The Economics of Healthcare Fraud

Jetson Leder-Luis* and Anup Malani†

January 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. We first offer a new economic definition of health care fraud that captures and connects the wide range of activities prosecuted as fraud. 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. Next, we examine the incentives for committing and for prosecuting fraud. We show how fraud is driven by a combination of inadequate (expected) penalties for fraud and imperfect reimbursement rates. Public anti-fraud litigation is driven by the relative monetary, political or careerist return to prosecuting fraud and prosecutorial budgets. Finally, we examine the prevalence of health care fraud prosecutions across types of fraud and types of care, and across the US, 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. Despite the large-and-growing size of the healthcare sector and prevalence of fraud and anti-fraud measures, healthcare fraud has received relatively little academic attention relative to other types of malfeasance and regulation in the healthcare industry, such as medical malpractice.

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

*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, Jim Poterba, Jimmy Roberts, Maggie Shi, 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.

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. 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 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). Yet 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 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. 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. Our framework considers (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 to any situation where the payor's choice, the provider's care, and the provider's bill differ. This framework captures three of the most common healthcare frauds prescribed by fraud regulations: upcoding, substandard care, and medical necessity fraud. We also discuss kickbacks and illegal financial relationships, which are prosecuted under fraud-adjacent statutes but are not about the care provided *per se.*

In Section 2.3, we modify 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.

We next alter our framework to address welfare. In Section 2.4, we contrast the care a payor deems appropriate with the care that a social planner would prefer. Our framework immediately 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 tackles related efficiency concerns, particularly providing the distinction between health care fraud and waste, two related but often conflated phenomena.

Having set forth a framework that connect the law on fraud with the economics of fraud, Section 3 provides a review of the economic literature on healthcare fraud. Our review of the literature shows how various case studies of healthcare fraud fit into our framework. This review also highlights the value of this paper: while a number of studies have discussed the policies used to combat healthcare fraud, no work has presented a cohesive, economic definition of the conduct or considered the welfare consequences of healthcare fraud at large.

In Sections 4 and 5, 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 litigate 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 fraud both when prices are far above cost (i.e., care is profitable regardless of need) and when prices paid for care are lower than cost (so that fraud enables cost recovery). In our discussion of regulator incentives we highlight the misalignment between the incentives of payors and the incentives of those tasked with regulation of fraud.

Finally, in Section 6, we provide 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. 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 specialty.

2 Economic Framework

2.1 Legal Standard for Healthcare Fraud

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

Relevant statutes. Legally, healthcare fraud is any conduct that violates a healthcare anti-fraud statute. There are two main statutes that define fraud against federally health insurance programs. First, the federal 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.

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 some authorize jail time (e.g., Cal. Penal Code § 550). Moreover, civil penalties are typically limited to return of fraudulently-obtained funds and forward-looking exclusion from coverage under an insurance plan, and do not include fines beyond this.

In addition to the statutes that prohibit fraud, there are a set of federal and state laws that prohibit so-called kickbacks or referral fees. These statutes prohibit providers for 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. 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). Administratively, providers convicted of criminal offenses, including paying kickbacks, are generally prohibited from any future billing.

Proving fraud. 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. To explain, we take a short detour on Medicare billing requirements and the elements of a fraud claim.

Medicare covers care that is "reasonable and necessary". 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.[1] 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.

In the next subsection we present an economic framework that captures the sort of fraud proscribed by healthcare fraud statutes, but not anti-kickback laws. Whereas fraud statutes prohibit lying about what medical care was provided, anti-kickback statutes on their face prohibit how rents from medical billing are divided amongst providers. While this can induce inefficient levels of care, it is a species of legal fraud that should economically be defined as something distinct from fraud. However, we will discuss the relationship between our framework to anti-kickback statutes in Section 4.

[1]Indeed, 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").

2.2 An Economic Framework for Fraud

Next, we consider an *economic* definition of healthcare fraud, which tries to separate different healthcare provider behaviors that all qualify as fraud, but with different consequences. We consider the 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 unpack our framework, consider the relationship between the provider's actions and the payor's rules. The medical provider, e.g., a doctor, lab or hospital, engages in two types of behavior: providing care to a patient and billing for care provided to the patients. In a simple model, the provider has a choice of providing a low level of care or high level of care for a patient.[2] The provider also has a choice of billing for a low level or a high level of care. The other actor 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. Similarly for simplicity, the payor determines that the patient should receive a low level L of care or a high level H of care.

