

WORKING PAPER · NO. 2023-08

Rationing Medicine Through Bureaucracy: Authorization Restrictions in Medicare

Zarek Brot-Goldberg, Samantha Burn, Timothy Layton, and Boris Vabson

JANUARY 2023

Rationing Medicine Through Bureaucracy: Authorization Restrictions in Medicare*

Zarek Brot-Goldberg[†]

Samantha Burn[‡]

Timothy Layton[§]

Boris Vabson[¶]

January 16, 2023

Abstract

High administrative costs in U.S. health care have provoked concern among policymakers over potential waste, but many of these costs are generated by managed care policies that trade off bureaucratic costs against reductions in moral hazard. We study this trade-off for prior authorization restriction policies in Medicare Part D, where low-income beneficiaries are randomly assigned to default plans. Beneficiaries who face restrictions on a drug reduce their use of it by 26.8%. Approximately half of marginal beneficiaries are diverted to another related drug, while the other half are diverted to no drug. These policies generated net financial savings, reducing drug spending by \$96 per beneficiary-year (3.6% of drug spending), while only generating approximately \$10 in paperwork costs. Revealed preference approaches suggest that the cost savings likely exceed beneficiaries' willingness to pay for foregone drugs.

*We thank Frank Gaunt for excellent research assistance. We thank Jason Abaluck, Giovanni Compiani, Stuart Craig, Julie Cullen, David Cutler, Itzik Fadlon, Amy Finkelstein, Josh Gottlieb, Bruce Landon, Jetson Leder-Luis, Lee Lockwood, Bapu Jena, Joe Newhouse, Matt Notowidigdo, Mark Shepard, Maggie Shi, Aaron Schwartz, Amanda Starc, Anna Zink, and seminar participants at Columbia, Chicago Booth, Chicago Fed, Notre Dame, NYU, UC Berkeley, UC San Diego, APPAM 2021, ASHEcon 2022, the 2022 Bates White Life Sciences Symposium, the 12th Annual Empirical Health Law Conference, IIOC 2022, the 2021 NBER Public Economics/Insurance Joint Meetings, and the 2022 SITE Conference on the IO of Healthcare and Consumer Finance Markets for helpful comments. We gratefully acknowledge support from Arnold Ventures, Becker-Friedman Institute, and the National Institute on Aging (under award number P01AG005842 (P29-D26)). All errors are our own.

[†]University of Chicago. Email: zarek@uchicago.edu

[‡]Harvard University. Email: sburn@g.harvard.edu

[§]Harvard University and NBER. Email: layton@hcp.med.harvard.edu

[¶]Harvard University. Email: vabson@hcp.med.harvard.edu

Administrative costs make up a substantial portion of healthcare spending in the United States. Estimates suggest that these costs account for between 20 and 34% of health care expenditures (Woolhandler et al. 2003, Drum 2019, Dunn et al. 2020, Himmelstein et al. 2020), roughly 1-4% of GDP. The academic and policy discussion of the bureaucracies that generate these costs typically characterizes them as wasteful institutions, causing the U.S. healthcare system to be “on a production possibility frontier that is interior to that of other countries” (Cutler and Ly 2011). Eliminating these costs is often seen as a key component of proposals for U.S. health care reform, with the savings often proposed as a way to ‘pay for’ eligibility expansions and increases in generosity of public programs.¹

However, half of administrative effort is spent on activities that aim to *reduce* healthcare utilization and spending, including policies such as auditing claims for fraud, overbilling, or wasteful care as well as enforcing compliance with managed care restrictions that limit access to costly providers, services, and drugs (Cutler 2020a, Chernew and Mintz 2021). While administrative costs can be reduced by making existing bureaucracy more efficient (Cutler et al. 2012), the outright elimination of administrative bureaucracy would also eliminate these activities, potentially resulting in utilization increases that would offset the savings.

In this paper, we take seriously the idea that bureaucracy has both costs and benefits. Bureaucratic rationing mechanisms trade off administrative burden for potential reductions in moral hazard and lower costs of insurance provision. We characterize this trade-off for prior authorization restrictions for prescription drugs. Under such policies, patients can only receive insurance coverage for certain drugs (typically high-cost, on-patent drugs) if they receive explicit authorization; otherwise they must pay the full cost out of pocket. Acquiring the necessary authorization requires the patient’s physician to fill out pre-specified paperwork making the case for why the patient should receive the drug. The goal of these policies is to restrict access to costly drugs to only those patients for whom those drugs provide the highest value. However, prior authorization comes with costs: Making authorization requests is a major source of administrative effort, requiring an average of 20.4 manpower hours per physician per week for physician practices in 2009, their second greatest administrative burden behind billing (Casalino et al. 2009). 34% of physicians report having at least one staff member who works *exclusively* on prior authorization requests (AMA 2017).

We conceptualize prior authorization as a tool for insurers to fight moral hazard problems, where generous insurance coverage may incentivize the use of low-value care (Pauly 1968). Prior authorization forms allow providers to directly communicate information to insurers about the patient’s suitability for the drug, helping resolve a key information asymmetry and allowing insurers to target coverage denials to low-value use. The effort required to fill out the associated paperwork also serves as an ordeal (Nichols and Zeckhauser 1982) that signals the provider’s beliefs about the patient’s suitability, beliefs that would not otherwise be credibly or objectively communicated. The welfare effects of this mechanism contrast the paperwork burden required for inframarginal patients who must go through the authorization process against the reductions in moral hazard for marginal patients who are deterred. Understanding the full welfare consequences of these policies therefore requires measuring the size and composition of these marginal and inframarginal groups.

We study prior authorization empirically in Medicare Part D, the public drug insurance program for the

¹One argument for single-payer reform is that traditional Medicare spends less per beneficiary in administration than private insurers (Archer 2011, Frakt 2018). Proponents of single-payer bills have argued that such reform would reduce administrative expenses by 50-60%, with this reduction having no effect on other outcomes (Pollin et al. 2018, Friedman 2019).

elderly in the United States. We focus on the Low-Income Subsidy (LIS) program. The LIS program has two appealing features: First, LIS beneficiaries receive large cost-sharing subsidies such that they effectively pay nothing out of pocket for covered drugs, making prior authorization the primary feature of the insurance contract that shapes drug demand. Second, LIS beneficiaries frequently face default rules which assign them to a randomly-chosen plan if they do not make an active plan choice, with these defaults typically binding (Brot-Goldberg et al. 2021). This provides us with exogenous variation in exposure to prior authorization restrictions (which differ across plans) at the person-drug level. Since cost-sharing cannot be applied to this population, the use of prior authorization policies is not uncommon: In 2015, prior authorization policies applied to roughly 4% of prescriptions and made up 20% of net drug spending.

We begin by measuring the effect of prior authorization on drug utilization. While we have random assignment to plans, assignment to prior authorization restrictions across drugs within a plan is nonrandom. Our research design therefore compares, *within a given drug, region, and year*, utilization for beneficiaries who are enrolled in plans that have authorization restrictions on that drug against those assigned to plans that cover the drug without restriction. We instrument for the authorization restriction actually faced by the beneficiary using the restriction status of the drug in the plan that the beneficiary was randomly assigned to. Our instrument is strong, with 91% of beneficiaries complying with their assigned plan. We estimate that prior authorization restrictions reduce the use of focal drugs by 26.8%, with slightly larger relative effects among non-white and older patients, and smaller relative effects on drugs in high-benefit classes.

Understanding the effect of prior authorization on drug spending requires us to understand how marginal patients substitute to alternative options. We explore substitution using a nested logit discrete choice model of drug demand, allowing the nesting parameter to govern the extent of substitution on the intensive margin (to therapeutic substitute drugs) versus the extensive margin (to no drug). We estimate that roughly half of patients substitute on each margin. Accounting for substitution, we estimate that the status quo use of prior authorization policies reduced total drug spending by 3.6%, or \$96 per beneficiary-year, compared to a counterfactual world in which all restricted drugs were instead covered without authorization requirements. This reduction in spending is comprised of a \$112 per beneficiary-year reduction in spending on restricted drugs and a \$16 per beneficiary-year increase in spending on cheaper, unrestricted drugs.

These results indicate that prior authorization policies clearly lower the cost of insurance provision. However, they also generate social costs due to the administrative burden they impose on providers who must fill out paperwork, as well as on payers who must process the paperwork. Our data do not permit us to directly measure these costs. Instead, we calibrate per-application administrative costs from prior studies and combine them with our demand model to estimate the size of the burden. Under our preferred calibration, we estimate that the administrative burden is roughly \$10 per beneficiary-year, less than 10% of the spending reduction. While the costs of bureaucracy are nontrivial, this suggests that they are second-order relative to the effects on utilization, and eliminating prior authorization would be cost-*increasing* rather than cost-decreasing. However, this result is not necessarily universal. In an alternative simulation, we show that if prior authorization restrictions were applied to *all* currently-unrestricted drugs, rather than just those where we observe restrictions applied in practice, the cost of the additional administrative burden would exceed the spending reduction these restrictions would achieve. In other words, insurers seem to impose prior authorization restrictions where they are most likely to be cost-reducing on net: high-cost,

niche therapeutics, where the group of inframarginal users is small relative to the size of the marginal group.

While our empirical results and calibrations suggest that prior authorization reduces net *financial* costs, the question remaining is whether moral hazard is large enough to justify such policies; that is, whether marginal patients' valuation for forgone drugs is below the money payers would have otherwise paid to procure those drugs. Estimating consumer valuation of forgone drug consumption is a difficult exercise, given previously-documented behavioral frictions which cause under-consumption of drugs (Baicker et al. 2015, Chandra et al. 2021). However, benchmarking cost savings against estimates of consumer willingness-to-pay for forgone drugs can still be useful for assessing the potential magnitude of the lost surplus. Estimating willingness-to-pay is impossible for the beneficiaries in our sample, as LIS beneficiaries face no out-of-pocket prices. Instead, we estimate price-responsiveness from an alternative, but similar, sample of beneficiaries who we observe transitioning into the LIS program from the unsubsidized component of Medicare Part D. This transition shifts out-of-pocket prices from positive amounts to approximately zero. We estimate a price semi-elasticity of demand of 0.15, making prior authorization policies approximately equivalent to charging \$227 more per prescription per year. This price exceeds the paperwork cost of prior authorization in a year by an order of magnitude.

