rates but are hard to verify. These fit into our framework as upcoding, because the needed level of care was administered, but the bill submitted to Medicare listed a higher level of care. More recently, Joiner et al. (2024) investigate upcoding across providers with data from Medicare audits, and show that 0.5% of hospitalizations and 2.5% of physician services have indicia for upcoding.

Shekhar et al. (2023) describe patterns in billing codes associated with insurance claims that are consistent with hospital upcoding behavior. They use the frequency of these patterns in claims data to rank hospitals by propensity to upcode. This paper is part of a larger computer science literature that uses healthcare data

to detect fraud, though that literature is largely focused on statistical methods rather than the extent of fraud. Kumaraswamy et al. (2022) provide a survey of the detection methods literature.

Leder-Luis (2025) and Gupta et al. (2024) discuss a different type of hospital upcoding involving the outlier payment system, a loophole that permits extra payment for expensive hospital patients. The former paper shows that anti-fraud enforcement produced large deterrence effects on upcoding, while the latter shows that for-profit hospitals that received fraudulent windfalls used these funds to enrich their executives and shareholders. This specific form of fraud maps to upcoding because there was no distortion of the type of care provided, but rather, just the type of bills that were submitted.

Finally, Sanghavi et al. (2021) discuss ambulance companies that bill payors for giving patients rides that never occurred, a phenomenon called ghost billing. This example highlights a potentially counter-intuitive category of upcoding cases where the service needed and provided is $L = 0$, but the provider billed for $H > 0$. Notably, this highlights a point in our framework, which is that it can be difficult to know the appropriate level of care; should these patients have received ambulance rides, i.e. was the correct level of $L = 0$? As another example, sending a visiting nurse but billing for a physician can be considered substandard care if the physician is needed, but could be thought of as upcoding if the nurse was sufficient but the doctor was billed for.

Medical necessity fraud. Much of the literature on health care fraud revolves around medical necessity fraud. Howard and McCarthy (2021) investigate the medically unnecessary use of implantable cardiac defibrillators, claims for which were prosecuted under the False Claims Act, and show that the prosecution led to large reductions in overuse. Shi (2023) discusses medically unnecessary hospital inpatient admission in the context of the Recovery Audit Contractors, private companies that were empowered to detect Medicare fraud through audits, and who targeted these unnecessary hospital admissions. Leder-Luis (2025) discusses inpatient spine surgeries that should have been performed on an outpatient basis, and shows that litigation reduced the over-utilization of inpatient care. Eliason et al. (2025) discuss the medically unnecessary use of ambulances for dialysis patients, which led to dozens of criminal indictments, as well as civil lawsuits and regulatory changes. O'Malley et al. (2023) document criminal home health fraud wherein patients received visiting nurse services that they did not require. Taken together, the literature on medically unnecessary care tends to show that enforcement against medical necessity fraud can generate high cost savings without patient harm, because the care was not needed in the first place.

A handful of papers shed further light on the welfare effects of medical necessity fraud. Leder-Luis (2025) discusses a whistleblower lawsuit concerning the promotion of Botox for headaches, for which the manufacturer (Allergan) had not received FDA approval. Under our model, this constitutes a type of medical necessity fraud, as the payor had determined that patients did not qualify. However, Botox eventually

Table 2: Examples of Healthcare Fraud with Categorization

Category	Example	Citation	Fraud Amount
Upcoding	Hospitals report expensive DRGs	Silverman and Skinner (2004) Dafny (2005) Becker et al. (2005) Joiner et al. (2024) Shekhar et al. (2023)	\$840 million Settlement w Columbia HCA + Many smaller cases
	“Outlier” hospital payment fraud	Leder-Luis (2025) Gupta et al. (2024)	\$788+ million settlements
	Home health care never provided	O’Malley et al. (2023)	
	Ghost ambulance rides	Sanghavi et al. (2021)	1.8% of rides
Medically unnecessary care	Unnecessary implantable cardiac defibrillators	Howard and McCarthy (2021)	\$280 million in settlements, \$2.8B in deterrence
	Hospices sued for enrolling unqualified (non-terminal) patients	Gruber et al. (2025)	\$330 million in settlements
	Patients receive spine surgery inpatient instead of outpatient	Leder-Luis (2025)	\$214 million in settlements
	Ambulances used as taxis for dialysis patients	Eliason et al. (2025)	\$4 billion estimated losses
	Unnecessary inpatient hospital admissions	Shi (2023)	
	Durable Medical Equipment not needed by beneficiaries	Diwan et al. (2025)	
Unnecessary home health services	Einav et al. (2025)		
Substandard Care	Doctors spend less time with patient than billed	Fang and Gong (2017)	
	South Florida Nursing home owner provides poor quality facilities		\$1.2 billion damages

received FDA approval for migraines, calling into question whether the care was actually inefficient. Gruber et al. (2025) discuss the use of anti-fraud statutes to target for-profit hospice facilities that admit patients that may not qualify for hospice. However, the authors conclude that this fraud may have actually *increased* welfare. For-profit hospice saves the government money by displacing more expensive care, and while hospice reduces patient lifespans, it also improves patient quality of life.

Finally, the literature on the over-prescription and diversion of opioids from medical to recreational does not claim to discuss healthcare fraud, but nonetheless provides examples of medical necessity fraud. Diversion supplies opioids to users who were unable to obtain legitimate prescriptions, i.e., involves treatment payors deem unnecessary. Maclean et al. (2020) provide a survey of the economics literature on opioids, though fraud is not a primary theme of the studies reviewed.

Substandard care fraud. Examples of substandard care are difficult to find in the economics literature. There is a substantial literature on examples of and on the extent of medical malpractice. Provider negligence

is often an indicator of substandard care because providers typically define appropriate care as non-negligent care. However, negligence is not sufficient to demonstrate fraud, which also requires that the provider billed for more or higher quality care than they actually provided. Similarly, there is a substantial literature on the (low) quality of healthcare. This includes the literature on quality issues like staffing ratios in nursing homes, such as Hackmann (2019) and Gandhi and Olenski (2024). However, like malpractice, low quality does not demonstrate fraud without evidence of billing for high quality care.

A second reason for the dearth of research on substandard care fraud is that, without evidence on what was needed, it is difficult to determine whether low quality care is substandard care or upcoding. For example, Fang and Gong (2017) discuss a method for detecting health care fraud by looking for doctors who bill for too many hours per week, indicating they spend less time than billed. This could be a form of substandard care fraud, if the patients needed the entire time billed for, as a form of doctor shirking. However, this could also be a case of upcoding, if the lesser time spent was medically appropriate but the bill was incorrect. This work does not address the welfare of patients.

One potential limitation to this literature is the difficulty in assessing quality of care using data available to researchers. Much health economics literature uses claims data, which contain information like billing codes and medical history. However, indications of low-quality care do not necessarily show up in claims data. The rise of research using electronic health records data, including data on physician time use, may provide an avenue for future exploration.

Next, we turn to the relationship between our model for fraud and a related legal and economic phenomenon: medical malpractice.

2.4 Medical Malpractice

Healthcare fraud arises from a breakdown of the principal-agent relationship between the doctor and insurer, and it is useful to contrast it to the corresponding principal-agent relationship between the doctor and patient: malpractice. The legal definition of malpractice is the provision of “unreasonable” level of care in a doctor-patient relationship, as determined by a court, with reference to the needs of the patient.³

We can map malpractice into a framework that is analogous to the one we use to define fraud, but with two differences. First, instead of considering the level of care that a payor deems appropriate for a patient, the malpractice framework focuses on the level of care a patient needs. Second, what a provider bills is not relevant for determining malpractice (though it may affect the level of damages the provider pays if found

³More precisely, a provider has the duty to provide the level of medical care that a reasonable provider would. What is reasonable may be determined by reference to data on what most doctors do in similar circumstance, guidelines provided by societies of doctors, or even expert determinations. Courts, typically with help from juries, determines the standard of care and whether it is met. Failure to fulfill this duty to provide reasonable care constitutes medical malpractice.

liable for malpractice). Table 3, which presents our malpractice framework, therefore replaces payor with patient need and drops the division of each cell that reflects billing.

Table 3: Medical Malpractice

		Provider	
		L	H
Patient Need	L	<Not malpractice>	Non-enforced malpractice/ defensive medicine
	H	Malpractice	<Not malpractice>

In our malpractice framework, misbehavior is found on the off-diagonal cells.⁴ If a provider delivers the correct level of care, she is not liable for malpractice. If a reasonable provider would deliver high care, but the actual provider delivers low care, she is liable for malpractice (red) (e.g., Danzon, 1984; Farber and White, 1991; Frakes, 2013; Mushinski et al., 2022); this is the canonical behavior targeted by the malpractice system.⁵ If low care is appropriate, but the provider delivers high care, the physician is technically liable for malpractice, but is unlikely to be sued. However, if the reason the physician provided high rather than low care is to reduce her probability of being sued for canonical malpractice, then behavior in this cell is called defensive medicine (yellow) (e.g., Kessler and McClellan, 1996).⁶

A comparison of Table 1 and Table 3 reveals the distinction between malpractice and fraud. When patient need and provider billing decisions align on the vertical axis, then these tables overlap, and we see that medical malpractice occurs on the off-diagonals – where either medical necessity fraud (top-right) or substandard care fraud (bottom-left) have occurred, provided the provider bills for the expensive care. In general, canonical medical malpractice occurs in the bottom-left (red) corner, and therefore aligns most closely with substandard care fraud.

⁴In this paragraph we focus on relevant papers on malpractice in the economics literature. However, there is also substantial evidence on malpractice and defensive medicine in the medical and health policy literature as well. See, e.g., Danzon (2000) and Kessler (2011) for reviews of the broader literature, and (e.g., Baicker et al., 2007; Mello et al., 2010; Lakdawalla and Seabury, 2012) for estimates of financial and welfare costs.

⁵Seminal papers by Brennan Troyen et al. (1991), Studdert et al. (2000), and Studdert David et al. (2006) examine rates of malpractice (i.e., treatment error) and rates of litigation given errors.

⁶These papers on defensive medicine focus on providers doing more to avoid liability, sometimes called positive defensive medicine. There is also a substantial literature on negative defensive medicine, which the act of doing less to avoid liability. For example, Klick and Stratmann (2007); Matsa (2007); Hyman et al. (2015), and Malani and Reif (2015) examine whether liability reduces physician supply.