Table 1 lays out the relationships between the three decisions. The rows track the payor's belief and indicate (a) whether the payor determines that the patient's condition warrants low care or high care. The columns track the provider and indicate (b) whether the provider delivered a low care or high care. Each cell tracks the billing decision of the provider; these cells have two rows, corresponding to (c) whether the the provider billed for a low level of care or high level of care.

Table 1: Economic Categories of Fraud

Payer		Provider: L	Provider: H
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 the bill – fully taxonomize all types of healthcare fraud, as well as a number of combinations

[2] We 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-cost care and high care as high-cost care, rather than assume that the levels of care corresponding to the medical intensity of intervention, e.g., whether to provide a arterial graft rather than a stent.

that are not fraud.

First, we discuss what is not fraud, as well as the set of actions that are unobserved. First, if the provider delivers and bills for the same level of care that the payor determines is appropriate for the patient, then there is not fraud. These categories are marked as <Not Fraud> in the table. There are two such cases: e.g. the provider delivers, provider bills, and payor deems appropriate a high level of care. Our taxonomy also produces 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 study them any further. Third, fraud that is rarely if ever prosecuted are 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.

Finally, we turn to the locus of fraud. Our framework highlights three types of healthcare fraud, which we can use to fully characterize the space of possible behaviors. 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 (and legal) fraud. A prominent example of upcoding was the Columbia-HCA lawsuit, wherein the hospital chain billed for patients as having been treated for severe pneumonia when they had been treated for simple pneumonia (Charatan, 2001; 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 is 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 the prosecution of nursing home owners for providing inadequate care in poor quality facilities, where the 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 the provider delivers and bills for high care, but the payor believes low care is appropriate. This fraud is called "medical necessity fraud". A clear example is the widespread use of expensive genetic cancer testing for patients who had no indication of need for these tests (U.S. Department of Justice, Office of Public Affairs, 2019a).

Table 2 provides real-world examples of the three types of healthcare frauds, using information from Department of Justice prosecutions. These examples highlight the variety of healthcare frauds that can occur in each type of medicine. We provider an empirical characterization of enforcement by category in our Section 6.

Our examples also highlight potentially counter-intuitive cases that fall into our categories. When care is not performed at all, as in the physical therapy case listed in Table 2, this is a special form of upcoding where the service provided (and needed) is $L = 0$. Moreover, we see that the line between certain categories of fraud is occasionally blurred; 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.

Kickbacks are not listed in our table, as they are not an economically meaningful type of fraud, but rather a penalty for the inducement of others to participate in a (potential) fraud. We discuss the relationship between kickbacks and our taxonomy when we discuss prices in Section 4, below, as well as in our empirical analysis of fraud enforcement in Section 6.

Table 2: Examples of Healthcare Fraud with Categorization

Fraud Category	Example	Dollar Amount
Upcoding	Columbia-HCA Hospitals report expensive DRGs	$840 Million settlement
	Ear Nose & Throat Physician reports higher Evaluation & Management codes	$750,000 settlement
	Tenet Health care has hospital patients labeled as "outliers" to receive additional payments	$900 Million settlement
	Physical therapy care never performed	$12 Million Settlement
Medically unnecessary care	Genetic tests for cancer screening	$2.1 Billion estimated losses
	Hospices sued for enrolling unqualified (non-terminal) patients	$330 Million in settlements
	Wheelchairs for patients with no mobility issues	Multiple $10M+ Settlements
	Patients receive spine surgery inpatient instead of outpatient	$214 Million in settlements
Substandard Care	South florida Nursing home owner provides poor quality facilities	$1.2 Billion damages
	Surgeon has students perform surgery, does multiple surgeries at a time, leaves patients under anesthesia	$8 Million Settlement
	Billing for physician home visits and sending an aide	$5 Million

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

Table 3: Medical Malpractice

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

In our malpractice framework, misbehavior is found on the off-diagonal cells. 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); this is the canonical behavior targeted by the malpractice system. 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).

The relationship between Table 1 and Table 3 show 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

[3]More precisely, a provider has the duty to provide the level of medical care that a reasonable provider would. What is reasonable may determine 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.

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.

2.4 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.[4] A fuller definition would allow that there are multiple levels of care where the MB 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 wants high, 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 providers behavior is waste, inefficient, or a transfer, in bold.

The resulting table provides four take-aways. 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

[4] In our theory section 4, 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 patients 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

Payer \ Provider	L	H
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 **Inefficient** **Transfer to Provider**	L: <Not observed> H: <Not fraud> **Potential Waste**

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, 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 inefficently low care imposes greater costs on the payor down the line.