We use this demand elasticity to infer patients' willingness-to-pay for forgone drugs and compute the consumer surplus lost from prior authorization. Since we cannot estimate willingness-to-pay specifically for beneficiaries who were marginal with respect to prior authorization restrictions (as we do not know where these beneficiaries fall on the demand curve), we instead generate estimates under various bounding assumptions about where on the demand curve the forgone consumption came from. First, we assume that screening is perfect in that the beneficiaries with the lowest willingness-to-pay for the drug are those that are screened out (the best-case scenario, inducing the lowest possible surplus loss). Second, we assume that a random set of beneficiaries are screened out. In these two scenarios, we compute willingness-to-pay for the forgone drugs at \$13 and \$81 per beneficiary-year, respectively. As long as screening is better than random, the net financial savings that prior authorization generates exceeds the amount beneficiaries are willing to pay for the forgone drugs. We then show that under even extreme assumptions about the relationship between willingness-to-pay and actual consumer valuation, under the perfect screening case, lost consumer surplus is likely to fall below financial savings. Under random screening, however, essentially any wedge between willingness-to-pay and consumer valuation will tend to cause lost surplus to exceed savings.

Finally, we estimate the effect of prior authorization on patient health. Since our variation is at the patient-drug level but health is measured at the patient level, this presents a challenge. We use two approaches: A case study on oral anticoagulants, where potential health effects are easy to isolate and can be fast-acting in response to poor management; and an approach where we aggregate exposure to prior authorization up to the patient level across all drugs the patient took in the prior year. Both exercises produce modest but noisy estimates. While the estimated health effects are statistically indistinguishable from zero, we also cannot reject that prior authorization generates large health effects in either direction.

Ultimately, we interpret our results as providing one clear implication and one murky implication. First, our results suggest that, even under some unfavorable assumptions, prior authorization policies clearly produce program savings that exceed the administrative costs they induce. Second, whether these savings are worth the loss in consumer surplus remains unclear. While the beneficiaries diverted to a clinical substi-

tute may experience little surplus loss, the fact that just over half of beneficiaries are diverted to no drug at all is concerning. While, under some assumptions, the amount beneficiaries are willing to pay for the forgone drugs falls far below financial savings, under other assumptions these two quantities are virtually indistinguishable.

Irrespective of whether bureaucracy in health care is an overall force for good or for bad, our results suggest that its *effects on quantities* are of greater orders of magnitude than the direct costs of its operation. This implies that the bureaucratic ‘waste’ that has been the main focus of prior research (Casalino et al. 2009, Cutler et al. 2012, Gottlieb et al. 2018, Dunn et al. 2020) needs to be counterbalanced with the direct effects of bureaucratic activities. Little other work exists to measure the latter; the closest to our study is Dunn et al. (2021), who show that more aggressive use of claim denials reduce the willingness of providers to contract with insurers.² While they focus on quantifying the harms from such quantity reductions, we quantify both the benefits and harms of those reductions, and our results suggest that (at least in our setting) the losses from harms to beneficiaries may not exceed the benefits of financial savings.

Our results also contribute to a broader literature on the trade-offs inherent in bureaucracy. The result that provider-facing bureaucratic review may generate positive social welfare effects is in line with recent work on authorization restrictions for non-emergency ambulance rides (Eliason et al. 2021), claims audits for inpatient hospitalization (Shi 2022), and opioid monitoring (Alpert et al. 2020).³ This prior work stands in contrast to recent work on beneficiary-facing bureaucracy (Deshpande and Li 2019, Finkelstein and Nowowidigdo 2019, Homonoff and Somerville 2021, Shepard and Wagner 2022), which has tended to find that bureaucratic hurdles screen out high-value uses.⁴ The differences between these parts of the literature suggest that ordeals are more likely to work well when they occur as a result of policies such as prior authorization that directly attempt to elicit information in a costly way, rather than screen purely through the burden itself.

Finally, we contribute to a literature on rationing mechanisms in health care. Since at least Pauly (1968), health economists have thought about what mechanisms best allocate health care in the face of potential moral hazard issues. Economists have typically focused on price-based mechanisms such as greater patient cost-sharing (Zeckhauser 1970). However, recent empirical work has suggested that cost-sharing serves as a poor rationing mechanism, often inefficiently screening out the use of high-value care for low-income households (Baicker et al. 2015, Brot-Goldberg et al. 2017, Chandra et al. 2021, Gross et al. forthcoming). Our work suggests that non-price rationing mechanisms may be provide a promising alternative.

²A related literature has studied the effects of bureaucratic institutions that deal with contractual incompleteness on prices and quantities in public procurement, see, e.g., Bajari et al. (2014).

³There is also a small existing literature on the quantity effects of prior authorization policies for prescription drugs (Seabury et al. 2014, Sarig 2020). Closest to our work is Dillender (2018), who estimates the effects of prior authorization for a small set of abuse-prone drugs in the Texas worker’s compensation insurance program. This literature has generally used time-series variation in the imposition of prior authorization restrictions, for which the estimated effect may be confounded by evolving patterns of drug utilization. Our approach focuses on random variation within a market, precluding this confound.

⁴This literature has generally only considered the screening value of bureaucracy and not the burden on inframarginal recipients, since any policy which reduces targeting efficiency is inefficient no matter how large the burden it imposes is. A smaller literature on in-kind transfer program design has considered trade-offs of improved targeting efficiency against reduced value for inframarginal recipients, but has primarily focused on the design of the transferred service (Lieber and Lockwood 2019, Waldinger 2021).

1 Prior Authorization Restrictions in Theory and Practice

1.1 Prior Authorization Restrictions in Practice

The vast majority of health insurance in the United States is provided by managed care organizations (MCOs), private firms that provide insurance coverage. These firms typically place restrictions on this coverage to keep costs down (Glied 2000). Nearly all insured Americans face managed care policies of some kind. Prior authorization restrictions are one policy in an MCO’s toolkit for reducing costs and ensuring appropriate care.

When a service or drug is under a prior authorization restriction, in order for the service or drug to be covered, the patient’s medical provider (rather than the patient herself) must fill out a form provided by the MCO. Authorization forms for prescription drugs generally require the provider to answer some yes-or-no questions regarding why they are choosing to prescribe a restricted drug, particularly when an unrestricted option is available, as well as the patient’s history of taking the restricted drug (possibly under a different insurer) as well as other drugs used to treat the condition in question. Generally, the provider will be asked to provide medical documentation of the assertions made in the form. In Appendix C we provide some examples of prior authorization forms used by MCOs. After the form is submitted, the provider and patient must wait until the MCO approves the request. Authorization requires an administrator at the MCO to review the application and respond accordingly. This generally takes between 1 and 5 business days (AMA 2017). If the authorization is approved, the patient can then receive the drug or service with standard insurance coverage. If not, they will not be able to use coverage unless their provider makes another request and receives authorization.

Prior authorization restrictions are generally applied to discrete services.⁵ Prescription drugs, especially specialty and high-cost branded drugs, are the most common treatment to face restrictions, with more than half of all prior authorization requests being drug-related (AMA 2017). Other commonly-restricted services include certain surgeries, durable medical equipment, and imaging, most of which are also highly discrete services (AHIP 2020). In Section 2.4 we describe how prior authorization is used in our empirical setting.

The stated purpose of prior authorization restrictions is to limit the use of expensive drugs and treatments to those patients for whom those drugs and treatments provide the highest value. In theory, these types of policies do this via two mechanisms. First, the responses to questions on the prior authorization forms explicitly transmit information about value to the patient from experts (physicians) to payers.⁶ Second, the physician’s willingness to complete the forms (possibly multiple times) implicitly signals to the payer that the value of the drug or treatment to the patient is high enough to justify going through the (costly) prior authorization process. Thus, while prior authorization acts as an ‘ordeal’ in the logic of Nichols and Zeckhauser (1982), it is more than that. Indeed, rather than being a *pure* ordeal with no benefit other than

⁵Since a single hospital stay or physician office visit is comprised of a bundle of many services, requiring prior authorization for some subset of those services would be unnecessarily disruptive, forcing providers to deliver care in a piecemeal way. For such categories, MCOs typically instead employ *retrospective* utilization review, rescinding payment for wasteful or fraudulent service provision. See e.g. Dunn et al. (2021) and Shi (2022) for studies of such mechanisms.

⁶The form also allows for communication in the opposite direction: By laying out explicit guidelines, the form also allows insurers to communicate their beliefs about cost-effectiveness to providers, helping guide them away from actions which the insurer might challenge ex post.

screening out those who will not go through the ordeal, it is an *informative* ordeal that potentially screens on both behavior and information transferred from the expert to the payer.

1.2 A Model of Prior Authorization Restrictions

To fix ideas and motivate our empirical analyses below, we present a simple model of prior authorization restrictions in the spirit of [Finkelstein and Notowidigdo \(2019\)](#).

Consider a patient-physician pair deciding whether the patient should receive a drug d . The patient i values the drug with valuation v_{id} . Let $\Delta v_{id} = v_{id} - v_{i(-d)}$ denote the incremental value that i has for d over their next best alternative $-d$, which includes taking another drug or taking no drug at all. Similarly, let Δc_{id} denote the incremental (social) cost. Finally, let $\theta_{id} \in [0, 1]$ be an index of beneficiary types, with the associated mappings $V_d(\theta) = \Delta v_{\theta d}$ and $C_d(\theta) = \Delta c_{\theta d}$. We assume that patients are fully insured and thus face no out-of-pocket price for taking any covered drug in the choice set.