2.5 Welfare

Our framework also allows us to consider how different frauds affect social welfare. To do so, we consider the level of care the social planner would have the provider deliver, rather than the level the payor would recommend. This helps determine when fraud is inefficient or wasteful and when it is a transfer that mainly affects distribution.

We define the level of care that a social planner recommends, minimally, as care for which the marginal benefit to the patient (or society if there are external effects) is greater than the (marginal) cost of care to the provider.⁷ A fuller definition would allow that there are multiple levels of care where the marginal benefit is greater than cost, but chooses the level of care where the gap is the greatest. Thus, in our simplified framework, when we say the planner prefers a high level of care, that means the surplus from high is greater than the surplus from low. Importantly, we reserve for later the discussion of price, which may be less than, equal to, or greater than the cost of care. Price will be relevant for understanding why providers engage in fraud or malpractice, but is not essential to judgments about allocative or distributive effects.

The social planner’s preferences do not necessarily align with either the payor’s determination (which defines fraud) nor with the patient’s need (which determines medical malpractice). First, the social planner’s preferences may diverge from what is categorized as fraud, as fraud concerns the insurer’s billing rules, but does not internalize whether the care provided was potentially efficient. Similarly, the “reasonable provider” standard malpractice courts use is not, in practice, equal to a “cost-effectiveness” or similar social-welfare standard, and so cases may be labeled as malpractice for failure to treat (or, respectively, for overtreating) even when that decision was more socially beneficial.

Table 4 describes the welfare consequences of fraud and malpractice. It replicates Table 1, but, assuming that the payor and patient agree, overlays our medical malpractice framework and highlights medical malpractice in red and defensive medicine in yellow. Next, we add the social planner’s beliefs about when a provider’s behavior is waste, inefficient, or a transfer, in bold.

The resulting table provides four takeaways. First, when there is a match between the payor’s preferred level of care and the care delivered by the doctor, there is no distortion of care, and therefore fraud only acts as a transfer. In upcoding fraud, where the doctor delivers low care and bills high, there is only a *transfer* from the insurer to the provider. (If there is a co-pay and it is collected, then there may also be a transfer from the patient to the provider.) However, this type of fraud does not cause the wrong care to be provided,

⁷In our theory presented in Section 3 and Appendix B, we use an even more exacting standard, and say that welfare recommends the care that generates the maximal surplus, defined as the difference between a fully-informed patient’s willingness to pay for a type of care and the provider’s cost of providing that care.

We elide issue of whether the proper standard is marginal cost or dynamic average cost. The latter may be appropriate where there are fixed costs of, e.g., innovation. This issue is not, however, germane to our discussion of fraud.

Table 4: Welfare Consequences of Fraud and Malpractice

		Provider	
		L	H
Payer	L	L: <Not fraud> H: Upcoding Transfer to Provider	L: <Not observed> H: Med. unnecessary care Efficiency Unclear
	H	L: <Not prosecuted> H: Substandard care Usually Inefficient Transfer to Provider	L: <Not observed> H: <Not fraud> Potential Waste

i.e. there is no distortion in the allocation of the medical effort.

We find allocative effects in the off-diagonal cells. When the provider should deliver high, but instead delivers low (see substandard care in the red-shaded cell), there is both inefficiency from distorted care, and also a transfer – providers gain, and patients lose. Payors also lose if, dynamically, the inefficiently low care imposes greater costs on the payor down the line. In theory, there could be circumstances where substandard care is efficient, such as where payers require a higher level of care than is socially optimal, but there are no clear examples of this efficiency in the literature.

The most uncertain examples for welfare arise when the payor thinks the doctor should deliver low, but instead the doctor delivers high (see medical necessity fraud in the yellow-shaded cell). More care is produced than is permitted under billing rules, so fraud has occurred. However, the welfare effects are unclear. For example, if the high care is highly valuable to the patient, or offsets more expensive care, the provider may increase surplus by providing that care, *even if they commit fraud in doing so*. In the economics literature, a canonical example of this comes from Gruber et al. (2025), who show that fraud in the hospice sector led to a relaxation of eligibility rules, but that the accompanying hospice expansion was potentially welfare *enhancing*, as it both saved money and benefited patients. Nevertheless, the conduct they describe was considered, and prosecuted, as healthcare fraud.

Finally, our table also allows us to distinguish between fraud and waste. Waste occurs when doctors give high care when, from a social welfare perspective, low care should have been given. This is irrespective of how the provider bills. Waste can occur in circumstances with or without fraud, and is about the relationship

between the provider’s choices and the social planner’s preferences, not about the distinction between the provider’s choice and the payor’s choice. In cases of medical necessity fraud, i.e. the provider gave and billed for higher care in contrast to the payor’s determination of low care, then waste can occur if the social planner agrees with the payor that low care was appropriate. In this circumstance, waste and fraud align. However, waste can also occur in the bottom right diagonal cell, without fraud: if by the social planner’s welfare calculations the provision of high care was inefficient, then waste occurs even if the provider did not violate the payor’s billing guidelines and therefore has not committed fraud.

In summary, the welfare analysis here allows us to highlight key features of health care fraud. Fraud is a legal classification, not an economic behavior, and therefore, different behaviors all classified as fraud have very different welfare effects, including potentially positive ones. Moreover, waste, which is a welfare criterion, does not always align with fraud because waste is about social welfare and not about insurance company rules.

3 Incentives to Commit and Enforce Against Fraud

3.1 Incentives to Commit Fraud

The literature on incentives to commit fraud is sparse. There is a debate among scholars about the main drivers of fraud. Feldman (2001) and Rai (2001) stress the pecuniary incentives – particularly prices and payment structures – to commit fraud. By contrast, Hyman (2001) argues that social norms are responsible for fraud.

In Appendix B, we sketch a model of the incentives of providers to commit fraud, based on their available treatments, reporting, and payment rates. This model provides intuition about the incentives to commit each of the different types of fraud in our framework, and when those frauds are (in)efficient.

Intuitively, our model highlights two important drivers of fraud: inadequate penalties for fraud and setting too high or low a price for a procedure. First, if a payor does not or cannot adequately punish providers who are revealed to have committed fraud after an audit or litigation, then providers will have an incentive to upcode—to file a claim for a more expensive procedure than the procedure that they actually committed. The logic follows from Becker’s model of crime: if one lowers the cost of fraud, rational providers will commit more fraud.

Second, if a payor sets reimbursement rates—prices—for a procedure too high or too low, they will encourage medically unnecessary care or substandard care. Ideally, payors set prices of different procedures a provider may deliver such that each procedure offers the provider the same profit. Then the provider

chooses the most appropriate procedure the payor recommends or the provider's altruism towards patients breaks ties between procedures. However, if the payor sets too high a reimbursement for a procedure such that it offers a higher profit than other procedures, then the physician will have an incentive to perform that procedure even when it is not appropriate care, e.g., will have an incentive to provide medically unnecessary care if unnecessary care offers higher profit. Likewise, if the payor sets too low a price for a procedure the payor deems appropriate such that it offers a lower profit than other procedures, the provider has an incentive not to perform the recommended procedure, e.g., to deliver substandard care if lower quality care offers higher profits than appropriate care.

The literature on sources of fraud is a potential area for future research. Diwan et al. (2025) highlight competition and market structure as a potential source of fraud, noting that price competition can drive prices down in a way that selects on low-cost, fraudulent firms. However, many questions remain open in this literature, and in particular, there may be an avenue to relate to the growing literature on vertical integration in health care.

3.2 Anti-Fraud Enforcement

A growing literature has addressed efforts by the federal government to curb health care fraud. Generally, these studies use examples of specific frauds, including some discussed in Section 2, to evaluate anti-fraud policies for their impacts on spending, rates of fraud, and patient health outcomes.

One area of literature has evaluated the effects of the False Claims Act, which provides civil penalties for health care fraud. In empirical legal work, Engstrom (2012) discusses the volume of anti-fraud civil cases under the False Claims Act, and shows that roughly 2/3 of all civil fraud cases are healthcare cases. Howard and McCarthy (2021) discuss the deterrence effects of anti-fraud civil litigation under the false claims act, using cardiac defibrillators as an example; they find large deterrence effects, about 10 time larger than the settlement amounts. Similarly, Leder-Luis (2025) examines case studies of large settlements totaling roughly \$1.9 billion under the False Claims Act and shows that these cases produce deterrence effects of about \$19 billion, mirroring the 10x result. The general takeaway from this literature is that private civil anti-fraud enforcement through the False Claims Act is effective at deterring fraud by large providers such as hospitals.

Comparative work has examined methods other than *ex post* litigation for eliminating fraud. Eliason et al. (2025) compare enforcement under civil and criminal litigation to up-front prepayment authorizations and argues that preventative regulation is more effective when fraud is diffuse, i.e., committed by many small providers. Shi (2023) examines the overuse of inpatient care and finds that privatized audits were effective at eliminating this conduct. To be clear, however, while some of this care was subject to anti-fraud scrutiny,

the focus of the Shi (2023) paper is on waste. Einav et al. (2025) compare the use of criminal enforcement versus a moratorium on new providers to deter home health care fraud, which is usually a medical necessity fraud.

Finally, some work has addressed state efforts to combat Medicaid fraud. Becker et al. (2005) examine the impact of state Medicaid anti-fraud expenditures on abusive Medicare billings that they have reason to believe would be responsive to state-level Medicaid fraud enforcement. They show that increased enforcement leads certain types of patients and hospitals to have lower billings, without adverse consequences for patients' health outcomes. In contrast, Perez and Wing (2019) and Nguyen and Perez (2020) examine state-level enforcement of Medicaid fraud and find limited effectiveness of decentralized enforcement by the states, though neither study distinguish which types of fraud are pursued or deterred.

Although our review has focused on the economics and quantitative empirical literature on fraud, there is a large legal literature that has examined how different legal rules affect the ease and efficiency of anti-fraud enforcement in law. Drake (2013) thoroughly surveys this literature with extensive citations to case law and legal scholarship.

There are two main takeaways from the empirical anti-fraud literature. First, different tools are effective on different types of fraud. While civil lawsuits can have big deterrence effects when employed against large providers, they are largely ineffective against small firms. In contrast, preventative measures are more effective against frauds committed by multiple small firms, because they can scale to regulate more firms and can overcome the problem of limited liability for *ex post* damages.

Second, taken together, the literature shows a lack of cohesive government strategy on health care fraud. Each of the studies on enforcement shows multiple attempts to combat the same behavior, and myriad examples of multi-billion dollar frauds being allowed to persist. This may arise either due to limited incentives for enforcement due to poor bureaucratic incentives, or because of organizational concerns due to the many arms of the US government that attempt to enforce anti-fraud policy with poor coordination. Moreover, it is not clear why certain frauds are pursued while others are left unchecked, which may also have to do more with ease of enforcement rather than the social cost of different frauds.