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.* We provide empirical examples from the economics literature in 3, including cases where medical necessity frauds result in increased efficiency.

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 occurred 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.

3 Literature Review

The economics literature on healthcare fraud can be divided into two categories. The first uncovers examples of the different types of fraud in our framework. The second examines incentives to commit fraud and efficacy of enforcement efforts.

3.1 Literature on Examples of Frauds

Upcoding. In early work, Dafny (2005) and Silverman and Skinner (2004) explore 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. More recently, Joiner et al. (2024) investigate upcoding across providers with data from Medicare audits and shows that 0.5% of hospitalizations and 2.5% of physician services have indicia for upcoding. Shekhar et al. (2023) describe patterns in International Statistical Classification of Diseases (ICD) and Diagnosis Related Group (DRG) codes in billing invoices that are consistent with hospital upcoding behavior, and use the frequency of these patterns in claims data to rank hospitals by propensity to upcode. This paper is part of a large, computer science literature that uses healthcare data to detect fraud data, though that literature is largely focused on statistical methods rather then the extent of fraud; Kumaraswamy et al. (2022) provide a survey of the detection methods literature. Finally, Sanghavi et al. (2021) discuss ambulance companies that bill payors for giving patients rides that never occurred, a phenomenon called ghost billing. In our framework, this is upcoding the level of care, where the low level actually given is specifically zero.

Medical necessity fraud. 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. Howard and McCarthy (2021) investigate the medically unnecessary use of implantable cardiac defibrillators, treatments that were prosecuted under the False Claims Act.

Finally, the literature on the over-prescription and diversion of opioids from medical to recreational use also provides examples of medical necessity fraud, because drug diversion supplies users that were unable to obtain legitimate prescriptions; Maclean et al. (2020) provides a survey of the economics literature on opioids, though fraud is not a primary theme of the studies reviewed. Eliason et al. (2019) consider the dialysis market, where in some circumstances epogen was unnecessarily used, particularly before it was included in the dialysis payment bundle.

A handful of papers shed light on the welfare effects of medical necessity fraud. Leder-Luis (2025) examines 4 whistleblower lawsuits under the False Claims Act and involving large settlements. One of these cases involved the use 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 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 not reduce 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.

Substandard care fraud. There is a substantial literature on examples of and on the extent of medical malpractice. While medical malpractice is not definitively substandard care fraud, it is a necessary condition for such fraud if insurers define appropriate care as non-negligent care. The remaining condition is that the provider billed for more or higher quality care than they actually provided.

3.2 Literature on Incentives to Commit Fraud and Effects of Enforcement

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. Other scholars discuss the prevalence of fraud rather than drivers of fraud. For instance, 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.

The literature on the effect of enforcement is also in its infancy, but has a number of important results. Eliason et al. (2025) compares enforcement under civil and criminal ex-post litigation to up-front prepayment authorizations and argues that the latter is more effective when fraud is diffuse, i.e. committed by many providers. Perez and Wing (2019) and Nguyen and Perez (2020) examine state-level enforcement of Medicaid

fraud and find limited effectiveness of decentralized enforcement, though neither study distinguish which types of fraud are pursued or deterred. Leder-Luis (2025) discusses federal False Claims Act whistleblower lawsuits via four case studies of large settlements and shows that these cases produce large deterrence effects. Finally, Becker et al. (2005) identifies which types of patients and hospitals have abusive Medicare billings that are responsive to law enforcement. For a 20 percent random sample of elderly Medicare beneficiaries hospitalized from 1994-98 with one or more of six illnesses that are prone to abuse, they obtain longitudinal claims data linked with Social Security death records, hospital characteristics, and state/year-level anti-fraud enforcement efforts. They show that increased enforcement leads certain types of types of patients and hospitals to have lower billings, without adverse consequences for patients' health outcomes. [5]

4 Incentives to Commit Fraud

Different preferences and information asymmetries across patients, providers and payors create the incentive for fraud by the provider. In addition, incomplete information across parties creates the potential for both efficient and inefficient fraud. Here we sketch some basic features of a model that provides intuition about the incentives to commit each of the different types of fraud in our rubric and when those fraud are (in)efficient.

Our model will highlight 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 is follow's from Becker's model of crime: if one lower 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, respectively. 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 it 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, i.e., will have an incentive to provide medically unnecessary care. Likewise, if the payor sets to low a price for a procedure such that offer a lower profit than other procedures, the provider has an incentive not to perform that procedure even when it is appropriate, i.e., to deliver substandard care.