We assume that the joint decision-making process of the patient and their physician has a utility representation, and that their incremental choice utility for d relative to the next-best option is $u(\theta_{id})$. This utility function will reflect some combination of the patient's and the physician's preferences over different drugs. The patient will receive the drug if $u(\theta_{id}) \geq 0$, and will receive private value $V_d(\theta) \times 1\{u(\theta_{id}) \geq 0\}$. While $u(\cdot)$ is the positive argument that determines behavior, $V(\cdot)$ is the normative argument that determines valuation and welfare.

We assume a utilitarian social welfare function, where incremental social welfare is the sum of private valuations, minus the social cost of procuring drugs for those who receive them.⁷ In this setting, that will be

$$W(0) = \int_{\Theta_0} [V_d(\theta) - C_d(\theta)] d\theta$$

with $\Theta_0 = \{\theta : u(\theta_{id}) \geq 0\}$, the set of θ_{id} -type patients who choose the drug.

One potential choice utility function is simply $u(\theta_{id}) = V_d(\theta_{id})$, i.e., patients get the drug if they have a positive incremental value for it. Since patients do not internalize social costs, under this choice utility function, patients for whom private value is positive but social value is not, $0 < V_d < C_d$, will inefficiently receive the drug, the classic case of moral hazard ([Pauly 1968](#)). As mentioned above, prior authorization on d serves as a tool for fighting this inefficiency. Under prior authorization restrictions, the patient will only get d if a constant effort cost a (to fill out prior authorization paperwork) is paid by the physician. Moreover, inappropriate requests may be rejected.

The presence of prior authorization restrictions affects social welfare in two ways. First, prior authorization changes who gets drugs. Specifically, these restrictions generate a new choice utility function $u_A(\theta_{id})$. This choice utility function may change because physicians now have a higher cost of prescribing d ; it may also change because physicians anticipate being rejected if they request authorization for a given patient. Second, authorization restrictions introduce a new administrative cost a that must be paid for each

⁷We do not include manufacturer profits in social welfare, which we believe is consistent with how regulators would view social welfare in this setting. A wider view of social welfare might include manufacturer profits (thus replacing the cost of procuring drugs with the cost of *producing* them, likely to be lower), but would also necessarily include the cost of procuring drugs as measured in the cost of procuring public funds to finance the drugs.

inframarginal patient, i.e. those who get the drug despite the presence of restrictions.⁸

Social welfare under prior authorization must account for both of these changes, and will thus be

$$W(1) = \int_{\Theta_1} [V_d(\theta) - C_d(\theta) - a] d\theta$$

with $\Theta_1 = \{\theta : u_A(\theta_{id}) \geq 0\}$ representing the set of inframarginal patients.

Given this setup, we can evaluate the welfare impact of prior authorization as (suppressing $d\theta$):

$$W(1) - W(0) = - \underbrace{\int_{\Theta_M} V_d(\theta)}_{\text{Reduction in patient surplus}} + \underbrace{\int_{\Theta_M} C_d(\theta)}_{\text{Reduction in program costs}} - \underbrace{\int_{\Theta_1} a}_{\text{Sludge for inframarginals}}$$

with $\Theta_M = \Theta_0 \setminus \Theta_1$ denoting the set of marginal patients who are deterred from the drug as a result of the restrictions.⁹

This welfare change has three components. First, patient surplus is reduced, since the program is moving them away from their most-preferred choice to another option that they value less.¹⁰ Second, the social cost of providing insurance will fall, proportional to the size of the marginal group and to what extent their alternatives are less costly. Finally, to implement prior authorization restrictions, every inframarginal patient must have paperwork done on their behalf, generating administrative sludge. This will lower social welfare in proportion to the size of the set of inframarginal patients Θ_1 .

Prior authorization, in this model, can act similarly to an efficient ordeal (Nichols and Zeckhauser 1982). As an example, take the case where choice utility is $u_A(\theta_{id}) = V_d(\theta_{id}) - a$. The patient will receive the drug if $V_d(\theta_{id}) \geq a$. If the authorization paperwork is exactly arduous enough such that a is equal to the expected social cost of procuring the drug, then prior authorization will efficiently screen out patients who value the drug below cost, while still allowing those who value the drug above cost to receive coverage for it. This need not be the choice utility function applied in practice, however. Physicians might not weigh patient valuation identically to their own costs; furthermore, ‘behavioral hazard’ (Baicker et al. 2015) may cause the patient and/or physician to overreact to the burden, generating a wedge between valuation and choice utility. In these cases, prior authorization may inefficiently screen out high-value uses of the drug.

Understanding the welfare impact of prior authorization restriction thus requires us to quantify the total reduction in program costs, the total administrative burden created by paperwork, and the reduction in patient surplus. In Sections 4, 5, and 6 we attempt to estimate these three quantities.

1.2.1 When Should Policymakers Restrict Drugs?

Before moving to estimation, we can first use this welfare arithmetic to discuss, in general terms, what drugs are the best candidates for restrictions under a utilitarian social welfare function. First, prior authorization is

⁸In reality, administrative costs also must be paid for marginal patients for whom the physician submits paperwork but whose requests are rejected by the insurer. For this section, we assume that no rejection occurs because physicians can perfectly predict who will be rejected. We revisit the role of rejection in inflating administrative costs in Section 5.

⁹We assume $\Theta_1 \subset \Theta_0$, i.e. that there are no ‘defiers’ who get the drug only when authorization restrictions are in place.

¹⁰Some prior authorization restrictions are done for safety reasons, where the patient and physician may not know that the drug is unsuitable for the patient. In this case, patient surplus may rise rather than fall.

unlikely to work well when there are many inframarginal users relative to the number of marginal users, as the administrative cost must be paid for every inframarginal patient. Instead, sludge costs can be minimized by targeting drugs that have relatively few inframarginal users, i.e., drugs that are relatively niche and meant for a specialized population. Second, our model indicates that prior authorization, like any rationing mechanism, is socially useful when moral hazard for a drug is high, i.e., when the incremental value of the drug is low relative to its incremental cost.

Both factors are relevant in the market for prescription drugs. There are many drugs that treat only small groups of patients, including drugs for specific types of cancer and other rare conditions. Incremental value and cost also differ greatly across drugs. A drug's incremental value will be lower when it is in a mature market with many existing substitutes; its incremental cost will be higher when the drug is expensive on a per-unit basis and when there are existing low-cost generic substitutes. Incremental value will be highest when there is no clinical alternative to the restricted drug and lowest when the restricted drug has close clinical substitutes (like a generic equivalent).

The ideal drug to restrict, from this perspective, is an expensive, niche branded drug, especially one that is a new entrant within an established therapeutic class. The worst are those like generic aspirin: Drugs which can be cheaply procured, have high incremental patient value (since the next alternative is likely to be nothing), and substantial numbers of inframarginal users. One caveat applies: If the incremental net social value of a drug is too small (e.g. expensive branded drugs with cheap bioequivalent generic substitutes, where there is little justification to purchase the branded option), prior authorization will be too weak a tool to use to improve social welfare since it may still permit uses of the drug, essentially all of which are inefficient. In that case, a policymaker should want to *exclude* the drug from coverage outright.

2 Setting & Data

2.1 Medicare Part D and the Low-Income Subsidy

Our empirical setting is Medicare Part D, the drug insurance component of Medicare. Under Part D, drug coverage is fully outsourced to private insurers contracted to provide coverage on the government's behalf. The Medicare program organizes a centralized market in which beneficiaries may select from one of these private plans, segmented by geographic service region. Plans have wide scope to differentiate themselves in terms of what drugs they offer insurance coverage for and to what extent they apply cost-sharing or utilization management policies (such as prior authorization) to each covered drug.¹¹ Consumers choose from the plans offered in their service region, each plan charging a monthly premium for enrollment.

Part D beneficiaries with financial need are granted additional subsidies through the Low-Income Subsidy (LIS) program, which offers supplemental drug premium and cost-sharing support. Around 30% of Medicare beneficiaries participate in the LIS program. 'Dual-eligibles,' who also qualify for their state's Medicaid program, are automatically enrolled in the LIS program when they qualify for Medicare, as are beneficiaries of the Medicare Savings Program. Others who meet income and asset eligibility criteria can

¹¹Plans must offer insurance coverage, with or without utilization management, for at least two drugs in each of 148 therapeutic classes.

enroll by applying directly.

Full LIS recipients receive a subsidized reduction in their plan premium payments up to the ‘benchmark’ amount, meaning that those enrolling in a subset of plans (known as ‘benchmark plans’) would not be charged for premiums.¹² Beneficiaries typically have access to between two and sixteen benchmark plans, with 92% of beneficiaries having at least 5 to choose from. We plot a histogram of this count in Appendix Figure A1. Full LIS recipients additionally receive substantial cost-sharing subsidies. For any drug that is covered by their plan’s formulary, they face a custom copayment schedule, with Medicare subsidizing any difference between their regulated copayment and the payment mandated by their plan. In 2020, they were charged a copayment of \$1.30 for all covered generic drugs and \$3.90 for all covered branded drugs, though in most cases these nominal copayments are not actually collected. This policy makes plans effectively uniform in their financial characteristics for full LIS recipients, nullifying any variation in cost-sharing.

Instead, for these beneficiaries, plans primarily differ in terms of the set of drugs covered by their formularies, along with the use of utilization management tools.¹³ This differentiation is substantial. Taking the popular anti-cholesterol drug Lipitor as an example, of the nine benchmark plans available in New York in 2009, six plans covered the drug on their formulary while three did not. Among the six plans that did cover the drug, two required prior authorization for beneficiaries to obtain coverage, while four did not. Beneficiaries aiming to take Lipitor would thus have vastly different experiences across plans.