In Appendix C, we sketch a model of the incentives for the US government to enforce against fraud. We model the tradeoff between enforcing against health care fraud and pursuing other areas of law enforcement, given that the Department of Justice has wide responsibility beyond this area. Ultimately, health care fraud enforcement levels are driven by the government's budget and relative political priorities.

4 Empirical Evidence on Healthcare Fraud and Enforcement

In this final section we provide descriptive statistics on the extent and nature of enforcement of healthcare fraud laws. Before we dive in, let us start with some caveats.

Enforcement is an imperfect way to measure fraud. Not all healthcare fraud is caught, and there is no code in claims data that labels claims as fraud. Moreover, as we explained in Section 3.2, not all fraud is prosecuted. Therefore enforcement is a highly non-specific test for or measure of fraud, because cases not prosecuted do not prove no fraud. However, successful prosecution (i.e., prosecution with a positive trial verdict or settlement) is a fairly sensitive test of fraud, meaning liable or guilty verdicts are very likely frauds. Even so, unless cases of fraud are randomly detected and prosecuted, cross-sectional comparisons of successfully prosecuted cases will conflate variation in fraud with selection in prosecution.

Successful prosecution also does not tell us about enforcement rates per se. For one thing, without the true rate and distribution of fraud, we cannot determine the fraction or nature of cases prosecuted. For another, we cannot determine incentives for healthcare fraud prosecution without knowing returns and costs of alternative non-fraud cases.

That said, there are interesting patterns in the prosecution. Moreover, given the importance of fraud in the healthcare system, even imperfect data is better than no data. Therefore, we examine new data on the universe of health care fraud that was subject to enforcement.

4.1 Data and methods

We collected data from the Department of Justice Press Release archive. Each press release details a substantial step in an anti-fraud case, generally a settlement or judgment, although press releases can sometimes be released when the government initiates a case, i.e., indicts, a defendant responsible for a particularly egregious fraud.

We assemble, clean, and categorize our data in five steps. First, we scrape 6,909 press releases that contain the word “Medicare” or “Medicaid” and are from 2013-2023.⁸ Second, we use search terms to exclude press releases not related to a fraud litigation,⁹ leaving 6,112 press releases.

Next, we categorize cases in press releases by the type of fraud involved. Each cases may involve multiple categories of fraud. To start, We use search terms to identify kickback cases. For the remaining types of fraud, we manually categorize 500 randomly-selected press releases, and use those to fine tune a prompt for

⁸The focus on Medicare and Medicaid may potentially miss fraud against the Veterans Affairs system, the Federal Employees Health Benefits Program, or the military TRICARE/CHAMPUS system. However, these cases often contain a Medicare fraud component as well.

⁹Specifically, we eliminate anti-trust cases and income tax violations from our dataset. We also eliminate press releases from end-of-year overviews or awards ceremonies that detail year-long activity across cases (even with regard to fraud cases) instead of individual cases.

a large-language model (LLM) that classifies remaining press releases. This process is able to label all but 956 press releases, 15.64% of our sample.

Fourth, we validate the non-categorized press releases. We randomly select 100 of these 956 cases and manually categorize them as fraud or not. We use this 100 to train a prompt for an LLM that identifies which of the 956 cases are actually not fraud and drop the non-fraud cases.

This process yields 5,663 press releases concerning fraud actions, with 507 lacking a categorization. This sample approximates the universe of all enforcement related to federal Medicare or Medicaid healthcare fraud. Finally, for each of these 5,663 press releases, we scrape the state in which each fraud occurred and the specialty of medical care involved. The latter is done by matching text in each press release with medical specialty terms derived from the AAMC specialties list (American Academy of Medical Coders, 2024). Appendices D and E contains details about our sample and the categorization process, respectively. Our LLM categorization closely matches our hand-coded sample, indicating that our AI-assisted categorization is reasonably accurate.

Our data must be interpreted with caution. First, we tabulate press releases rather than cases. It is possible the same case is the subject of two press releases, e.g., one at the time of indictment and another on conviction. This means that the number of releases may be greater than the number of cases. However, if this error is not systematically related to covariates such as time, location or subject, then it will not affect relative comparisons. Second, the prompt-optimization method we use to categorize types of fraud may have a small rate of error, but we cannot be sure that error is entirely random. Rates of sensitivity, i.e., the probability that cases of fraud are labeled as fraud, vary from roughly 70-90% across types of fraud.

4.2 Locus of Fraud

First, we provide an overview of the patterns in healthcare enforcement over time, geography, and medical specialty. Figure 1 shows that reported enforcement of health care fraud has been stable around 500 press releases per year over the past decade.

Medicare and Medicaid fraud appears to be more common in Southeastern and Northeastern states. Figure 2 shows the geographic distribution of fraud across states. Panel (a) reports the number of enforcement actions per state.¹⁰ Because we are focused on reported cases of fraud against Medicare and Medicaid, states with larger populations covered by those programs will likely have more cases, even if the rate of fraud per instance of care is the same across states. To address this, we adjust case frequencies for the population covered by Medicare or Medicaid in a state in Panel (b).¹¹ Now we see that health care fraud enforcement

¹⁰Washington is omitted as we are unable to distinguish between the state and the District of Columbia.

¹¹Specifically, we report the share of all press releases that are in a state divided by the share all Medicare beneficiaries that




Figure 1: Number of DOJ press releases on healthcare fraud cases over time

cases are more common on a per-beneficiary level in the American Southeast and Northeast.¹² An interesting comparison that illustrates the value of adjusting for Medicare and Medicaid exposure is California and Louisiana. The former has more cases but also a larger Medicare population. Once we adjust for Medicare population, the latter actually has a higher rate of fraud cases. Future work should explore whether case reporting is a function of the nature of procedures performed in a state or the size of DOJ offices in a state.

There is substantial variability in the number of reported cases of fraud enforcement actions across are in that state.

¹²Three states that stand out are Kansas, Maine and New Jersey. Kansas is outside the two regions with the highest rates of fraud. Maine and especially New Jersey standout as the states with the two states with the highest rate of fraud.

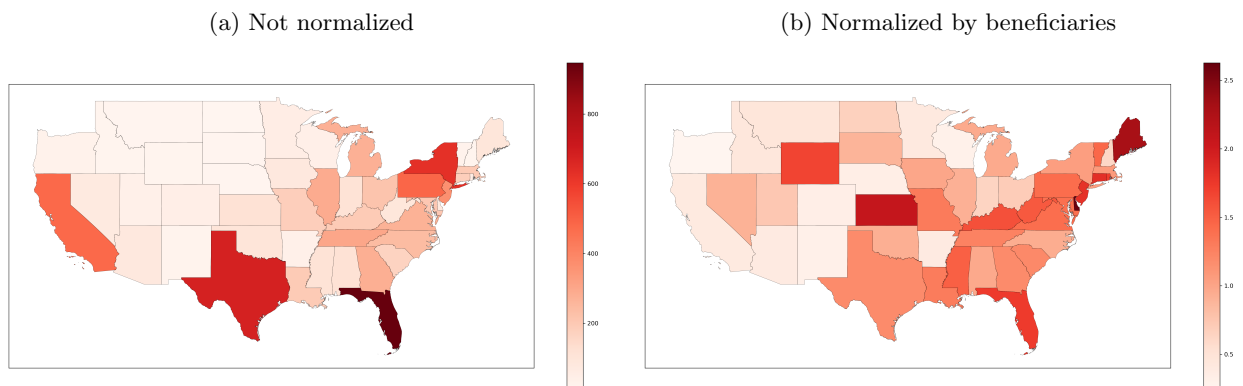


Figure 2: Geography of reported healthcare fraud cases in DOJ press releases, normalized by number of Medicare or Medicaid beneficiaries in each state

medical specialty. Table 5 shows the distribution of press releases among the 10 medical specialties most likely to be named in a press release. Pain is by far the more frequent specialty, which may relate to the extensive prosecution of opioid-related cases. Other top categories include emergency medicine, which has been explored by Leder-Luis (2025) and Dafny (2005), and hospice, which received attention in Gruber et al. (2025).

Table 5: Top 10 types of medical care named in DOJ press releases

Type of Medicine	Frequency
Pain	658
Emergency	223
Hospice	163
Psychiatry	163
Genetic	161
Cardiology	129
Oncology	99
Vascular	99
Urology	97
Pediatric	79

4.3 Types of Fraud

Next, we examine the types of fraud from these press releases. Upcoding and medical necessity cases are a larger fraction of prosecuted frauds than substandard care. Figure 3 plots the fraction of press releases in each of 4 categories of fraud: upcoding, substandard care, medical necessity frauds, and kickbacks. Note that percentages do not add up to 100% because a given press release can have multiple types of fraud. Upcoding and medical necessity fraud rates track one another because roughly 65.6% of medical necessity cases are also cases involving upcoding. The low rate of substandard care is not entirely due to low levels of sensitivity in our LLM categorization of such cases. In our randomly-selected test sample, where we manually code substandard care, such care constituted only about 18.6% of cases. Finally, Figure 3 also shows the proportion of press releases that involve a kickback.

Because types of frauds are not mutually exclusive, a given case can include, e.g., upcoding and medical necessity fraud, or medical necessity fraud and also kickbacks. Figure 4 shows the rate of overlap of different types of fraud in the same press relates. Specifically, each cell describes what fraction of the press releases concerning the fraud indicated in the row header also mention the the fraud indicated in column header. Overlap is greatest for medical necessity cases, which constitute roughly 60 to 80% of press releases discussing other frauds. Overlap is lowest for kickbacks, which constitute between 25 and 35% of the other three types of fraud.