[5] Although our focus is on the economics and quantitative empirical literature, a large legal literature 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.

We formally model insurance fraud in the context of interaction between a patient and a provider drawn from the population wherein the patient has insurance covered by a payor. We build 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 delivers 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.[6] 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.[7] Second, the provider diagnoses the patient with a specific condition x and provides a treatment y at a cost (to the doctor) of $c(y,x)$.[8] Third, the provider files a claim with the payor stating that the provider delivered service $\tilde{y}$, potentially different than y.[9] 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.

The payor's problem is choose (a) an audit policy $a(\tilde{y})$ and (b) a reimbursement schedule $p(\tilde{y}, a(\tilde{y}) \cdot y)$

[6]Homogeneous 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.

[7]More 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.

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

[9]Here 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.

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 payor what care she delivered:

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

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 $t(x)$:

$$p(t(x)) - c(t(x),x) \geq p(y') - c(y',x) \tag{2}$$

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.[10]

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 increase her revenue by $p_H - p_L$. The audit constraint is intended to stop this upcoding. But audits have

[10] Our 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.

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 provider only audits high-cost treatments, then the threat of audit for those treatments should deter upcoding involving those treatments should be avoided. 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.[11] 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).[12]

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 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

[11] Of 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.

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

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 more than 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:

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

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 $c(t(x), x; \theta) - c(y', x'\theta) > (E_\theta[c(t(x), x; \theta) - c(y', x'\theta)]) + Z$, 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 $t(x)$:[13]

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

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 unnecessary care.[14]

Unknown or heterogeneous costs can explain why inappropriate care is delivered, but what explains

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

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

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

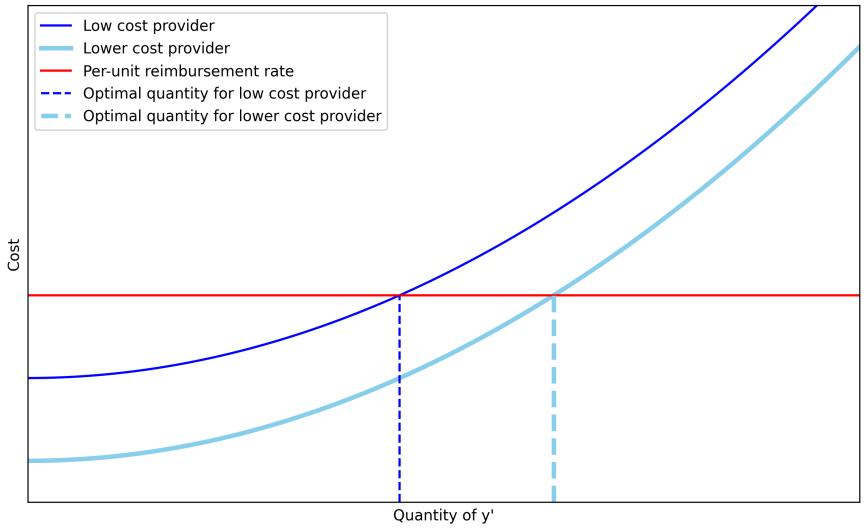


Figure 1: Lower cost providers have higher volumes (given adequate supply of patients)

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 1. 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.

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.[15] 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.

Efficient fraud. Finally, our model can also be embellished to create a distinction between fraud and

hypothetical provider.

[14] An 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.

[15] Likewise, 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.

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.

5 Incentives to Enforce Laws against Fraud

Recall that there are both federal criminal and civil penalties for committing fraud against the federal healthcare programs and state criminal and civil penalties for fraud against private insurance companies. In general, criminal laws can only be enforced by the government, while suits for civil penalties can also

be brought by private parties harmed by fraud. In the case of civil laws targeting fraud against federal government, private whistleblowers can bring suits, despite not being harmed. Unsurprisingly, given the amounts at stake, the majority of all anti-fraud civil cases alleging fraud on the federal government are brought by whistleblowers. However, the Department of Justice (DOJ) can join or take-over such suits.

In this section we explore the federal Department of Justice's (DOF) incentive to initiate criminal or civil anti-fraud actions or to joint a or takeover suits. We do this partly because it allows analysis of criminal cases. Moreover, the federal government tends to join the biggest, if not the majority of private whistleblower actions. Finally, data on cases that the federal government initiates or joins is more readily available than data on independent whistleblower cases.

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 it's resource constraints.