Beneficiaries who qualify for the LIS program are automatically assigned to a benchmark plan by default if they do not actively choose a plan when they initially enroll in Medicare. This plan is uniformly-randomly chosen from the set of benchmark plans available in the beneficiary’s service region. Moreover, if a beneficiary was previously automatically enrolled in a plan whose premium, in a later year, rises above the premium subsidy and therefore is no longer a benchmark plan, that beneficiary is automatically reassigned to a randomly-chosen benchmark plan by default if they do not make an active choice. We direct interested readers to [Brot-Goldberg et al. \(2021\)](#) for a more detailed description and study of the default assignment mechanism in this setting.

2.2 Data

We use several administrative datasets from the Centers for Medicare and Medicaid Services (CMS). These data contain information on beneficiary program enrollment status, medical utilization, and prescription drug utilization within the Medicare program. The data are nationwide in scope and extend from 2007 to 2015, tracking drug utilization for all Medicare beneficiaries and medical utilization for all beneficiaries outside of Medicare Advantage.

Beneficiary Demographics, Enrollment, and Choice Status. We obtain information on beneficiary demographic characteristics and plan as well as program enrollment from the Medicare Beneficiary Summary File. This file provides demographic information such as age, gender, and geographic location. It additionally tracks enrollment status at a beneficiary-month level for different Medicare coverage programs,

¹²A different group of ‘partial LIS’ beneficiaries receive lesser subsidies, but are omitted from our analysis.

¹³Note that, in this context, formulary exclusion of a drug means a beneficiary would have to pay the full sticker price of that drug out-of-pocket if they opt to purchase the drug, even if they are in the LIS program.

including Part D, as well as enrollment in the LIS program.

We combine these data with the plan election type file. For all Part D enrollment spells, this file tracks whether enrollment was initiated through active choice or the default auto-assignment mechanism. In addition to listing the plan a beneficiary was enrolled in during each month, the file also includes the default plan that was assigned to the beneficiary, even if the beneficiary opted out of that default. This allows us to observe the *assigned* plan as well as the *enrolled* plan for each beneficiary, even when the two differ.

Plan Characteristics and Formulary Data. We obtain information on plan characteristics from publicly available CMS datasets, which cover all Part D plans offered during our sample period. For each plan, in each region-year pair where it was offered we observe the monthly premium that the plan charged and the plan’s benchmark status.

We use public drug-level formulary data for each Part D plan. This dataset tracks the set of drugs covered by each plan’s formulary each year. For each covered drug, the data indicates the type of utilization restrictions imposed by the plan on the covered drug, including prior authorization, step therapy, or quantity limits. We group prior authorization and step therapy together since they are often applied similarly, and ignore quantity limits, since these are infrequently used.

The original CMS dataset defines drugs by their National Drug Code (NDC), which identifies the strength, dosage form, formulation and package size. We map NDCs to drug active ingredient using RxNorm, the National Library of Medicine repository of clinical drugs. For our analysis, we instead define drug at the combination of active ingredient (e.g., atorvastatin; warfarin) and brand/generic status. In doing so, we effectively treat different doses and different modes of administration as equivalent. We define a drug’s formulary status by the ‘maximum’ coverage across all listed NDCs: If any such NDC is covered without restriction, the drug is considered unrestricted. If any such NDC is covered with an authorization restriction but none are covered without restriction, we consider the drug to be restricted. Finally, if no NDCs are listed on the formulary, we consider the drug to be excluded.¹⁴ This approach also means that we treat identical generic substitutes as equivalent, and treat the full set of generic substitutes as covered so long as at least one is covered by a plan.

Outpatient Prescription Drug Data. We track outpatient prescription drug fills for a random 20% sample of Part D enrollees whose claim-level data are available in the Part D Event files. Each claim represents an event where a beneficiary filled a single prescription of a given drug. For each claim, we observe the specific drug prescribed and filled (at the NDC code level), the quantity/days supply for the fill, as well as the date the fill occurred, and the cost paid directly to the pharmacy by all payers.

Other Drug Information. We use the Micromedex Red Book data, a drug pricing database, to classify drugs. As our main measure of therapeutic class, we use the definition provided therein. Where one active

¹⁴We opt for this definition because not all NDCs are explicitly listed by plans as covered. We observe many claims for NDCs not listed in the formulary, but where an extremely similar NDC is listed as being covered. With these adjustments, this problem is much less common. Disagreement about formulary status within our drug definition across NDCs is uncommon: only 2.9% of drug-plan pairs have at least one NDC that is fully covered and at least one NDC that faces an authorization restriction.

ingredient maps to multiple therapeutic classes, we assign the drug to the class accounting for most prescriptions. Additionally, we use data from SSR Health, which estimates the size of rebates paid to insurers by comparing gross and net revenue from public filings.¹⁵ For each drug, we estimate price net of rebates by deflating list price expenditure using the average rebate for the drug in that year from the SSR Health data.¹⁶ We direct interested readers to [Kakani et al. \(2022\)](#) for more information on the SSR Health dataset.

For our main analyses, we restrict only to drugs that were listed as covered by at least one Medicare Part D plan formulary in that calendar year. This is meant to remove uniformly uncovered drugs from our sample, for which there would be no coverage variation, and additionally to remove miscellaneous drug types whose coverage status we would not be able to track in formularies whatsoever.¹⁷ We additionally restrict to drugs that have a therapeutic class listed in the Red Book database.

2.3 Sample Selection

For our main analyses, we employ a single subsample of LIS beneficiaries. We restrict to those enrolled in Medicare Parts A, B, and D, and not enrolled in Medicare Advantage. Sampled beneficiaries must qualify for the full LIS subsidy. We sample at the beneficiary-year level and require these restrictions to be true for every month in a year in which we include a beneficiary in our sample.

We additionally restrict to beneficiaries who faced the automatic reassignment mechanism described in Section 2.1: Those who were previously automatically-enrolled in a benchmark plan, whose plan subsequently lost benchmark status by charging a monthly premium above the premium subsidy. We focus on these beneficiaries, rather than new Medicare enrollees, since we can observe pre-assignment data for them. We exclude beneficiaries whose reassignment is expected to be non-random based on program rules.¹⁸ Finally, for beneficiaries whose assigned plan retained benchmark status for the year after the beneficiary’s reassignment, we include data for the second year post-reassignment. For beneficiaries whose assigned plan lost benchmark status in the second year post-reassignment, we drop the second year and only keep observations from the first year. We drop observations from 2007 where we cannot observe data from before reassignment.

Table 1 shows summary statistics for our final sample and for the entire LIS population. Our sample is broadly similar to the LIS population in general, except that it is slightly younger and healthier but spends slightly more on both drug and non-drug medical spending than other LIS beneficiaries. Table 2 shows *plan*-level summary statistics for the plans included in our sample, distinguishing between plans that beneficiaries are randomly assigned to by default (in the first column), and those that beneficiaries enroll in, which also includes non-benchmark plans. We define a plan at the region-year level, such that otherwise-identical plans

¹⁵Drug manufacturers pay rebates to insurers, intermediated through their pharmacy benefit managers, as an incentive to give drugs preferred placement on their formularies. Rebates are often paid on a per-prescription basis. This offsets the true price of procuring a drug in a way that is not otherwise reflected in our claims data.

¹⁶The SSR Health data contains average rebates across all payers rather than insurer-specific rebates. This has two limitations. First, the rebates Part D insurers receive may be systematically different from other market segments. Second, insurer-specific rebates may be related to prior authorization schedules, for example if an insurer covers a drug without restrictions in return for a larger rebate from its manufacturer.

¹⁷For example, our formulary dataset generally does not track coverage status for over-the-counter drugs.

¹⁸For example, reassignment will not be randomized if the sponsor of the beneficiary’s incumbent plan also offers another benchmark plan in the region. In that case, all reassignees will instead be auto-assigned to that plan.

offered by the same carrier in different regions are considered to be different plans. The average assigned plan requires prior authorization for 12% of drugs, and excludes 28%, with the remaining 60% covered without restriction. Plans vary in their use of prior authorization, however, with the 10th and 90th percentile plans requiring authorization for 6% and 16% of drugs, respectively. Plans that beneficiaries actually enroll in generally look similar, in aggregate, to plans that they are assigned to.

2.4 Prior Authorization in Medicare Part D

Before proceeding to our main empirical analysis, we describe the use of prior authorization restrictions in Medicare Part D over time and across drug types. This provides some insight into the extent to which authorization restrictions, as applied, reflect the optimal conditions we described in Section 1.2.1. Figure 1 shows the use of prior authorization restrictions in claims for beneficiaries in our sample. The use of prior authorization increased over this period. By 2015, 3.6% of filled Part D claims in our sample involved a prior authorization requirement, accounting for 22% of overall gross spending, and 20% of overall spending net of rebates.

Table 3 shows summary statistics for the drugs we retain in our sample. The average drug (using our drug definition, unweighted by actual utilization) is under prior authorization for 13% of plan-years, but varies considerably across drug type. We divide drugs into three categories: Generic drugs, branded drugs with generic bioequivalents, and branded drugs without generic bioequivalents. Of these three categories, branded drugs without bioequivalents are the most-frequently restricted, with the average drug being restricted in 23.3% of plan-years. These drugs tend to be expensive, with net prices of roughly \$56 per day (compared to \$6 for generics), and niche, with the average drug being used by 0.3% of the population (compared to 1.7% for generics). The least-restricted drugs are branded drugs with generic bioequivalents. This is because, as suggested in Section 1.2.1, prior authorization is too weak a restriction for such drugs. The average drug in this category is, instead, excluded in 57.2% of plan-years.

We examine how prior authorization differs across features expected to predict it. Figure 2 plots the average share of plan-years with prior authorization applied for drugs binned into ventiles of price (defined by average price per day supply in our sample, plotted in log scale). Prior authorization frequency is monotonic in the price, with the top ventile of branded drugs by price being under restriction in 59% of plan-years. In Figure 3 we construct a similar figure but cut drugs into ventiles based on the share of beneficiary-years where the beneficiary filled the drug at least once, with less-used drugs being more likely to face restrictions than highly-used drugs.