Our data also shed light on the top 5 specialties most commonly mentioned in fraud press releases over




Figure 3: Categories of healthcare Fraud from DOJ Press Releases


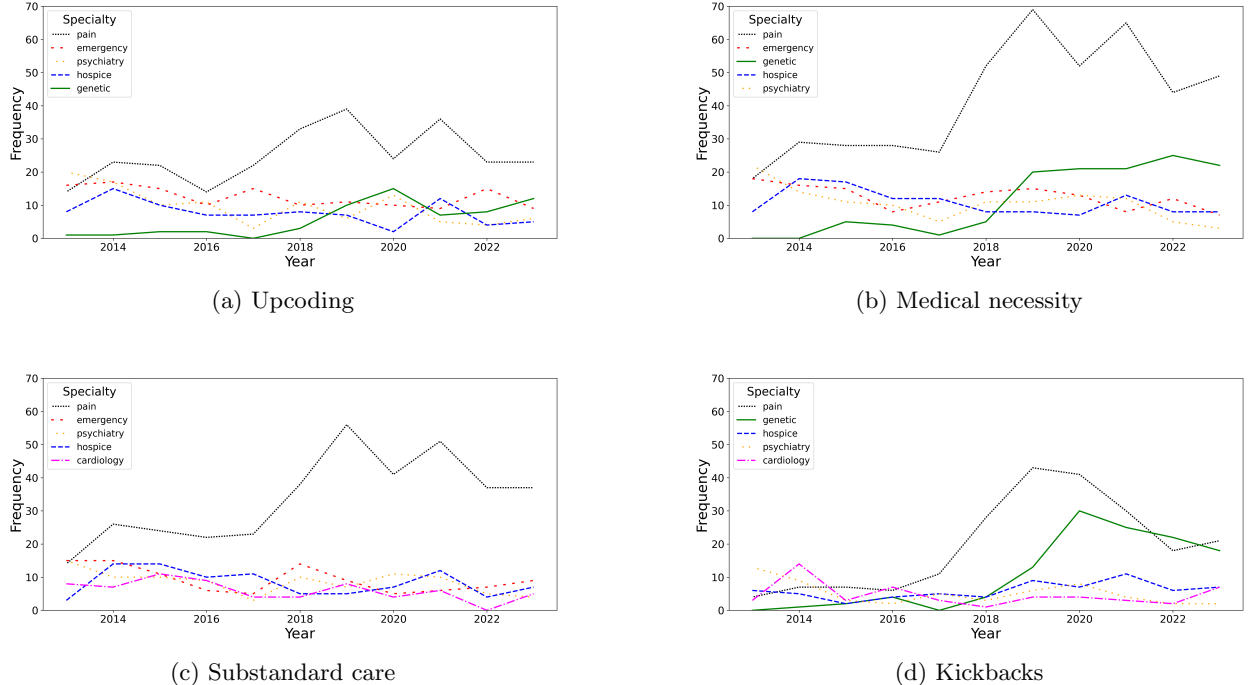


Figure 4: Overlap between different types of fraud discussed in DOJ Press Releases

time (Figure 5). Pain, the most common specialty represented in fraud press releases, is also the most common specialty cited within each category of fraud. However, it is most common in medical necessity cases. Moreover, pain-related fraud press releases peaked during the second half of our sample, even as other

data indicate that prescription rates for opioids were declining (Wang et al., 2024). Genetics is probably the second most common specialty represented in fraud press releases. It is most common in medical necessity and kickback cases. Like pain, genetics-related cases peak in the second half of our sample, reflecting recent enforcement priorities by the DOJ (Sun et al., 2020). The remainder of the top 5 specialties, including hospice, emergency, and psychiatry were more common in the first half of the sample, and do not stand out in any particular class of fraud press release.

Figure 5: Specialties involved in healthcare fraud cases reported in DOJ press releases, by type of fraud and year.



Note: Units are number of press releases. Data based on DOJ press releases from 2013-2023. Specialties are determined using AAMC classification.

Overall, our empirical results allow us to understand the relationship between our model and the enforcement of real-world fraud, highlighting the types of medicine and geographic regions most likely to experience health care fraud. This also provides opportunities for further research, highlighting gaps such as psychiatry fraud, which has not to our knowledge been addressed by the health economics literature.

5 Conclusion

Health care fraud comprises a variety of behaviors with vastly different economic consequences. In this paper, we develop a new economic definition of health care fraud that separates these behaviors, enabling

a clearer analysis of the welfare concerns associated with different kinds of fraud. We identify three major types of fraud: upcoding, medical necessity fraud, and substandard care. These frauds arise from different conditions and have distinct interpretations for efficiency. We relate these categories of health care fraud to a growing literature on the topic, as well as to numerous examples from all areas of medicine.

We connect our model to the incentives under which providers conduct health care, as well as the incentives of enforcement agencies to combat fraud. For providers to produce health care without fraud, there are a set of constraints that we describe, including truthful reporting, incentive compatibility, and individual rationality. The failures of these conditions correspond to our different types of fraud. Moreover, we discuss the incentives and limitations on anti-fraud enforcement, notably the budgets of the enforcement agencies.

Using data from the Department of Justice, we provide new empirical evidence on the relative prevalence of our different frauds, as well as their relationship to geography and medical specialty. Medical necessity fraud is the most prevalent, and the Eastern United States sees the highest rates of fraud per beneficiary. While health care fraud has touched nearly every area of medicine, some areas show particularly high frequency, including pain medicine and genetics.

Overall, this paper outlines a cohesive way of thinking about health care fraud, and also highlights avenues for future research. The conditions that allow different frauds to arise also provide potential avenues for anti-fraud tools, and also indicate which tools may be more successful under different circumstances. Moreover, changes in insurance design, beyond the scope of this paper, may help mitigate incentives to commit fraud.

References

- Alt, J. E. and D. D. Lassen (2012, 11). Enforcement and public corruption: Evidence from the american states. *The Journal of Law, Economics, and Organization* 30(2), 306–338.
- American Academy of Medical Coders (2024). Specialty profiles.
- Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *The American Economic Review* 53(5), 941–973.
- Baicker, K., E. S. Fisher, and A. Chandra (2007). Malpractice liability costs and the practice of medicine in the medicare program. *Health Affairs* 26(3), 841–852. doi: 10.1377/hlthaff.26.3.841.
- Becker, D., D. Kessler, and M. McClellan (2005). Detecting medicare abuse. *Journal of Health Economics* 24(1), 189–210.
- Brennan Troyen, A., L. Leape Lucian, M. Laird Nan, L. Hebert, A. R. Localio, G. Lawthers Ann, P. Newhouse Joseph, C. Weiler Paul, and H. Hiatt Howard (1991). Incidence of adverse events and negligence in hospitalized patients. *New England Journal of Medicine* 324(6), 370–376. doi: 10.1056/NEJM199102073240604.
- Chandra, A. and D. Staiger (2007). Productivity spillovers in healthcare: evidence from the treatment of heart attacks. *The Journal of Political Economy* 115, 103.
- Dafny, L. S. (2005). How do Hospitals Respond to Price Changes? *American Economic Review* 95(5), 1525–1547.
- Danzon, P. (1984). The frequency and severity of medical malpractice claims. *The Journal of Law and Economics* 27(1), 115–148. doi: 10.1086/467060.
- Danzon, P. M. (2000). *Chapter 26 Liability for medical malpractice*, Volume 1, pp. 1339–1404. Elsevier.
- Diwan, R., P. J. Eliason, R. J. League, J. Leder-Luis, R. C. McDevitt, and J. W. Roberts (2025). Competition and fraud in health care. *Work in Progress*.
- Drake, Tim Kanu, A. S. N. (2013). American criminal law review. *American Criminal Law Review* 50, 1131.
- Einav, L., A. Finkelstein, Y. Ji, and N. Mahoney (2025). Targeting overuse of home health care: Evidence from multiple policy instruments. *Work in Progress*.
- Eliason, P. J., R. J. League, J. Leder-Luis, R. C. McDevitt, and J. W. Roberts (2025). Ambulance taxis: The impact of regulation and litigation on health care fraud. *Journal of Political Economy* *Forthcoming*.

- Engstrom, D. F. (2012). Harnessing the private attorney general: Evidence from qui tam litigation. *Columbia Law Review* 112(6), 1244–1325.
- Fang, H. and Q. Gong (2017, February). Detecting potential overbilling in medicare reimbursement via hours worked. *American Economic Review* 107(2), 562–91.
- Farber, H. S. and M. J. White (1991). Medical malpractice: an empirical examination of the litigation process. *Rand J Econ* 22(2), 199–217. Farber, H S White, M J Journal Article Research Support, Non-U.S. Gov't Research Support, U.S. Gov't, Non-P.H.S. United States 1992/02/03 Rand J Econ. 1991 Summer;22(2):199-217.
- Feldman, R. (2001). An economic explanation for fraud and abuse in public medical care programs. *The Journal of Legal Studies* 30(S2), 569–577.
- Frakes, M. (2013). The impact of medical liability standards on regional variations in physician behavior: Evidence from the adoption of national-standard rules. *American Economic Review* 103(1), 257–76.
- Galizzi, M. M., G. Godager, J. Li, I. Linnosmaa, T. Tammi, and D. Wiesen (2023). *Economics of healthcare provider altruism*, pp. 1–23. Springer.
- Gandhi, A. and A. Olenski (2024, March). Tunneling and hidden profits in health care. Working Paper 32258, National Bureau of Economic Research.
- Gruber, J., D. H. Howard, J. Leder-Luis, and T. L. Caputi (2025, January). Dying or lying? for-profit hospices and end-of-life care. *American Economic Review* 115(1), 263–94.
- Gupta, A., A. La Forgia, and A. Sacarny (2024, June). Turbocharging profits? contract gaming and revenue allocation in healthcare. Working Paper 32564, National Bureau of Economic Research.
- Hackmann, M. B. (2019, May). Incentivizing better quality of care: The role of medicaid and competition in the nursing home industry. *American Economic Review* 109(5), 1684–1716.
- Holmström, B. (1979). Moral hazard and observability. *The Bell Journal of Economics* 10(1), 74–91.
- Howard, D. H. and I. McCarthy (2021). Deterrence effects of antifraud and abuse enforcement in health care. *Journal of Health Economics* 75, 102405.
- Hyman, D. A. (2001). Health care fraud and abuse: Market change, social norms, and the trust “reposed in the workmen”. *The Journal of Legal Studies* 30(S2), 531–567.