To simplify, we treat the DOJ a unitary actor that obtains value from prosecuting different types of healthcare fraud cases $\mathbf{h} = (h_1, ..., h_4)$ and other types of cases o: $V(\mathbf{h}, o; \mathbf{x})$. We assume agency preferences are well-beheaved and different types of cases are substitutes, implying indifference curves with 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 types. 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 as-if 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 2. Even this rudimentary framework, we obtain some basic insights. Given the agency's budget, a critical driver of the 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

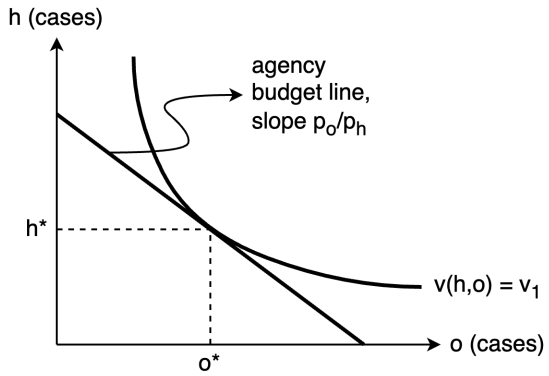


Figure 2: An agencies optimal portfolio of fraud (h) and non-fraud (o) cases (assuming only 1 type of healthcare fraud case)

depends on the ratio of the marginal value of bringing a fraud versus another suit. It is relative, not absolute value that matters.

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.

6 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. (There is no code

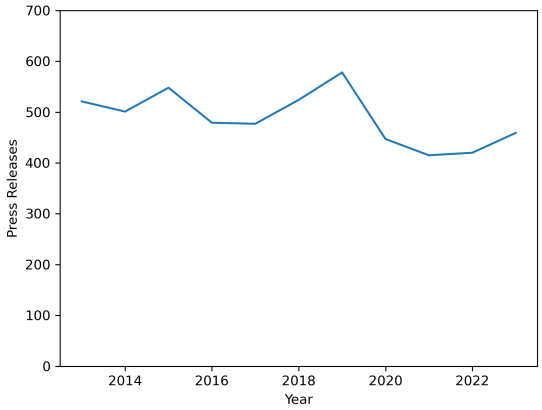


Figure 3: Number of healthcare fraud cases reported in DOJ press releases over time

in claims data that labels claims as fraud!) Moreover, as we explained in Section 5, not all fraud is prosecuted. Therefore enforcement is a highly non-specific test for or measure of fraud, i.e., cases not prosecuted do not prove no fraud (unless fraud is very rare). However, successful prosecution (i.e., prosecution with a positive trial verdict or settlement) is a fairly sensitive test of fraud, i.e., 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.

Nor does successful prosecution tell us about enforcement 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 potential magnitude of fraud in the healthcare system, even imperfect data is better than no data.

6.1 Data and methods

We collected data from the Department of Justice Press Release archive. Each press release details substantial action in an anti-fraud case, generally a settlement or judgment, although press releases can sometimes be released when the government initiates a case, such as indicting a particularly egregious defendant. We filter for all press releases that contain the word "Medicare", which also includes all references to the Centers for Medicare and Medicaid Services, and thus provides an approximate universe of all enforcement related to federal healthcare fraud.[16]

[16]We eliminate anti-trust cases and income tax violations from our dataset, as well as press releases from end-of-year overviews or awards ceremonies that detail year-long behavior instead of individuals cases. The focus on Medicare may potentially miss fraud against the Veterans Affairs system, the Federal Employees Health Benefits Program, or the military

Our data contain 5,369 Press Releases from 2013-2023. We hand-coded 500 press releases for their type of fraud, and then use large-language models (LLMs) to classify each press release by the type of fraud it contains. We further scrape the state in which each fraud occurred and the types of medicine represented. The latter is done by matching text in each press release with medical specialty terms derived from the AAMC specialties list (of Medical Coders, 2024). Appendices A and B contains details about our sample and the categorization process, respectively. Our LLM categorization closely matches our hand-coded sample, indicating that accuracy is not an issue.

6.2 Locus of Fraud

First, we give an overview of the patterns in healthcare enforcement over time, geography, and medical specialty. Figure 3 reveals that reported enforcement of health care fraud has been stable around 500 press releases per year over the past decade.