Use of prior authorization also differs substantially by therapeutic class. Appendix Table A1 shows the frequency of prior authorization restrictions for the top 30 therapeutic classes by gross Part D drug expenditure during 2008-2015. These classes together make up 83% of gross drug spending. Among the highest spending classes, prior authorization is particularly common for biological response modifiers (affecting 70% of total claims spending), immunosuppressants (66%), and anti-neoplastic (cancer-treating) drugs (58%). Prior authorization is also regularly applied in non-insulin treatments for diabetes (15%) and in anticoagulants (15%), which are used for patients who have had or are at high risk for strokes. On the other hand, prior authorization is less common for important classes like the antihyperlipidemic drugs and

insulins.

Importantly for our identification strategy, prior authorization varies significantly across plans for a given drug. For each drug, in each region and year, we compute the share of offered benchmark plans that restricted that drug. Figure 4 displays the distribution of this share across drug-region-years, omitting cases where the share is 0 or 1, which comprise 74.2% and 2.6% of drug-region-year tuples, respectively. We observe full support across the $[0, 1]$ interval.

3 The Effect of Authorization Restrictions on Drug Utilization

We begin our analysis by estimating the effect of prior authorization restrictions on drug utilization at the person-drug level. We specifically consider the treatment effect of moving a drug from being covered with no restrictions to being covered with restrictions, all else equal.

3.1 Research Design

In estimating the effect of prior authorization on drug utilization, we face two challenges to identification. First, whether a beneficiary faces prior authorization for a drug depends on whether they are enrolled in an insurance plan that restricts that drug; because they are free to choose plans, beneficiaries intending to take specific drugs may be inclined to avoid plans that restrict the drugs they want, introducing reverse causality between (latent) propensity to use a drug and whether a beneficiary faces prior authorization. Second, plans do not randomly select which drugs to restrict. As Section 2.4 showed, the propensity of a drug to face an authorization restriction depends on its price and baseline level of utilization.

Our approach to dealing with issues of beneficiary selection is simple. For beneficiaries who face random assignment to default plans, their assigned plan is, by construction, orthogonal to any underlying drug preferences of the beneficiary themselves. Therefore, we restrict to *only* beneficiaries who faced this randomization. We then use, for each beneficiary-drug pair, an indicator for whether the drug was restricted under the beneficiary’s *assigned* plan as an instrument for whether the drug was restricted under the beneficiary’s *enrolled* plan. This instrument is exogenous by construction and is likely to be strong given the high compliance with default assignment in this population (Brot-Goldberg et al. 2021). Since assignment is random *within a market* (a service region and year pair), we conduct our primary analysis within-market by interacting all fixed effects with market fixed effects.

We deal with selection of plan-drug combinations into restricted status by using rich controls. First, to account for differences in formulary treatment across drugs, we include drug-by-market fixed effects, which absorb any secular differences across drugs in both underlying preferences for that drug, and its propensity to face restrictions. Therefore, all of our analysis is within-drug, rather than across-drug. Second, we include assigned-plan-by-market fixed effects to account for the fact that the *general* restrictiveness of a plan’s formulary may be correlated with other plan design choices that secularly affect utilization levels across all drugs.¹⁹ Finally, we include a control for whether the drug was excluded in the plan the beneficiary

¹⁹We construct this for the beneficiary’s assigned plan so that we can assume that it is exogenously assigned; while we could instrument for the enrolled plan indicators with assigned plan indicators, this would potentially induce a weak instruments problem for plans that are rarely enrolled in by this population. Any bias introduced by this choice should be very small due to the fact that

was enrolled in, instrumented for by an indicator for whether the drug was excluded in the beneficiary's assigned plan. Thus, our primary coefficient of interest estimates the difference between prior authorization vs. coverage with no restrictions.

Given this set-up, the primary remaining threat would be any unobserved or uncontrolled plan design choices that have larger effects on drugs that are more or less likely to also face prior authorization. For example, for two drugs d, d' that treat similar illnesses, if formularies are designed strategically, their formulary statuses will likely be correlated,²⁰ and restrictions on d' will encourage use of d . We account for this by explicitly including controls for the formulary status of therapeutic substitutes. We operationalize this by constructing a single control that takes the weighted share of other drugs in the same therapeutic class as the focal drug that face an authorization restriction, with weights equal to the drug's market share in the entire sample in that year. We also include a similar control for formulary exclusion of substitute drugs.

Our final system of estimating equations is

$$Y_{idt} = \beta_1 \text{Auth}_{idt}^{\text{Enrolled}} + \beta_2 \text{Excl}_{idt}^{\text{Enrolled}} + \kappa_{dm(it)} + \lambda_{j(it)m(it)} + \gamma_1 \text{Auth}_{j(it)dt}^{\text{Sub,Assigned}} + \gamma_2 \text{Excl}_{j(it)dt}^{\text{Sub,Assigned}} + \nu_{idt} \quad (1)$$

$$\begin{bmatrix} \text{Auth}_{idt}^{\text{Enrolled}} \\ \text{Excl}_{idt}^{\text{Enrolled}} \end{bmatrix} = \begin{bmatrix} \delta_1 \text{Auth}_{j(it)dt}^{\text{Assigned}} + \delta_2 \text{Excl}_{j(it)dt}^{\text{Assigned}} + K_{dm(it)} + \Lambda_{j(it)m(it)} \\ \Gamma_1 \text{Auth}_{j(it)dt}^{\text{Sub,Assigned}} + \Gamma_2 \text{Excl}_{j(it)dt}^{\text{Sub,Assigned}} + u_{idt} \end{bmatrix} \quad (2)$$

that is, for every beneficiary i , drug d , and year t , where i was in market m and assigned to a default plan j , we estimate a regression of utilization at the beneficiary-drug-year level on dummies for whether the drug faced a prior authorization restriction or exclusion in the plan that beneficiary was enrolled in during that year, drug-by-market and assigned-plan-by-market fixed effects, and our set of substitution controls. We instrument for the formulary status in the enrolled plan with the formulary status in the assigned plan, along with all of the other controls used.

To identify β_1 , our coefficient of interest, our instrument must be valid. This requires two assumptions: First, that default assignment is conditionally random within-market. This is known to be true institutionally, and we verify it with balance tests below. Second, that the formulary status of a drug on a given plan is exogenous to underlying utilization patterns conditional on controls, i.e., assigned plans do not engage in unobserved actions that differentially affect utilization of specific drugs. Ultimately, this must simply be assumed, though we note that, beyond prior authorization and exclusion, the tools available to Part D plans for modifying drug utilization by LIS beneficiaries are few. We test this assumption to the extent possible below by testing the sensitivity of our estimates to additional plan-by-drug controls.

beneficiaries only rarely opt out of their assigned plans (see Table 4).

²⁰This correlation might be negative if the plan wants to steer patients to a particular drug, or positive if the plan wants to deter beneficiaries requiring such treatment to enroll in their plans (Geruso et al. 2019).

3.2 First Stage

Our research design uses the default plan assignment as an instrument for the formulary a beneficiary faces. Beneficiaries need not comply with this default—they can switch plans at any time post-assignment. The strength of this instrument depends on the level of compliance. In our analysis, we measure the beneficiary’s enrolled plan as of December 31 of year t (the last day of the year following assignment).

Our first stage regression is given in Equation 2 above. We report the results from this regression in the first two columns of Table 4. The first stage is extremely strong, with F-statistics in the tens of thousands. Assignment to a plan that restricts a given drug makes a beneficiary approximately 91% more likely to be enrolled in a plan that restricts that drug, consistent with the fact that (as reported in Table 1) 91% of beneficiaries enroll in the plan that they are assigned to.

As mentioned in Section 2.3, the average beneficiary only takes 10.8 unique prescription drugs in a year, although there are thousands of unique drugs available. Therefore, the majority of beneficiary-drug pairs are irrelevant, and we might be worried about noncompliance with the assigned formulary for pairs where use is likely. Beneficiaries might be more concerned with switching to a plan with formulary that differs in terms of its coverage of the drugs the beneficiary is already taking with little concern for the plan’s coverage of other drugs. To evaluate this, we re-estimate the first stage regressions on a subset of beneficiary-drug pairs where the beneficiary filled a prescription for the drug at least once in the prior year. These beneficiaries should be especially likely to take the drug again in the following year. We present results from these regressions in the third and fourth columns of Table 4. Reassuringly, the associated coefficients are only slightly smaller than those estimated without the restriction. This is unsurprising given that [Brot-Goldberg et al. \(2021\)](#) previously showed that plan assignment in this setting is extremely sticky, with beneficiaries rarely actively choosing plans in response to defaults that *exclude* their previously-used drugs.²¹

We also perform three sets of balance tests to verify that beneficiary formulary assignment is conditionally random. First, in Appendix Table A3, we estimate a placebo first stage regression, estimating whether *contemporaneous* assignment predicts enrollment *in the prior year* in a plan that restricted or excluded a given drug. Second, in Appendix Table A4, we estimate the ‘effect’ of prior authorization restrictions on utilization outcomes (which we discuss further in the next section) *in year prior to assignment*. Finally, in Appendix Table A5 we estimate the ‘effect’ of prior authorization restrictions on beneficiary characteristics (gender, race, age, and Elixhauser Comorbidity Index), which measures whether beneficiaries who face more restrictions are differentially likely to have those characteristics. Reassuringly, in all three cases we can reject even extremely small non-zero effects, suggesting the assignment we observe is indeed orthogonal to beneficiary characteristics.

²¹In Appendix Table A2, we report the results from an exercise where we estimate the first stage with controls, then sequentially add the controls we use until we get to the specification in Table 4. We see that the first stage is weakened when we add drug-level controls. This reflects the fact that many drugs are restricted close to 0 or 100% of the time and thus their formulary status will typically be the same for any two plans, inflating our first stage coefficient. Adding in these fixed effects weights drugs by how close they are to being restricted 50% of the time, which moves the estimates closer to the overall default compliance rate.