- Hyman, D. A., C. Silver, B. Black, and M. Paik (2015). Does tort reform affect physician supply? evidence from texas. *International Review of Law and Economics* 42, 203–218.
- Jack, W. (2005). Purchasing health care services from providers with unknown altruism. *Journal of Health Economics* 24(1), 73–93.
- Joiner, K. A., J. Lin, and J. Pantano (2024). Upcoding in medicare: Where does it matter most? *Health Economics Review* 14.
- Kessler, D. and M. McClellan (1996). Do doctors practice defensive medicine?*. *The Quarterly Journal of Economics* 111(2), 353–390.
- Kessler, D. P. (2011). Evaluating the medical malpractice system and options for reform. *Journal of Economic Perspectives* 25(2), 93–110.
- Klick, J. and T. Stratmann (2007). Medical malpractice reform and physicians in high-risk specialties. *The Journal of Legal Studies* 36(S2), S121–S142. doi: 10.1086/520416.
- Kumaraswamy, N., M. K. Markey, T. Ekin, J. C. Barner, and K. Rascati (2022). Healthcare fraud data mining methods: A look back and look ahead. *Perspectives in Health Information Management* 19(1).
- Lakdawalla, D. N. and S. A. Seabury (2012). The welfare effects of medical malpractice liability. *International Review of Law and Economics* 32(4), 356–369.
- Leder-Luis, J. (2025). Can whistleblowers root out public expenditure fraud? evidence from medicare. *Review of Economics and Statistics* Forthcoming.
- Maclean, J. C., J. Mallatt, C. J. Ruhm, and K. Simon (2020, November). Economic studies on the opioid crisis: A review. Working Paper 28067, National Bureau of Economic Research.
- Malani, A. and J. Reif (2015). Interpreting pre-trends as anticipation: Impact on estimated treatment effects from tort reform. *Journal of Public Economics* 124, 1–17.
- Matsa, D. A. (2007). Does malpractice liability keep the doctor away? evidence from tort reform damage caps. *The Journal of Legal Studies* 36(S2), S143–S182. doi: 10.1086/519466.
- Mello, M. M., A. Chandra, A. A. Gawande, and D. M. Studdert (2010). National costs of the medical liability system. *Health affairs* 29(9), 1569–1577.
- Mushinski, D., S. Zahran, and A. Frazier (2022). Physician behaviour, malpractice risk and defensive medicine: an investigation of cesarean deliveries. *Health Economics, Policy and Law* 17(3), 247–265.

- Nguyen, T. and V. Perez (2020). Privatizing plaintiffs: How medicaid, the false claims act, and decentralized fraud detection affect public fraud enforcement efforts. *Journal of Risk and Insurance* 87(4), 1063–1091.
- Nyman, J. (1999). The value of health insurance: the access motive. *Journal of Health Economics* 18(2), 141–152.
- O’Malley, A. J., T. A. Bubolz, and J. S. Skinner (2023). The diffusion of health care fraud: A bipartite network analysis. *Social Science Medicine* 327, 115927.
- Perez, V. and C. Wing (2019). Should we do more to police medicaid fraud? evidence on the intended and unintended consequences of expanded enforcement. *American Journal of Health Economics* 5(4), 481–508.
- Picard, P. (2013). Economic analysis of insurance fraud. *Handbook of insurance*, 349–395.
- Rai, A. K. (2001). Health care fraud and abuse: A tale of behavior induced by payment structure. *The Journal of Legal Studies* 30(S2), 579–587.
- Rose-Ackerman, S. (1996). Altruism, nonprofits, and economic theory. *Journal of economic literature* 34(2), 701–728.
- Sanghavi, P., A. B. Jena, J. P. Newhouse, and A. M. Zaslavsky (2021). Identifying outlier patterns of inconsistent ambulance billing in medicare. *Health Services Research* 56(2), 188–192.
- Shekhar, S., J. Leder-Luis, and L. Akoglu (2023, February). Unsupervised machine learning for explainable health care fraud detection. Working Paper 30946, National Bureau of Economic Research.
- Shi, M. (2023, 09). Monitoring for waste: Evidence from medicare audits*. *The Quarterly Journal of Economics* 139(2), 993–1049.
- Silverman, E. and J. Skinner (2004). Medicare Upcoding and Hospital Ownership. *Journal of Health Economics* 23(2), 369–389.
- Studdert, D. M., E. J. Thomas, H. R. Burstin, B. I. W. Zbar, E. J. Orav, and T. A. Brennan (2000). Negligent care and malpractice claiming behavior in utah and colorado. *Medical Care* 38(3).
- Studdert David, M., M. Mello Michelle, A. Gawande Atul, K. Gandhi Tejal, A. Kachalia, C. Yoon, L. Puopolo Ann, and A. Brennan Troyen (2006). Claims, errors, and compensation payments in medical malpractice litigation. *New England Journal of Medicine* 354(19), 2024–2033. doi: 10.1056/NEJMsa054479.

Sun, R., M. Whirl-Carrillo, A. Kurian, and T. Hernandez-Bouchard (2020). Genetic testing fraud waste and abuse.

United States Government Accountability Office (2024). Medicare and medicaid: Additional actions needed to enhance program integrity and save billions.

U.S. Department of Justice, Office of Public Affairs (2000). Hca - the health care company and subsidiaries – to pay \$840 million in criminal fines and civil damages and penalties, largest government fraud settlement in u.s. history.

U.S. Department of Justice, Office of Public Affairs (2004, May). Warner-lambert to pay \$430 million to resolve criminal & civil health care liability relating to off-label promotion.

U.S. Department of Justice, Office of Public Affairs (2012). Houston-area doctor convicted in \$17.3 million medicare fraud scheme involving fraudulent claims for home health care services.

U.S. Department of Justice, Office of Public Affairs (2019a, September). Federal law enforcement action involving fraudulent genetic testing results in charges against 35 individuals responsible for over \$2.1 billion in losses in one of the largest health care fraud schemes ever charged.

U.S. Department of Justice, Office of Public Affairs (2019b, April). South florida health care facility owner convicted for role in largest health care fraud scheme ever charged by the department of justice, involving \$1.3 billion in fraudulent claims.

U.S. Department of Justice, Office of Public Affairs (2021, Feb). Justice department’s false claims act settlements and judgments exceed \$5.6 billion in fiscal year 2021.

U.S. Department of Justice, Office of Public Affairs (2023). Missouri laboratory owners agree to pay \$1.9 million and relinquish \$7 million in escrow in settlement of civil fraud claims.

Wang, S., M. E. Rossheim, S. T. Walters, R. R. Nandy, and K. Norheim (2024). Geographic trends in overall and long-acting opioid prescriptions under medicaid and medicare part d in the united states, 2013-2021. *Am J Drug Alcohol Abuse* 50(5), 690–702.

Appendices

A Legal Standards for Health Care Fraud Liability

Medicare and other government health programs cover care that is “reasonable and necessary”. For Medicare, what is reasonable and necessary is determined by Medicare requirements for billing (Medicare Act, 42 U.S.C. § 1395y(a)(1)(A)). These requirements—often set forth in local coverage determinations by Medicare contractors—are a combination of objective and subjective criteria. Objective criteria are, e.g., that the provider be a physician, that a lab test be conducted, the patient be admitted for two nights to bill for an inpatient hospitalization, etc. Subjective criteria involve a provider’s professional judgment, e.g., that a scan shows the patient has arterial blockage or a patient is likely to die in 6 months. Like Medicare, other payors may impose conditions on when care may be reimbursed. Such conditions, whether from Medicare or other payor, can be the basis for a fraud suit.

A common feature of civil and criminal fraud statutes is that the plaintiff must show that the defendant, among other things, (a) filed a claim that was false and (b) knew the claim was false (*United States v. Aseracare, et al.*, 938 F.3d 1278 (11th Cir. 2019)). Showing (a) a claim is false means showing that the provider did not meet Medicare’s requirement that care be reasonable and necessary. What is required to show (b) knowledge varies depending on whether the Medicare’s requirements are objective or subjective and whether the case is civil or criminal. For objective criteria, all that is required is that the provider knew that they did not meet that requirement, even if they thought that the care, in their good faith subjective judgment, was appropriate.¹³ For subjective criteria, courts will defer more to the provider’s good faith judgment. However, plaintiffs have leeway to show that the provider lied about her beliefs concerning appropriate care, for example, because she did not review a patient’s medical record. Finally, in civil cases, knowledge can be demonstrated by showing reckless disregard for facts (False Claims Act, § 3729(b)(1)(A)(iii)) and a provider’s lack of knowledge about requirements does not disprove (b) (*United States ex rel. Schutte v. Supervalu Inc.*, 143 S.Ct 1391, 1400 (2023)). In short, there are circumstances when the provider can be liable for fraud even if she did not intend to provide inappropriate care or even affirmatively know she did.

¹³Indeed, when Medicare requires the provider certify certain facts are true and the provider does not know those facts are true, a jury can find the provider knew that she did not know the facts and thus knew she did not meet Medicare requirements. *Universal Health Services, Inc. v. United States ex rel. Escobar*, 136 S.Ct 1989 (2016) (discussing “implied false certification theory”).

B Model: Incentives to Commit Fraud

More formally, our model of insurance fraud unfolds in the context of an interaction between a patient—who has insurance from a payor—and a provider drawn from the population. This model builds on previous models of fraud by a beneficiary against an insurance company (Picard, 2013; Holmström, 1979) by replacing the beneficiary with the medical provider. In our model, the patient is insured by a payor that pays the provider directly for services delivered to the patient. We assume that there is no information asymmetry or misalignment of interests between the payor and patient, patients are homogeneous, that there are no externalities from medical care, that insurance serves only to smooth patient consumption, and that insurance is actuarially fair. We will also assume, for the moment, that healthcare providers are homogeneous and the payor knows the provider’s cost of each treatment.¹⁴ This will allow us to focus on the payor’s optimal contract with the provider.

We define optimal care to be care that maximizes the joint surplus of the patient and provider. The surplus from providing care y when the patient has condition x is $S(y, x) := v(y, x) - c(y, x)$, where $v(y, x)$ is the fully-informed patient’s willingness to pay for treatment y given her health condition x and $c(y, x)$ is the provider’s cost of providing treatment y given x . The socially optimal care t is the treatment y for which the patient has largest, positive surplus: $t(x) = \arg \max_{y: S(y, x) > 0} S(y|x)$.

To motivate our model, we imagine the following sequence of events. First, a patient visits a provider. Neither she nor the payor knows her health condition.¹⁵ Second, the provider diagnoses the patient with a specific condition x and provides a treatment y at a cost (to the doctor) of $c(y, x)$.¹⁶ Third, the provider files a claim with the payor stating that the provider delivered service \tilde{y} , potentially different than y .¹⁷ Fourth, the payor potentially conducts an audit $a(\tilde{y}) \in [0, 1]$ of the provider’s claim at cost e . (In the real world, the audit could be a utilization review or litigation for fraud.) We assume the audit is perfectly informative, revealing the patient’s condition x and actual care y provided. If the audit does not confirm the provider’s claim, the provider pays a penalty of F . Fifth, the payor reimburses the provider an amount $p(\tilde{y}, a(\tilde{y}) \cdot y)$, which may depend on both the provider’s claim and the results of the optional audit.