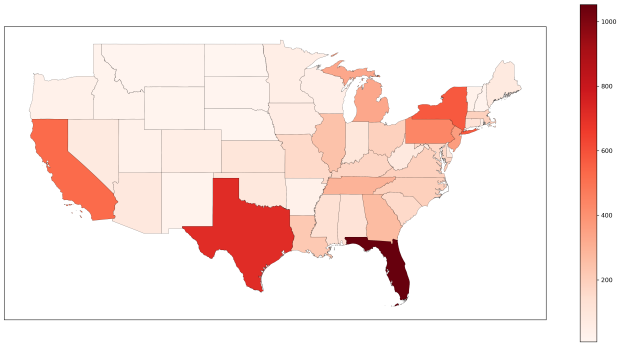
Medicare fraud appears to be more common in Southeastern states. Figure 4 shows the geographic distribution of fraud across states. Panel (a) reports the number of enforcement actions per state. Because we are focused on reported cases of fraud against Medicare, states with larger Medicare populations 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 by the Medicare population in a state in Panel (b), so our scale is <share press releases in a state> / <share medicare beneficiaries in state>.[17] This figure shows that health care fraud enforcement cases are more common on a per-beneficiary level in the American southeast. An interesting comparison 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. In future work we will 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 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).

TRICARE/CHAMPUS system, but these cases often contain a Medicare fraud component as well.

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

(a) Not normalized



(b) Normalized by Medicare beneficiaries

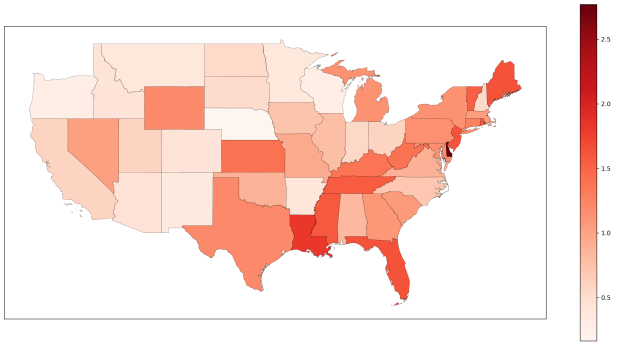


Figure 4: Geography of reported healthcare fraud cases in DOJ press releases, noralized by number of Medicare beneficiaries in each state

6.3 Types of Fraud

Upcoding and medical necessity cases are a larger fraction of prosecuted frauds than substandard care. Figure 5 plots the fraction of press releases in each of 3 categories of fraud: upcoding, substandard care, and medical necessity frauds. Note that percentages do not add up to 100% because a given press release can have multiple types of fraud. We find that upcoding and medical necessity fraud are the most likely fraud to be prosecuted, while substandard care fraud is the least likely to be prosecuted. Upcoding and medical necessity fraud rates track one another because roughly 75% 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 12% of cases.

Finally, Figure 5 also shows the proportion of press releases that have text that tags the cases as a

Type of Medicine	Frequency
Pain	622
Emergency	197
Hospice	179
Genetic	171
Cardiology	140
Psychiatry	131
Vascular	105
Oncology	104
Urology	97
Podiatry	81