3.3 Main Estimates

We now estimate the effect of prior authorization restrictions on utilization as laid out in Equation 1. Our primary coefficient of interest in that equation is β_1 , reflecting the treatment effect of prior authorization restrictions on drug utilization relative to unrestricted insurance coverage. We focus on one primary utilization outcome: A binary indicator for whether the beneficiary filled the drug at least once in the year. We multiply this outcome by 100 so that the regression coefficient represents percentage point changes. We estimate a sequence of regressions that progressively add the controls described in Section 3.1 to demonstrate how their inclusion affects our estimates of the effect of prior authorization. We cluster standard errors at the assigned plan and year level.

We present the results of this exercise in Table 5. Generally, the absolute magnitude of the estimated effects is quite small; nearly all of our effects imply changes in utilization of less than one percentage point. However, given that drugs under prior authorization restrictions tend to be niche and not widely used, benchmarking these effects against baseline levels of utilization is important for assessing the true effect size. We benchmark against the mean utilization for beneficiary-drug pairs which had coverage without utilization restrictions and list this as “Control Mean”. This is roughly 1.3, indicating that the average beneficiary only ever consumes 1.3% of possible drugs. However, this is still not the correct baseline level of utilization to reference. With drug-market fixed effects, regression coefficients produce an estimate equal to the *weighted* average treatment effect, with weights equal to the drug-market-specific variance in usage of prior authorization across plans (Gibbons et al. 2019). We therefore construct a reweighted control mean by taking the weighted average of drug-market-specific control means, with weights equal to $\text{Var}[\text{Auth}_{idt}^{\text{Assigned}}|d, m]$, such that the aggregate control mean reflects the implicit weighting within the regression. The reweighted mean utilization is even smaller, around 0.4, reflecting the fact that heavily-restricted drugs tend to fill small niches. To recover valid percent change effects of prior authorization restrictions, relative to baseline levels of utilization, we divide our coefficient estimates by this reweighted control mean. These implied percent changes are also reported in Table 5 (as “PA % Effect”).

Column (6) in Table 5 presents our preferred estimate of the effect of prior authorization restrictions on any use of the restricted drug in the year, including all of our preferred controls. Under this estimate, prior authorization reduces the use of the restricted drug by 26.8%. These effects are quite large, and refute claims that high prior authorization approval rates mean that prior authorization has little significance for actual utilization. The other columns in Table 5 build up to column (6) by adding successive controls. The most important controls are the drug and drug-year fixed effects, which account for the bias that would otherwise be introduced from comparing frequently restricted drugs (which tend to be less-used to begin with) against frequently unrestricted drugs.

In Appendix Table A6, we explore further specifications. In Columns (7) and (8), we test robustness to our assumption that plans do not differentially influence specific drug utilization by replacing plan-market fixed effects with plan-market-therapeutic class fixed effects and plan-market-price ventile fixed effects, respectively. In column (9), we estimate a specification that accounts for the threat of contamination bias, as highlighted in Goldsmith-Pinkham et al. (2022), by dropping any beneficiary-drug-year observations where the assigned plan excluded the drug. All of these specifications estimate effects that are quantitatively similar

in magnitude to our preferred estimate.

In Appendix Table A7, we replicate our regression in column (6) with alternative utilization measures: The count of prescriptions of the drug filled during the year by the beneficiary, the total days supply of the drug filled during the year, and total allowed spending on the drug. Effects on these outcomes are comparable to effects on our main utilization measure. The similarity of our main effects, which represent ‘any use,’ and these, which represent ‘total use,’ suggest that most of the effect of prior authorization restrictions occurs on the extensive margin of use (filling any prescription) rather than the intensive margin (number of prescriptions filled).

3.4 Heterogeneous Effects

The goal of prior authorization is to deter low-value care. However, it is an empirical question as to whether this occurs in practice. One approach to determining whether prior authorization restrictions are deterring the ‘correct’ care is to examine who is deterred, and what sort of care is deterred. While we do not observe the ‘value’ of each forgone drug, we do observe a variety of characteristics of beneficiaries and drugs. We thus stratify effects on those characteristics to test for differences across groups.

We begin by examining heterogeneous responses by beneficiary demographics. Prior authorization requires physicians to exert effort on behalf of their patients; if they are generally less willing to exert effort on behalf certain groups for reasons unrelated to the value of the drug (e.g. due to the patient’s race or gender), prior authorization may inefficiently deter care for those patients. Similarly, if patients of different races or genders tend to see physicians with different levels of willingness to exert effort on behalf of *all* of their patients, prior authorization may inefficiently cause disparities in use across populations with similar levels of need. We replicate our primary regressions for subsamples of beneficiaries identified by their demographics: White vs. non-white, female vs. male, and by four groups of age. We report effects (and their confidence intervals) for each sub-group in terms of the percent change for that sub-group in Figure 5. We see that there are statistically significantly larger relative effects of prior authorization for older and non-white patients. While men experience larger proportional effects relative to women, this difference is not statistically significant.

In the same figure, we measure differential effects by health status, segmenting beneficiaries by their score in the Elixhauser Comorbidity Index. This index counts the total number of chronic conditions the beneficiary had, as measured by diagnoses that appeared on their medical claims in the year before they were reassigned. We see smaller effects for healthier beneficiaries who have no chronic conditions compared to sicker beneficiaries who do have chronic conditions. We also estimate separate effects for beneficiaries who we observed filling the drug at least once in the year before reassignment, compared to ‘naive’ beneficiaries who would be taking the drug for the first time. We see that restrictions bind less tightly for prior drug users, instead largely discouraging new initiations.

In addition to studying heterogeneous effects by beneficiary type, we also study how effects differ across drug categories. These estimates are displayed in Figure 6. In some cases, prior authorization is used for safety reasons rather than cost effectiveness reasons. Specifically, virtually all generic drugs under prior authorization restrictions are restricted for safety. Other drugs restricted for safety motivations are typi-

cally ‘schedule drugs’ (those indicated as a controlled substance by the U.S. Drug Enforcement Administration). Ultimately, we estimate smaller effects for these categories, consistent with the hypothesis that prior authorization restrictions are less binding when the motivation is safety than when the motivation is cost-effectiveness, and showing that overall estimates are not driven by these types of drugs.

We also investigate heterogeneous effects by whether the drug is used to treat a chronic versus an acute condition, with a ‘chronic’ drug defined as one where the median beneficiary observed filling the drug in a given year did so at least three times in that year. Prior authorization deters chronic-use and acute-use drugs in equal proportion. We also estimate effects for a subset of drugs in classes where we expect benefits to be high, as previously defined by [Brot-Goldberg et al. \(2021\)](#).²² Encouragingly, restrictions bind less tightly on these drugs, suggesting that prior authorization restrictions may be at least modestly well-targeted as a rationing mechanism.

Finally, we also estimate effects for a subset of drugs evaluated by the National Institute for Health and Care Effectiveness (NICE), an organization in the United Kingdom that evaluates prescription drugs on their cost-effectiveness to determine regulation under the U.K. National Health Service. NICE has three categories: ‘Recommended,’ meaning that NICE generally recommends use of the drug for its intended purpose; ‘Limited recommendation,’ meaning that NICE only recommends the drug for certain patients; and ‘Not recommended,’ meaning that NICE does not recommend that physicians ever prescribe the drug. While, unintuitively, effects are slightly *larger* for drugs that are more recommended, the standard errors on these effects are sufficiently large that it is difficult to come to any strong conclusion.

We additionally evaluate the heterogeneous effects of prior authorization by deciles of price per day supply in Appendix Figure A2 (to investigate whether the effect scales with the cost of the drug) and by deciles of the share of plan-years restricting the drug in Appendix Figure A3 (to investigate whether the effect is larger or smaller for drugs that are more commonly restricted). While effects are quite heterogeneous across these groups, they do not follow any clear pattern. Finally, in Appendix Figure A4, we plot class-specific estimates. These results are also quite heterogeneous, but not in any systematic way.

4 Substitution Patterns and Spending Effects

While our above results show that prior authorization restrictions reduce the use of restricted drugs, this is insufficient for assessing the effects of these restrictions on overall spending. To estimate total spending effects, we also need to know how beneficiaries who are deterred from using a restricted drug substitute to alternative options. While some beneficiaries will respond to prior authorization restrictions on a given drug by taking no drug at all, others will substitute to an unrestricted alternative. Using the terminology of our model in Section 1.2, we need to know not only the *share* of marginal beneficiaries (which our above analysis estimates), but we also need to know what their costs would be if the restricted drug was taken out of their choice set, i.e. the *incremental* cost. These substitution patterns cannot be estimated in a reduced-form way, as that would require that we estimate nearly 4 million cross-drug substitution parameters. Therefore, we must impose some parametric restrictions on drug demand to tractably estimate substitution patterns.

²²These include anticonvulsants, antidiabetic agents, antihyperlipidemic drugs, cardiac drugs, oral anticoagulants, antipsychotics, and antidepressants.

4.1 Estimating Substitution Patterns

We model the drug consumption process as a discrete choice of a single drug (or no drug) within a therapeutic class for a given year. This rules out any patterns of substitution or complementarity across classes, and assumes that any drug within the same class is a potential substitute (and *not* a complement). We assume that the beneficiary and their prescribing medical provider choose a drug via a joint decision-making process which admits a stable utility function representation (Brot-Goldberg and de Vaan 2018), with the form:

$$u_{idt} = \underbrace{\beta_C \text{Auth}_{idt} + \delta_C \text{Excl}_{idt} + \kappa_{dm(it)}}_{V_{idt}} + \xi_{iCt} \mathbf{1}\{d \neq 0\} + \lambda_C \epsilon_{idt}$$

We allow beneficiary-provider pairs to have preferences that vary in mean terms across drugs, and allow pairs in different markets to have different preferences across drugs ($\kappa_{dm(it)}$).²³ Beneficiary-provider pairs face a barrier to prescribing due to prior authorization (β) and formulary exclusion (δ), the effects of which are assumed to be constant within a class C but allowed to vary across classes. We normalize u_{i0t} , the mean utility of the outside option of getting no drug ($d = 0$), to zero.