¹⁴Homogeneous patients allow us to ignore adverse selection in insurance market; homogeneous providers allow us to ignore variation in provider production functions; information parity allows us to ignore patient moral hazard vis-a-vis payors; insurance for smoothing allows us to ignore insurance that enables patients to afford care that costs more than lifetime income (Nyman, 1999); and fair insurance allows us to ignore incomplete insurance.

¹⁵More realistically, we could assume the payor knows some aspects of the patient’s condition x , but not all of it, and not as much as the provider will learn. Qualitatively, that will not change our model. The incentive compatibility condition will require some aspects of x to be unobserved outside an audit; otherwise, the problem becomes trivial and the payor simply pays $c(t(x), x)$ if the provider delivers optimal care $t(x)$, and nothing otherwise. In the main, we do not observe such trivial contracts in the real world.

¹⁶It is assumed that the patient does not fully know or understand what y was provided, and so cannot simply reveal that to the payor.

¹⁷Here we present a model where the provider only reports the care she conducted. However, it is possible to write an isomorphic model where (i) the provider reports both the care delivered and some aspects of the patient’s condition, but not the patient’s entire condition, or (ii) the provider reports the patient’s complete condition and some aspects of care, but not all the care provided.

The payor’s problem is to choose (a) an audit policy $a(\tilde{y})$ and (b) a reimbursement schedule $p(\tilde{y}, a(\tilde{y}) \cdot y)$ before the patient visits a provider so as to maximize expected joint surplus net of audit costs, $\int_X [v(y, x) - c(y, x) - a(y)e - P]dG(x)$, where $G(x)$ is the distribution of the patient’s health state and P is the patient’s insurance premium, subject to four constraints. The first is a truthful reporting (TR) constraint that ensures that the provider honestly reports to the payor what care she delivered:

$$p(y) - c(y, x) \geq p(y') - c(y, x) - a(y')F \iff p(y) + a(y')F \geq p(y') \quad (1)$$

for all $x' \in X$, $y' \in Y$. Given the probability of an audit and the fine for failing the audit, this constraint says the provider prefers to report the care she actually provided rather than report any other level of care. The second is an incentive compatibility constrain (ICC) that ensures the provider delivers optimal care $t(x)$:

$$p(t(x)) - c(t(x), x) \geq p(y') - c(y', x) \quad (2)$$

for all $y' \in Y$. Given the patient’s condition x , the provider incremental reimbursement from providing $t(x)$ exceeds her cost savings from delivering any other care y' . The third constraint is a provider individual rationality (IR) constraint that ensures that the provider’s expected utility is greater than her outside option, so she is willing to serve the patient in the first place. The final constraint is a non-zero profit-condition wherein the patient pays a premium that ensures that the insurance company’s expected profits are non-negative.¹⁸

The three types of fraud in our framework—upcoding, substandard care, and medically unnecessary care—can each be understood as failure to satisfy one of these constraints for some states of the world, i.e., some patient conditions. In this sense, fraud is a failure in the design of insurance to account for incentives to distort or misreport care in a way that is optimal, perhaps due to insufficient penalty for failing an audit or to prices creating bad incentives in cases that are not audited.

Upcoding. Upcoding is a product of the failure of the insurance design to satisfy the truthful reporting constraint due to audit costs. As a preliminary matter, note that the incentive compatibility constraint requires a higher cost treatment y_H be reimbursed at a higher rate p_H than a lower-cost treatment y_L , otherwise providers will not deliver the higher cost treatment even when it increases joint surplus. This higher reimbursement creates an incentive for the provider to perform y_L but report that she did y_H to

¹⁸Our formulation is somewhat different than Picard (2013) in that we maximize patient surplus subject to a participation constraint on other actors, rather than maximizing provider utility subject to a participation constraint on other actors. The economics are the same, but in our model the surplus flows to the patient, while in the conventional model it flows to the provider.

increase her revenue by $p_H - p_L$. The audit constraint is intended to stop this upcoding. But audits have costs, so the payor may rationally choose to audit only a subset of claims. For audited procedures, the payor observes the patient condition and actual treatment, and can punish whenever the latter is not exactly $t(x)$. But for non-audited claims, the payor must rely on the incentive compatibility constraint to induce the provider to deliver appropriate treatment.

This logic explains why there is upcoding. But a separate question is why one ever *observes* litigated cases with upcoding. If the payor only audits high-cost treatments, then the threat of audit for those treatments should deter upcoding involving those treatments. Moreover, upcoding of treatments that are not high-cost treatments would never be audited, and thus not observed in litigation. One reason why one might still observe upcoding in audits, however, is that the punishment for not reporting truthfully may be inadequate to deter some providers.¹⁹ For example, if some providers are judgment proof, then audits may not deter filing of high-cost claims. In other words, the Becker model of crime can explain why we observe upcoding: the expected punishment for the crime (upcoding) is smaller than the expected gain from the crime (higher reimbursement).²⁰

Substandard care and medically unnecessary care. These two frauds, in contrast to upcoding, are a product of failing to satisfy the incentive compatibility constraint in non-audited cases. As we explained, the incentive compatibility constraint requires that higher-cost treatments be reimbursed more than lower-cost treatments. But it also requires that the profits from appropriate care, $t(x)$, for a condition x must be weakly greater than the profits from inappropriate care, $y \neq t(x)$, for that condition. If the payor observed x , the payment for optimal care could be set strictly greater than any other care. But when x is unobservable, the only way the ICC is satisfied for multiple conditions is that appropriate procedures across conditions each yield the provider roughly the same net revenue. When combined with an assumption that the provider will do what the payor recommends when the provider is indifferent or that the provider is altruistic, then an equal-profit reimbursement schedule will ensure the provider chooses appropriate treatment for each patient. Note that an equal-profit schedule is not the same as an equal-price schedule, because costs of different procedures may vary.

If equal-profit reimbursement schedules eliminate providers' incentives to commit fraud, why do we observe examples of substandard care or medically unnecessary care in litigation? Presently, our model cannot explain inappropriate care because we assume that all providers have identical cost for any given treatment and that the provider knows these costs. The payor can deliver identical profit reimbursements as

¹⁹Of course, mistakes in crafting an audit policy may also reveal fraud. If the payor employs an audit policy wherein the probability of litigation is positive and increasing in size of claim, then it is possible the payor sets the probability too low from some claims. In that case, it would be rational for the provider to upcode and it will occasionally be discovered.

²⁰A fine goes further to deter a provider if she is risk-averse, as our model assumes, than if she is risk-neutral.

long as they know these costs. But if the provider does not know the cost of a treatment, or if provider have different costs for any given treatment and the payor cannot have provider-specific reimbursement schedules, then substandard and medically unnecessary care can emerge as a result from mis-pricing.

The first condition that might generate fraud—lack of knowledge about costs—is likely to hold in the modern, complex healthcare system. Payors have to set prices for nearly 10,000 procedures in the Healthcare Common Procedure Coding System (HCPCS). It is reasonable to suppose that they might not know provider costs for every single one of these procedures. If they pay too much for a procedure so that it offers a higher net profit, then they will incentivize doctors to to perform that procedure even when it is not necessary, i.e., unnecessary care fraud. If they pay too little for a procedure, then they will incentivize substandard care because the procedure offers too little profit even when it is appropriate.

A second condition that could generate medically unnecessary or substandard care, that providers have different costs for the same procedure, is also likely, especially in a sector where there is both learning by doing and specialization (Chandra and Staiger, 2007). Consider, for example, the case where providers have heterogeneous cost functions, but payors know only the the distribution of cost functions, so the payor must employ a common reimbursement scheme. The payor sets profits for appropriate care weakly greater than profits for other care, given *expected* costs of providers:

$$p(t(x)) - E_{\theta}[c(t(x), x; \theta)] \geq p(y') - E_{\theta}[c(y', x; \theta)] \quad (3)$$

where θ is a cost-shifter that varies across providers. Consider a hypothetical provider with cost $c(y, x; \theta) \neq E_{\theta}(y, x; \theta)$. If the difference between a provider's cost of $t(x)$ and some alternative treatment y' is sufficiently greater than the difference between the expected cost of each across providers (specifically, if $c(t(x), x; \theta) - c(y', x; \theta) > (E_{\theta}[c(t(x), x; \theta)] - E_{\theta}[c(y', x; \theta)]) + Z$, where $Z = p(t(x)) - E_{\theta}[c(t(x), x; \theta)] - p(y') - E_{\theta}[c(y', x; \theta)]$), then provider will have an incentive to provide y' rather than $t(x)$:²¹

$$p(t(x)) - c(t(x), x; \theta) < [p(y') - c(y', x; \theta)] \quad (4)$$

If the alternative treatment y' is of lower intensity than $t(x)$, then the hypothetical provider will want to deliver substandard care. If y' is higher intensity than optimal care, then the provider will deliver medically

²¹To see this, note that Eq. (3) implies $Z > 0$. The incremental gain to the hypothetical provider delivering $t(x)$ is

$$\begin{aligned} p(t(x)) - c(t(x), x; \theta) - [p(y') - c(y', x; \theta)] &= p(t(x)) - E_{\theta}[c(t(x), x; \theta)] - (p(y') - E_{\theta}[c(y', x; \theta)]) \\ &\quad - (c(t(x), x) - E_{\theta}[c(t(x), x; \theta)]) + (c(y', x) - E_{\theta}[c(y', x; \theta)]) \\ &= Z + (c(y', x) - E_{\theta}[c(y', x; \theta)]) - (c(t(x), x) - E_{\theta}[c(t(x), x; \theta)]) < 0 \end{aligned}$$

where the last inequality comes from our condition on the difference in costs between optimal and alternative care for the hypothetical provider.