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

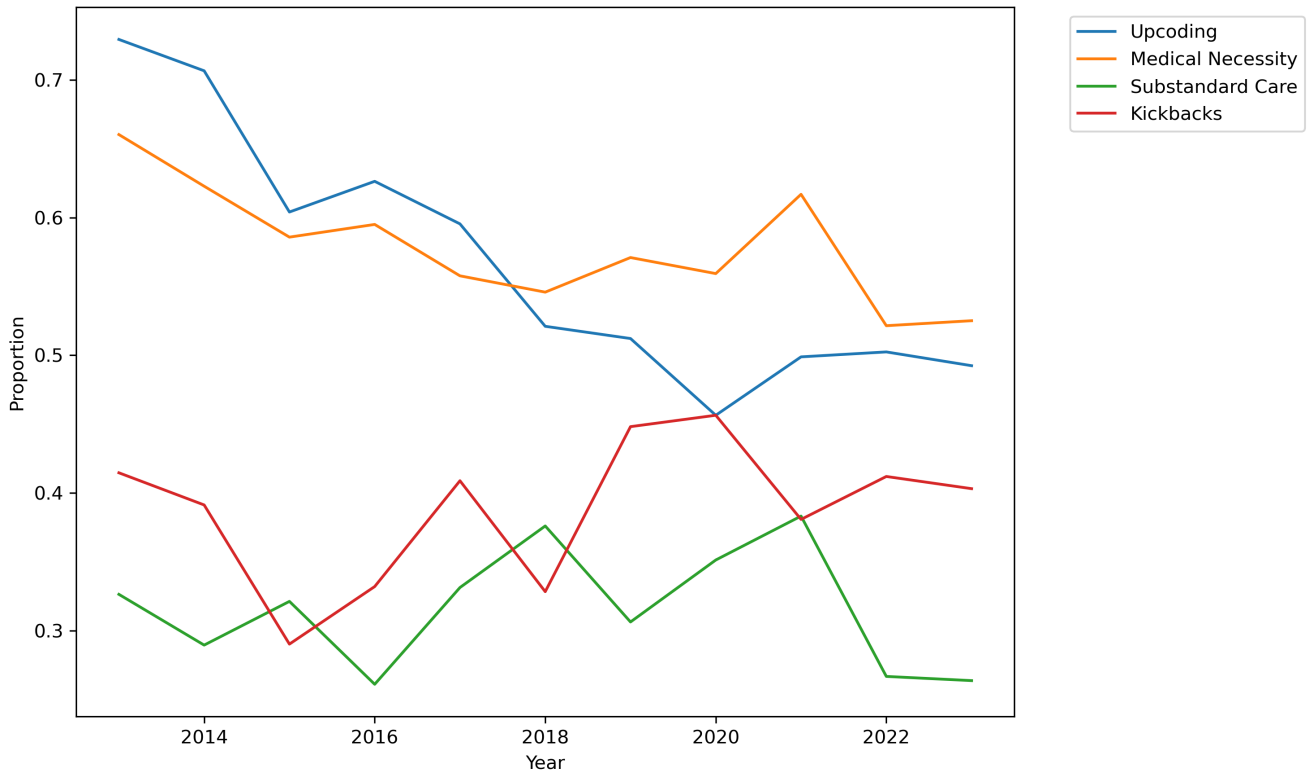


Figure 5: Categories of healthcare Fraud from DOJ Press Releases

kickback. Because types of frauds are not mutually exclusive, a given case can include upcoding and medical necessity fraud, or medical necessity fraud and kickbacks. However, the three types of fraud discussed in the last paragraph generally overlap with each other less than each type overlaps with kickbacks.

7 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 describe 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 limitation 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 southeastern 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 freuqency, including pain meidcine, emergency care, and hospice.

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

Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *The American Economic Review 53*(5), 941–973.

Becker, D., D. Kessler, and M. McClellan (2005). Detecting medicare abuse. *Journal of Health Economics 24*(1), 189–210.

Chandra, A. and D. Staiger (2007). Productivity spillovers in healthcare: evidence from the treatment of heart attacks. *The Journal of Political Economy 115*, 103.

Charatan, F. (2001, Jan). Us settles biggest ever healthcare fraud case. *BMJ 322*.

Dafny, L. S. (2005). How do Hospitals Respond to Price Changes? *American Economic Review 95*(5), 1525–1547.

Drake, Tim Kanu, A. S. N. (2013). American criminal law review. *American Criminal Law Review 50*, 1131.

Eliason, P. J., B. Heebsh, R. C. McDevitt, and J. W. Roberts (2019, 11). How Acquisitions Affect Firm Behavior and Performance: Evidence from the Dialysis Industry*. *The Quarterly Journal of Economics 135*(1), 221–267.

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.

Feldman, R. (2001). An economic explanation for fraud and abuse in public medical care programs. *The Journal of Legal Studies 30*(S2), 569–577.

Galizzi, M. M., G. Godager, J. Li, I. Linnosmaa, T. Tammi, and D. Wiesen (2023). *Economics of healthcare provider altruism*, pp. 1–23. Springer.

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.

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.

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*.

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).

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.

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.

of Medical Coders, A. A. (2024). Specialty profiles.

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.

Silverman, E. and J. Skinner (2004). Medicare Upcoding and Hospital Ownership. *Journal of Health Economics 23*(2), 369–389.

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.

Appendices

A Sample of press releases

Scraping DOJ website and filtering for Medicare-related cases. We scraped a total of 7051 documents from the Department of Justice website, including the years 2009-2023 for main justice and 2013-2023 for the district-level US attorney's offices. These years were chosen because of the availability of data in structured (i.e., scrapable) formats from those years. Each of these press releases contain the word "Medicare". We keep all data 2013 onward, where both sources have coverage, which limits us to 6063 documents.