Finally, we assume that unobserved preferences for drugs in the choice set, ϵ_{idt} , are drawn from a standard Gumbel distribution, independent and identically distributed across beneficiary-drug pairs. We also allow an unobserved uniform preference for any drug option (i.e. any option other than no drug) ξ_{iCt} , whose distribution depends on a parameter λ_C such that $\xi_{iCt} \mathbf{1}\{d \neq 0\} + \lambda_C \epsilon_{idt}$ is also distributed standard Gumbel, independent and identically distributed across beneficiaries. This utility formulation implies a nested logit demand system, with a single nest for all ‘inside goods’ (all options that involve taking a drug rather than no drug), with $1 - \lambda_C$ governing the within-nest correlation of unobserved preferences. That is, the probability of taking a drug d is

$$P_{id} = \frac{\exp \frac{V_{idt}}{\lambda_C} \left(\sum_{k \in C} \exp \frac{V_{ikt}}{\lambda_C} \right)^{\lambda_C - 1}}{1 + \left(\sum_{k \in C} \exp \frac{V_{ikt}}{\lambda_C} \right)^{\lambda_C}}$$

and the probability of taking no drug is

$$P_{i0} = \frac{1}{1 + \left(\sum_{k \in C} \exp \frac{V_{ikt}}{\lambda_C} \right)^{\lambda_C}}$$

This nesting structure is essential, because the standard conditional logit substitution effects depend on the share of relevant beneficiaries taking each option. Since most beneficiaries take no drug within any given class, omitting the nesting structure would lead us to incorrectly predict substantial extensive-margin substitution in response to authorization restrictions. Instead, the extent to which beneficiaries substitute on the intensive margin (to another drug) or the extensive margin (to no drug) depends on the nesting parameter $\lambda_C \in [0, 1]$. Lower λ_C values imply that relatively more substitution is on the intensive margin for class C .

²³We omit plan fixed effects here. These would be computationally burdensome to estimate since they would require us to estimate all of our class-specific demand models simultaneously, and our results in Table 5 suggest they have little explanatory power and that their inclusion does not materially affect estimates of effects on utilization of restricted drugs. While plan-by-class fixed effects have more explanatory power, they are not identified separately from λ_C .

Identification of this model largely relies on assumptions similar to those required for our reduced-form approach. Our primary parameters of interest are β_C , the effect of prior authorization on choice utility, and λ_C , which governs the extent of intensive margin substitution. β_C is identified from differences in a drug’s market share among beneficiaries enrolled in plans that restrict it versus the market share among beneficiaries in plans that do not restrict it, holding the formulary status of all therapeutic substitutes fixed. As in Section 3, we instrument for the formulary status in the plan a beneficiary enrolled in with the formulary status in the default plan that they were randomly assigned, to eliminate potential selection bias. λ_C is identified similarly; to estimate it, we need to simply compare the relative market shares of *all other drugs* from beneficiaries in plans that restrict a given drug, compared to the market shares of those drugs from beneficiaries in plans that do not restrict a given drug, holding all else equal (including the formulary status of those drugs).

For example, say that there are two drugs, 1 and 2, and two plans, a and b , where both cover drug 1 without restriction but where plan a restricts drug 2 and plan b does not. The effect of prior authorization on utilization (β) is identified from the difference in the share of drug 2 between the two plans. The nesting parameter λ_C , is identified from the difference in the share of **drug 1** between the two plans. If $\lambda_C = 1$, corresponding to the standard logit demand model, substitution is proportional to baseline market shares, and since most drugs have miniscule market shares over all consumers, then we would expect a tiny change in the market share of drug 1 due to restrictions on drug 2. In contrast, if $\lambda_C = 0$, then all substitution is on the intensive margin, and so we would expect the change in the share of drug 1 to be *exactly equal* to the change in the share of drug 2, but with opposite sign. λ_C is thus identified from the share of marginal beneficiaries who switch away from drug 2 who substitute to drug 1.

In contrast to the underlying dataset we use in Section 3, here we must assign each beneficiary to a unique choice of drug within a given class. To construct the relevant data set, we define a beneficiary as taking a drug within a class if they ever fill a prescription for that drug during the year. For beneficiaries who filled prescriptions for multiple drugs within a class in a year, we assign them to the drug they filled with the highest days supply during the year, and break ties randomly.²⁴ We limit to therapeutic classes where, for at least 10% of region-year pairs, we observe that (1) at least two drugs were ever taken, and (2) at least one drug faced variable prior authorization status (i.e., was restricted in at least one plan and unrestricted in at least one plan). Restriction (2) is required for us to identify direct effects of prior authorization, whereas (1) is required for the identification of λ_C . Our final dataset includes classes making up 98.3% of gross spending. We provide more detail on this restriction in Appendix D.

Our parameterization requires us to estimate hundreds of thousands of fixed effects across many demand systems. Estimating this once via maximum likelihood takes multiple days. Moreover, the log-linear inversion approach of Berry (1994) is unavailable in our case, since there are many drug-plan pairs for which *no* beneficiaries in the plan are taking the drug in a given year; therefore the log of the drug-plan market share is undefined. Instead, we exploit the equivalence between the likelihood functions of the conditional logit and the Poisson generalized linear model (Guimarães et al. 2003) as well as recent improvements in

²⁴On average, 15.07% of beneficiaries who filled a prescription for any drug during the year for a given class received two or more unique drugs. For these beneficiaries, 63.9% of the days supply for drugs in that class were made up for by the drug we pick. For all beneficiaries, the primary drug makes up 90.4% of total days supply.

high-dimensional Poisson pseudo-maximum-likelihood estimation (Correia et al. 2020). We instrument for formulary status in the enrolled plan with formulary status in the assigned plan, and implement this using the control function approach of Petrin and Train (2010). We compute standard errors using a Bayesian bootstrap procedure, clustering at the assigned plan level. This estimation approach is described in detail in Appendix D. In that Appendix, we additionally describe the intuition for the equivalence of the maximum likelihood and Poisson regression approaches to estimating logit demand models.

4.2 Spending Effects of Prior Authorization Restrictions

With the parameters of the model estimated, we can use them to evaluate how the use of authorization restrictions affected total spending and utilization. Using our model, we simulate demand for drugs under: 1) the status quo of beneficiary plan assignment and plan formularies; and 2) an alternative arrangement where drugs that were previously under prior authorization restrictions are now unrestricted, holding all else fixed. We then compute spending in these simulations by assuming that any beneficiary who chooses a given drug spends an amount on it equal to the empirical average amount spent on that drug by beneficiaries in our sample who consumed it in the same year.

We then measure the effects of moving from simulation (2) to (1), the effect of adding prior authorization. We measure the effects on spending and utilization of *all* drugs, not just the restricted drugs. We also break these overall consumption effects down into the effects on consumption of restricted drugs, effects on consumption of drugs not facing prior authorization restrictions, and effects on the share of beneficiaries taking no drug.²⁵ Finally, we measure the diversion ratios, the share of marginal beneficiaries who are diverted to taking another, substitute drug versus to no drug at all.

The results from this exercise are reported in Table 6.²⁶ Our results suggest that prior authorization reduced drug spending by 3.6%, or approximately \$96 per beneficiary-year. This spending reduction is composed of a 21.8% reduction in spending on restricted drugs (\$112 reduction in spending per beneficiary-year), and an offsetting 0.6% increase in spending on (much cheaper) unrestricted drugs (\$15.7 increase in spending per beneficiary-year). Spending increases due to substitution to unrestricted drugs thus do not come close to offsetting savings from reductions in use of restricted drugs.

The size of the spending offset from substitution is affected by a) the differential price of restricted and unrestricted drugs (unrestricted drugs tend to be much cheaper); and b) the portion of beneficiaries substituting to no drug rather than to an unrestricted drug. Estimates of the effects of prior authorization restrictions on the number of users per capita of the restricted drug, unrestricted substitutes, and no drug help us assess the extent to which (b) versus (a) is responsible for the small size of the offset. Table 6 shows the effects of prior authorization on the number of users per capita for each category. First, we find that prior authorization deters 28.9% (0.12 users per capita) of users from taking the restricted drug. Second, of the 0.12 users deterred by prior authorization from taking a restricted drug, 46.2% of them substitute to an unrestricted drug and 53.8% substituted to no drug. The extent of the extensive margin substitution we

²⁵The use of restricted drugs is defined at the beneficiary-drug level, i.e., a drug may be ‘restricted’ for some beneficiaries and ‘unrestricted’ for others. We define a drug as being restricted for a given beneficiary in terms of its formulary status on the plan the beneficiary was enrolled in as of December 31.

²⁶We suppress standard errors for readability. We report the same table with standard errors in Appendix Table A10.

see is substantial, although prior work has suggested that extensive margin substitution makes up an even greater share of the effects of patient cost-sharing (Newhouse and the Insurance Experiment Group 1993, Brot-Goldberg et al. 2017). Our explanation for these results is that part of the extensive margin effect may arise due to patients realizing a prior authorization request is necessary when attempting to fill their prescription at the pharmacy, and not returning to their provider to request that they complete the required paperwork or prescribe an alternative drug.²⁷ Substitution in this way (to no drug) is likely to be undesirable unless treatment for the associated condition was otherwise of very low value.

5 Administrative Cost Burdens

As discussed in Section 1.2, one of the three quantities that is critical for understanding the welfare consequences of introducing prior authorization restrictions is the size of the administrative cost burden generated by their introduction. In this section, we attempt to quantify the size of that burden and compare it to the spending reductions caused by prior authorization restrictions.