unnecessary care.²²

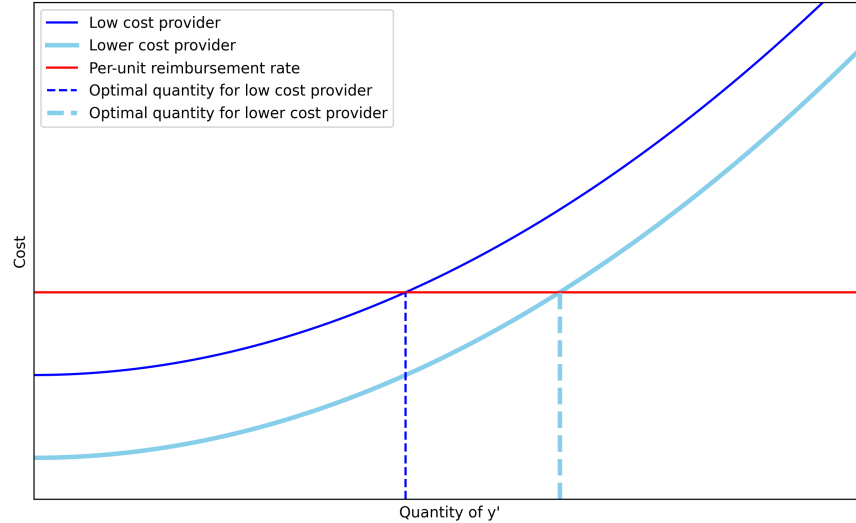


Figure A1: Lower cost providers have higher volumes (given adequate supply of patients)

Unknown or heterogeneous costs can explain why inappropriate care is delivered, but what explains why such care is observed in litigation, i.e., via audits? The answer is that, if the payor knows that some providers have an incentive to commit fraud because they have low costs, then it may conduct audits not just against high-value claims, but also low-cost claims. But how can a payor identify low-cost providers if it does not know their costs? One way is to use quantity. If providers have convex costs, and low-cost providers of a given treatment have a vertically down-shifted cost function, then the low-cost providers will deliver a higher volume of procedures (if they have the patient flow to do so), as illustrated in Figure A1. Thus, our rudimentary model of fraud also predicts (1) providers with lower costs for alternative care will commit more substandard or medically unnecessary care fraud, and (2) amongst those providers, the ones with the greatest volume will face litigation and some be found to have provided substandard or medically unnecessary care.

Altruism. So far we have ignored altruism, an important driver of provider behavior in the healthcare sector (Rose-Ackerman, 1996; Galizzi et al., 2023). Altruism can be defined as an internal payoff to providers for delivering appropriate care. This can introduce slack into the incentive compatibility constraint, reducing the probability of substandard or unnecessary care.²³ Of course, providers also care about their own income and there is heterogeneity in altruism (Jack, 2005), so the mere existence of altruism does not eliminate all provider moral hazard.

²²An interesting implication of this model is that medical malpractice, which is defined as providing substandard care, may also be the result of a payor's inability to observe individual physicians costs.

²³Likewise, if physicians face a psychic cost from lying, then this psychic cost can reduce the rate of auditing required to achieve any given level of upcoding.

Efficient fraud. Finally, our model can also be embellished to create a distinction between fraud and efficiency by introducing incomplete information. Suppose that the patient’s true condition is x , but audits by the payor reveal $x_A = x + e_A$ and examinations by the provider reveal $x_P = x + e_P$, where errors are mean zero. Suppose that the provider – in compliance with its contract with the payor – performs $t(x_P)$ and reports $(x_P, t(x_P))$. It is possible that an audit reveals a different condition $x_A \neq x_P$ and so optimal care differs in the minds of the payor and the provider. In this case, the provider will be liable for substandard or unnecessary care, which are defined in reference to what the payor thinks is appropriate care. (After all, contractually the payor exchanging money for the service that it wants, not what the social planner wants.) However, the parties measurement errors could be such that the surplus from the provider’s mis-measurement is greater than that under the payor’s mis-measurement: $S(t(x_P), x_P) > S(t(x_A), x_A)$. When that happens, we can conclude that fraud was relatively efficient. Of course, the opposite could also happen, and fraud could be inefficient. The lesson is that imperfect measurement of health condition could result in efficient fraud, even under the optimal, second-best contract.

Kickbacks. Although our framework does not focus on kickbacks, a large share of fraud cases fall into that category. Our sketch model can be modified to explain kickbacks. One can define a kickback in our model as a situation where the provider splits its surplus from providing and reporting (x, y) with either other providers or with patients. Splitting profits with other providers, often referred to as providing referral fees, can be modeled as care provided by a group of providers. In the case whether the group does not collude, the payor only has to ensure that each group member has a profit incentive to provide their specific component of the optimal care $t(x)$. In the case where they can collude, the payor has to impose the additional constraint that the sum of profits across group members for providing $t(x)$ is weakly greater than for any other treatment. Anti-kickback laws mitigate the need for that extra constraint. Or, if we consider the law as part of the optimal contract, then the anti-kickback law satisfies that extra constraint.

Splitting profits with patients also qualifies as a kickback. And this is made possible if it is the case that $p(y', x') - c(y', x') > v(t(x), x) - v(y', x)$ for $y' \neq t(x)$, i.e., the patient is compensated for her worse outcomes when the provider does not administer optimal care given the patient’s true condition x . This is most easily achieved if optimal care is $t(x) = 0$ and the patient has free disposal of y' , i.e., the patient can choose not to use the prescribed drug or device.

To summarize, our model describes the relationship between audit effectiveness, reimbursement rates, and the different categories of fraud we identify. Imperfect audits lead to upcoding. Medical necessity fraud is motivated by relatively high reimbursement rates that create excess profit opportunities for intensive but inappropriate care, while substandard care is a product of relatively inadequate reimbursements for intensive but appropriate care.

C Model: Incentives to Enforce Against Fraud

In this section we explore the federal government’s incentives to combat fraud. Many organizations contribute to anti-fraud efforts; primary is the Department of Justice (DOJ), but investigatory support comes from the Office of the Inspector General of Health and Human Services (HHS-OIG), the Federal Bureau of Investigation, the Center for Program Integrity within the Centers for Medicare and Medicaid Services, and the state Medicaid fraud control units (MFCUs). We focus on the DOJ, which has primary law enforcement authority through the US legal system.

The DOJ’s ambit includes not just different kinds of healthcare frauds, but also fraud against other, non-healthcare federal programs, as well as illegal activity that does not concern fraud, e.g., antitrust violations, civil rights violations by the government, discrimination, etc. So, a natural way to think about the DOJ’s incentives to prosecute fraud is to balance its incentives to enforce different law given its resource constraints.

To simplify, we treat the DOJ as a unitary actor that obtains value from prosecuting different types of healthcare fraud cases $\mathbf{h} = (h_1, \dots, h_4)$ and other types of cases o : $V(\mathbf{h}, o; \mathbf{x})$. We assume agency preferences are well-behaved and different types of cases are substitutes, implying indifference curves with the usual curved shape. The marginal value from each type of case may be a function of other factors, \mathbf{x} , which can include the amount of fraud, the patient harm from fraud, the social welfare benefits of other cases, political priorities of the administration, and the political goals of the agency within the administration.

The agency has a budget, measured in financial resources, labor, and time, to expend on prosecuting cases. Each type of case has a cost, including resources to identify a case, gather evidence, and prosecute the case. While some costs may be common, such as filing a motion, others might vary by case type. For example, electronic medical records (EMR) reduce the cost of filing a healthcare fraud case, but do not affect civil rights cases. And a ruling that makes antitrust cases easier to pursue decreases the relative price of other cases relative to fraud cases. We can capture these cost considerations with a budget line: $\mathbf{p}_h \mathbf{h} + p_o o \leq M$, where M is the agency’s budget and \mathbf{p}_i is the cost of prosecuting healthcare cases and p_o the cost of other cases. We assume the budget is one dimensional (e.g., just money or just labor) to capture the intuition behind trade-offs.

In this toy model, the DOJ’s choice of number of different fraud and non-fraud cases to pursue is dictated by the tangency between the agency’s indifference curve and budget line. See Figure A2. Even in this rudimentary framework, we obtain some basic insights. Given the agency’s budget, a critical driver of the number of prosecutions for a given type of fraud is the cost of bringing such a fraud suit. And the cost of that suit is the opportunity cost of bringing other suits. Conversely, the marginal rate of substitution depends on the ratio of the marginal value of bringing a fraud versus another suit. It is relative, not absolute,

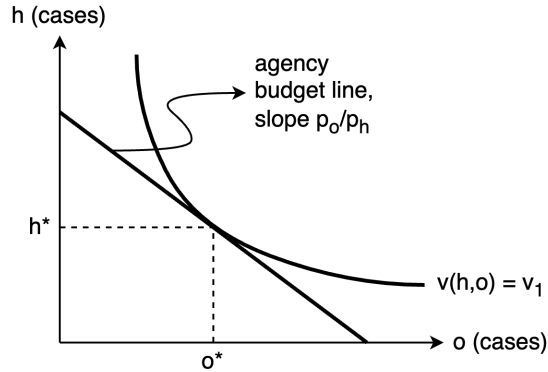


Figure A2: An agency’s optimal portfolio of fraud (h) and non-fraud (o) cases (assuming only 1 type of healthcare fraud case)

value that matters. This echoes a related empirical literature on U.S. corruption enforcement, which shows corruption enforcement responds to greater resources (Alt and Lassen, 2012)

Comparative statics reveal additional insights, which should be obvious conditional on our framing. Increasing the level or cost of a type of fraud, *ceteris paribus*, increases the relative marginal value from prosecuting fraud, and will increase prosecutions of that type of case. A political promise from the President to prosecute healthcare fraud has similar effects an increase in the relative marginal value of prosecuting that fraud. Likewise, increasing the welfare loss from another type of fraud or, e.g., the anti-competitive activities of technology companies might decrease the amount of prosecutions of a given type of fraud because of shifts in DOJ’s indifference curve. One must be careful not to overestimate the cross-case-type elasticity from one specific type of case to another. However, the elasticity from one case type to all other types of cases may be non-trivial.

Changing our focus to the cost side of the ledger, a favorable legal change or new investigative tool that makes all DOJ cases easier to prosecute will function like an increase in budget, and have income effects that increase all types of prosecutions. But a law or ruling, a budget earmark, or an investigative technology that is case-type specific will have the usual income and substitution effects. For example, a new law that makes it easier to prosecute kickbacks will have a substitution effect that reduces prosecutions of other frauds and non-fraud cases, but also a small income effect that pulls back that substitution effect to some degree.

D Details on Press Release Sample Construction

We gathered our sample of press releases in four steps. First, we scraped press releases from the Department of Justice (DOJ) and U.S. Attorneys Office (USAO) websites, specifically searching for press releases containing the words “Medicare” or “Medicaid”. We initially retrieved 12,984 press releases, but this includes duplicates

and press releases that mention Medicare or Medicaid, but do not concern fraud. We drop the duplicates first based on the URL of press releases, and then on the body text of press releases. We need both steps because some press releases might be duplicates but have different URLs (the same URL but with ‘-0’ at the end of it). Although the DOJ website contains data from 2009 to date, the USAO website contains data from 2013 to present. To keep the time period of our data consistent, we only keep press releases from 2013-2023. We are left with 6,909 press releases.