Next, we drop any documents that contain terms that indicate it is not a health care fraud case. These terms are: antitrust, awards ceremony, outreach campaign and income tax. We then filter strike force actions – a term used by the DOJ to announce multiple criminal arrests for related conduct but without details about each one – using the terms Takedown and Strike Force, which filters out 118 documents. To ensure we are capturing legal actions by the DOJ, we then split the dataset into two, one including the phrases {Sentenced, Settlement, Settle, Convict, Scheme, Guilty, Resolve, Allegation, Allegations, Arrested, Charg, Pay, Interven, False Claims Act, Indict, Doctor, Physician, Suit, Sues, Sued, Lawsuit, Complaint} and one that does not. These phrases indicate that the specific press release they are associated with are individual cases about healthcare fraud. We then manually check the dataset of press releases that do not contain those phrases and switch 95 cases that were individual cases back to the first dataset. This leaves us with 5369 cases.

B Classification by type of fraud

We 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.

B.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. This approach yields an accuracy of 93.6 as compared to manual classification%.

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.

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

To classify these three types of fraud we fine-tune a prompt with a random sample of cases and DSPy, which is a prompt and weight optimization package that builds on top of ChatGPT. To train our prompt, we randomly select 500 case-specific press releases and hand-code them into the 4 categories of fraud using the framework we set forth in Section 2.2. We use GPT 4o mini for our language model. Initially we load in our manually classified set of 500 and turn it from a "DSPy.example" object with our text as input. We also use a different train/test set for each type of these frauds and stratify their respective columns – for upcoding we stratify the manually coded upcoding column, etc.... For each of these three types of fraud we make a signature that takes text as an input and outputs 1 for true (this fraud has been committed) and 0 for false (has not been committed). DSPy trains the LLM on our train set and validates it on our validation set and optimizes the prompt based on that.

DSPy generates different prompts for each type of fraud. The optimized 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.

The optimized prompt for medically unnecessary 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. 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.

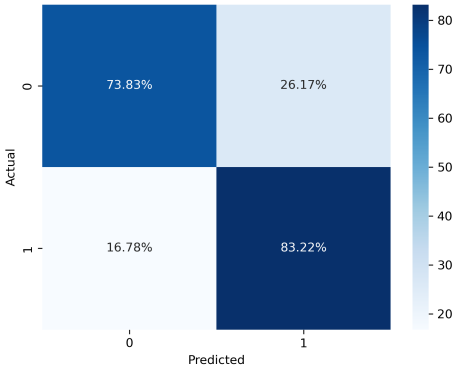
The optimized prompt for substandard care was:

> 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 their is a none zero 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.

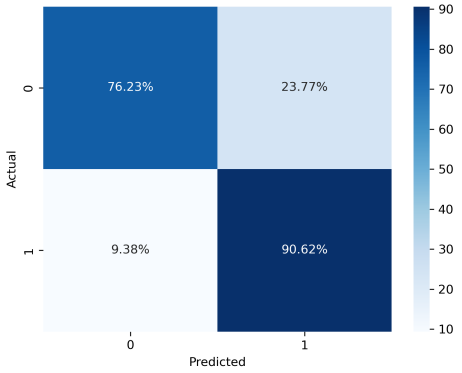
When setting up our DSPy program, we use a ChainofThought module since it is best at thinking step-by-step, and a BootStrapFewShotWithRandomSearch optimizer with different parameters for each fraud type based on trial and error and checking what parameters yield the highest accuracy. We then compile the module using a train set of 100 observations, validation set of 150 observations and a test set of 250 observations.

Finally, we run the optimized module on our test sets, measure the accuracy and make a confusion matrix. For Upcoding fraud, we get a sensitivity of 83.22%, specificity of 73.83% and overall accuracy of 79.2%. For Medical Necessity fraud, we get a sensitivity of 90.62%, specificity of 76.23% and overall accuracy of 83.6%. Finally, for Substandard Care fraud we get a sensitivity of 69.23%, specificity of 79.29% and overall accuracy of 77.2%.

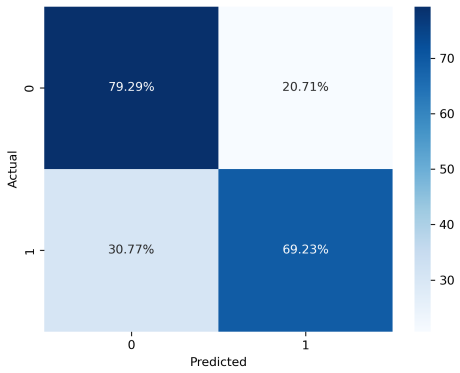
Appendix Figure A1 shows the confusion matrices for the 3 types of health care fraud we categorize using our machine learning tools as compared to the hand-coded sample.



Upcoding



Medical necessity fraud



Substandard care

Figure A1: Accuracy of fraud classification