Second, we filter out press releases that are not related to fraud. We do this with a combination of exclusion and inclusion criteria based on search terms. To start, we drop any press releases containing the following terms antitrust, outreach campaign, awards ceremony, income tax in either the title or body text from our sample. This leaves us with 6,497 observations. We also exclude cases containing takedown, strike force in the title, because this indicates that a press release is about multiple cases and doesn’t focus on any specific fraud. This step leaves us with 6,374 press releases. To ensure press releases focus on specific legal actions we keep cases that contain the following terms sentenced, settlement, settle, convict, scheme, guilty, resolve, allegation, allegations, arrested, charg, pay, interven, false claims act, indict, doctor, physician, suit, sues, sued, lawsuit, complaint in the title. We create a separate data frame that does not include those phrases in the title. The data frame contains 410 press releases. We manually check these and add back 148 press releases which we concluded to be fraud cases. We are left with 6,112 press releases after this second step.

Third, we classified these 6,112 press releases by the type of fraud they contain using large language models. We describe this step in detail in the next section.

Fourth, after our categorization step, we found that 956 press releases were tagged as false for all 4 categories of fraud (upcoding, medical necessity, substandard care, and kickbacks). We were concerned that some of these press releases were not merely cases of fraud that were difficult to categorize, but cases with no fraud in them. To distinguish these two types of press releases, we drew a random sample of 95 of the 956 press releases, manually categorized them as discussing fraud or not, and trained a prompt to identify which of the 956 press releases pertained to fraud. To check the accuracy of our categorization of these 956 cases, we compared our manual coding of the 95 random sample cases with the LLM categorization of these cases, treating the manual categorization as a measure of “actual” frauds. We estimate that our LLM classification has a sensitivity of 94.12% and a specificity of 72.73%, suggesting that our LLM approach was better at catching actual frauds than eliminating non-frauds. As such, we may slightly overestimate the number of frauds, and thus underestimate the rate of each specific type of fraud.²⁴ Based on this fourth

²⁴Out of the 95 randomly-selected press releases we manually code to to train our classifier, 32 are fraud, i.e., 33.68%. However, given our estimates of sensitivity and specificity, we code 49.79% of the 956 cases as fraud. This estimate follows from a sensitivity of 94.12% on perhaps an actual fraud rate of 33.68% and 17.18% false positive rate on 63.31% that are possibly

step, we added back 507 of the 956 press releases back to our sample. This left us with a sample of 5,663 press releases.

E Classification by type of fraud

We use large language models to classify cases into 4 types of fraud: upcoding, substandard care, medically unnecessary care, and kickbacks. These classifications are not mutually exclusive, so one case can involve multiple categories of fraud.

E.1 Kickbacks

We use two different techniques for classification, depending on the nature of fraud. For kickbacks, we used text-tagging to identify cases. Any press release that contains the phrases {kickback, stark law, inducements} are labeled kickback cases. The phrase ‘stark’ is because of one of the anti-kickback statutes is called the Stark Law; that law prohibits physicians from referring patients to other providers if they have a financial interest in that provider. It is worth noting that we use substring matching for these phrases so the phrase ‘recruit patient’ will tag ‘to recruit patients’ even if it doesn’t contain the ‘s’ at the end. Generally, recruiting patients is an action wherein a healthcare professional offers patients economic benefits or bribes to undergo treatments or to allow their patient ids to be used in false billing.

We are able to judge the accuracy with which our text-tagging of kickbacks identifies actual kickback cases by seeing how it performs on a random sample of 500 cases we draw in order to train an LLM to categorize other types of fraud, a training process we explain in the next section. We manually characterized the type of fraud, including kickbacks, in those 500 cases. Treating our manual categorization as ground truth or actual fraud, we estimate that text-tagging has 92% sensitivity and 96.92% specificity, a very high rate of accuracy.

Recall that cases may have multiple frauds, so the kickback cases are not removed from the sample. Instead, cases with kickbacks are labeled as at least having a kickback, and included in the sample for classification for other types of fraud.

E.2 Upcoding, substandard care, and medical necessity fraud

To classify these three types of fraud we fine-tuned a prompt with a random sample of cases and DSPy, which is a prompt and weight optimization package²⁵. We then use our optimized prompt on ChatGPT,

not frauds. Without step four, we would have a denominator of 6,112. With step four, our denominator is 5,632. But if we extrapolated solely from our manually-coded sample where 33.68% of cases were fraud, our denominator would be 5,478 cases.

²⁵We use version 2.5.31 of DSPy, the latest available model at the time.

with a press release as context. To train our prompt, we randomly selected 500 case-specific press releases and hand-coded them into the 4 categories of fraud using the framework we set forth in Section 2.2. We used GPT 4o mini for our language model. Next, we converted our sample into DSPy.example objects with our title and body text as input, and columns for each fraud category are dummy variables labeled 1 for true and 0 for false. We used 100 observations to train our model, 150 to validate it and 250 to test it. We trained a separate prompt for each category of fraud. This allowed us to use different sample splits for each type of fraud, and to stratify the splits based on each category of fraud (e.g., for the upcoding sample split, the 500 press release sample is first stratified using the upcoding column, and then randomly assigned to the training, validation, and and testing subsamples). Our models consisted of a ChainOfThought module with each fraud model having a different prompt. We used a BootStrapFewShotWithRandomSearch optimizer to iterate through different combinations of labeled and few-shot demos to find the best model given our metric, which is exact match.

To repeat, DSPy generates different prompts for each type of fraud. The prompt for upcoding was:

You are a tool that reads text from the Department of Justice and classifies the text into categories of different types of conduct committed by the defendant. Each observation is one press release describing a lawsuit from the Department of Justice against a company or person that committed health care fraud. There are different kinds of health care fraud and the goal is to identify when cases match a specific type of fraud based on the text given. Upcoding is a type of health care fraud where a provider bills for more expensive care than they delivered, or bills for care that was not delivered at all. The goal is to classify text based on whether upcoding has been committed. Using the text, find whether a medical professional has billed for more expensive treatment than provided or billed for treatment that was not provided. If services are not eligible for reimbursement, but a provider bills for them anyway, this is also upcoding. If people are put in a higher reimbursement level than they should be at, then this is upcoding because the bill for the hours of treatment that they needed will be more expensive than it should be. If medically unnecessary care is performed and billed for, then this is not upcoding; it only becomes upcoding if they billed for services not performed or higher than the services that they actually performed. The manner in which the services were provided does not matter, only whether the services billed for were provided or not. A proper medical justification for the services provided is also irrelevant, as it only matters whether the services billed for were provided or not. Whether the services billed for were warranted based on medical need is irrelevant, as it only matters whether the services billed for match the services provided.

—

Follow the following format.

Text: text from department of justice press releases

Title: Title of the press release

Reasoning: Let's think step by step in order to produce the answer. We ...

Answer: Produce 1 if upcoding fraud was committed and 0 if not

—

[3 Labeled OR Bootstrapped examples inputs and output]

The optimized prompt for medical necessity fraud was:

You are a tool that reads text from the Department of Justice and classifies the text into categories of different types of conduct committed by the defendant. Each observation is one press release describing a lawsuit from the Department of Justice against a company or person that committed health care fraud. There are different kinds of health care fraud and the goal is to identify when cases match a specific type of fraud based on the text given. Medical necessity fraud is a type of health care fraud where a doctor performs care on a patient that does not qualify for the service due to insurance rules. Classify based on whether a medical provider has performed care that insurance has determined is medically unnecessary. If the document is unclear and the chances that medically unnecessary care was performed are between 0 and 1 then classify as true. If medically unnecessary equipment was provided, this also qualifies as medical necessity fraud.

—

Follow the following format.

Text: text from department of justice press releases

Title: Title of the press release

Reasoning: Let's think step by step in order to produce the answer. We ...

Answer: produce 1 for true 0 for false

—

[7 Labeled OR Bootstrapped examples with inputs and output]

The optimized prompt for substandard care was:

You are a tool that reads text from the Department of Justice and classifies the text into categories of different types of conduct committed by the defendant. Each observation is one press release describing a lawsuit from the Department of Justice against a company or person that committed health care fraud. There are different kinds of health care fraud and the goal is to identify when cases match a specific type of fraud based on the text given. Substandard Care fraud occurs when a health care provider bills for providing care to a patient that needed it according to the payer, but the provider performed lower-quality care. To be clear, the care is only substandard if there is a nonzero chance that patient actually needed the care that was billed for but not performed. If an unqualified, uncertified, unlicensed, or disbarred healthcare provider or company performs care that is supposed to be performed by someone more qualified, then that is also Substandard Care fraud. If a healthcare provider or company that is disqualified from performing care to patients, performs it, then this is also Substandard Care fraud. If a healthcare provider gives a patient a lower dose (including splitting doses) of medicine than the patient needs, this is Substandard Care fraud. Also, if a company misrepresents its software or equipment as being higher-quality than it actually is in order get certified in a certification that the software does not qualify for, then this is also Substandard Care fraud. If a drug is prescribed to someone for a condition that the drug is NOT FDA approved to treat (misbranded), then this is also Substandard Care fraud. If a provider uses or distributes for use devices that are defective or lower quality than required then this is also Substandard Care fraud. If a provider performs care that they are not qualified to perform, then this is Substandard Care fraud. Classify text based on whether Substandard Care fraud was committed.

—

Follow the following format.

Text: text from department of justice press releases

Title: Title of the press release

Reasoning: Let's think step by step in order to produce the answer. We ...

Answer: produce 1 for true 0 for false

—

[8 Labeled OR Bootstrapped examples with inputs and output]

Our use of a long prompt with multiple examples of fraud is intended to clarify the difference between the semantic meaning of our categories of fraud and our definition of these categories of fraud.

We measure the accuracy of our LLM classifier using the 500 manually-coded press releases as a measure of ground truth. We compare our manual coding to the classification produced by our prompt-optimized LLM. For Upcoding fraud, we get a sensitivity of 86.26% and a specificity of 84.03%. For Medical Necessity fraud, we get a sensitivity of 86.96% and a specificity of 73.33%. Finally, for Substandard Care fraud we get a sensitivity of 80.43% and specificity of 72.06%. These results suggest our model is less likely to have false positives than false negatives, and that our approach is more accurate for upcoding and medical necessity frauds than substandard care frauds.