Unfortunately, we have no data on the steps in the bureaucratic process of prior authorization. Therefore, unlike prior studies, we cannot directly estimate the cost of compliance from accounting data (Shi 2022) or revealed preference (Dunn et al. 2021); nor can we compute rejection rates. We therefore take an alternative approach, wherein we calibrate relevant parameters (per-application costs and rejection rates) and combine them with our demand system estimates to estimate the total paperwork burden generated by compliance with prior authorization restrictions.

Specifically, we assume that, for any beneficiary-drug pair, authorization must be received once in a given year if it is required and if the patient wishes to obtain the drug.²⁸ We assume that making a request incurs some constant joint cost to both the requesting physician as well as the insurer, which we call a . The number of requests is also unobserved. We assume that any patient we observe taking a restricted drug must have made an authorization request. We also assume that there are some who made a request but were rejected, who we do not observe taking the restricted drug. We assume that there is a constant rejection rate r across all drugs and years. If we observe N patients taking the drug, with a rejection rate r , there will have been $\frac{N}{1-r}$ requests.²⁹ Therefore, the total administrative costs are $\frac{aN}{1-r}$.

We calibrate a and r from prior studies in the health policy literature, described below. For N , we simply use our demand system from Section 4 to estimate the number of beneficiaries consuming restricted drugs in the status quo simulation, summed across classes. We estimate that the average beneficiary, under the status quo, fills prescriptions for 0.299 unique restricted drugs per year across all classes in our demand estimates.

²⁷In theory, a prior authorization request can either be initiated by the provider prospectively when the drug is prescribed, or initiated retrospectively due to a patient facing an authorization barrier as in this example. A survey by CoverMyMeds (2020) finds that only 17% of authorization requests are prospective, with the other 83% retrospective.

²⁸In general, authorization is required once per treatment course, but this is heterogeneous across drugs and not documented in a systematic way. In practice, authorization may be required less or more than once a year.

²⁹We abstract from repeat interactions between the requesting physician and the insurer. Additionally, we are abstracting from real heterogeneity in both the costs and rejection rates associated with making requests for different kinds of drugs.

5.1 Calibrating Application Costs and Rejection Rates

When considering costs of prior authorization, there are two parties who incur costs for each authorization request: Medical providers, who need to submit requests, and insurers, who need to process and respond to them. We draw from case studies and industry reports to calibrate measures of each of these costs. In a systematic literature review, we found four studies that had estimated provider-side paperwork costs of prior authorization: Bukstein et al. (2006), Raper et al. (2010), CAQH (2013), and Carlisle et al. (2020). We describe the studies and their methods and estimates in Appendix Table A8. Their per-application estimates range from \$7.67 to \$27.35.³⁰ Our preferred estimate is from CAQH (2013), the study covering the largest number of providers. Their estimate is \$18.53 per application.³¹

We were only able to find one study that estimated insurer costs of fulfilling prior authorization requests, by CAQH (2013), who survey insurers. They estimated estimated manual processing costs of \$3.95 for insurers in 2012.³² Adding these insurer costs to the preceding estimates of provider costs, that gives us a range of total cost-per-application estimates from \$11.62 to \$31.30, with our preferred estimate being \$22.48, reflecting the two CAQH estimates. We also experiment with a handful of more extreme values: \$50, \$100, and \$200.

The literature provides many more estimates of prior authorization request rejection rates, although not all of them are directly comparable, and none precisely get at the exact quantity of interest—the number of (unobserved) requests per (observed) successful fill. Nonetheless, we take a handful of measures from this literature. We report the rates from the universe of studies we found in Appendix Table A9. Unfortunately, none of them are easily comparable to our setting. The studies are either too narrow in that they cover a single, potentially unrepresentative area of care, or too broad in that they include unrelated services (e.g. hospital services and physician-administered drugs). We use five values: 1.5%, 4%, 7.5%, and 15%, which cover the range of estimates found in the literature, as well as 0%.

5.2 Computing Net Financial Savings from Authorization Restrictions

With these calibrated values in hand, we can compute the total administrative burden generated by prior authorization. As in Section 4, we consider the burden generated by moving between the historical status quo and a counterfactual world in which prior authorization was removed but exclusion left intact. For every pair of calibrated values of a and r , we report in Table 7 the estimated total administrative costs from prior authorization per beneficiary-year. Subtracting the value here from \$96, our estimated spending reductions per beneficiary-year, measures the net financial savings from prior authorization restriction policies. Therefore, a value in this table below \$96 implies that prior authorization generates net financial savings, while a value above \$96 implies that it generates net financial losses. Note that our denominator is, as in the exercise from

³⁰ Another paper, Delate et al. (2005), does not measure administrative costs, but reports that a Medicaid program that institutes prior authorization policies for proton-pump inhibitors compensated providers by \$20 per request for their time, consistent with the magnitude of the estimates from the other studies.

³¹ We prefer their estimate for manually-submitted requests. In Appendix Table A8 we also report their estimate of costs for doing so through an IT system, but the majority of requests (110 million out of 130 million) were filed manually. Their cost estimates for manual filing decreased in later reports, with \$14.07 for calendar year 2013, \$7.17 for 2014, and \$7.50 for 2015.

³² Manual insurer-facing costs are stable across time in the CAQH survey and never exceed \$3.95 per request.

Section 4, all beneficiaries, so beneficiaries with many restricted drugs incur greater administrative costs, and those with no drugs incur zero costs.

Unsurprisingly, higher calibrated values of a and r increase the administrative burden and reduce the net financial savings from prior authorization. However, our estimates indicate that prior authorization policies generate *some* net financial savings unless per-application costs are at implausibly high levels such as \$200. Our preferred calibrated measures of a and r are given in bold in Table 7, reflecting the CAQH estimates of application costs and an intermediate rejection rate. At this calibration, administrative costs are \$9.76 per beneficiary-year. Our estimates thus suggest that the savings due to prior authorization are approximately 10 times larger than the paperwork costs. Ultimately, under this calibration we find that prior authorization policies generate net financial savings of approximately \$86 per beneficiary year.³³

We would expect the net financial savings to vary across classes. Measuring these savings in absolute terms is not easily comparable across classes since baseline spending in each class is so different. Instead, we construct the ratio of spending reductions to administrative costs. In Appendix Table A12 we provide this ratio for the set of all drugs; values above 1 imply net financial savings, while values below 1 imply net financial losses. For our class-level ratios, we use the calibration where $a = \$22.48$ and $r = 4\%$, for which this ratio is 10 for all drugs. We plot class-specific savings-to-admin-cost ratios in Figure 7, with 95% confidence intervals given by the black brackets and the red vertical line at the value of 1. For the majority of classes, we can reject that prior authorization generates net financial losses. The class with the largest (statistically significant) estimated savings per administrative dollar is the class of biologic response modifiers, a class where very few beneficiaries receive any drug at all, and where each individual drug is quite expensive, consistent with the type of class that our model in Section 1.2.1 predicts to be most well-suited for prior authorization.

Ultimately, these exercises indicate that prior authorization restrictions tend to generate financial savings vastly exceeding the associated administrative cost. This result is not trivially implied by revealed preference on behalf of the insurers. While we should not be surprised that insurers would institute policies that reduce their own *private* costs, there is no guarantee that the policies they institute would generate spending reductions that outweigh the administrative costs born both by themselves *and* external parties.

We also note that just because prior authorization restrictions resulted in net financial savings for the drugs selected by insurers for these restrictions, this does not imply that prior authorization would achieve similar savings for drugs not selected for restrictions. We explore this point explicitly in Appendix Tables A13 and A14 by replicating Tables 6 and 7 for a different counterfactual simulation exercise where we evaluate what would happen if we moved from the status quo to an alternative where all *unrestricted* drugs received prior authorization restrictions, holding the formulary status of previously-restricted and excluded drugs fixed. We find that, while this policy would indeed reduce drug spending considerably, under reasonable calibrations of a and r it no longer generates savings large enough to exceed the associated administrative costs. This result comes from the fact that many unrestricted drugs have large numbers of inframarginal consumers, generating significant administrative costs. While this exercise requires us to extrapolate far out-of-sample (many of these unrestricted drugs are never restricted and we thus have to assume that the

³³We omit standard errors to make the table easier to read. In Appendix Table A11 we generate this table with standard errors. Since a and r are not estimated, standard errors are similar across the cells.

effects of prior authorization on these drugs are similar to the effects on drugs observed to be restricted), we see it as an important demonstration of the idea that prior authorization policies generate net spending reductions *only if targeted appropriately*. This exercise also suggests that historically, prior authorization was indeed generally targeted well at the drug level.

6 Welfare Effects of Prior Authorization on Beneficiaries

Our results in Section 5 suggest that prior authorization restrictions generate net financial savings of roughly \$86 per beneficiary-year even when taking administrative costs born by inframarginal patients' providers and insurers into account. We can conclude that, in contrast to previous discourse about bureaucracy, the actual paperwork costs are second-order relative to the amount of spending moved around by these policies.

This leads us to conclude that the first-order effects of prior authorization restrictions are instead the actual effects on utilization of drugs. The stated goal of prior authorization is to reduce moral hazard, decreasing the use of drugs where beneficiary valuation falls below the cost of that utilization. Thus, if there is substantial moral hazard, this first-order effect of prior authorization may be efficient. However, if there is no moral hazard (i.e., the value of all drug utilization exceeds the social cost of that utilization), then these policies are pure waste: They burn an (albeit relatively small) amount of administrative effort *and*, even worse, move beneficiaries away from drugs they value highly to clinical alternatives that they value less, or to no drug at all.

In this final section we thus attempt to estimate the (incremental) value of the drug consumption that is marginal to prior authorization restrictions relative to the next alternative. We present two exercises to assess this value. First, we present a revealed preference exercise, in which we calculate beneficiary willingness-to-pay for drug consumption, estimated from the demand response to an out-of-pocket price change. This approach relies on strong assumptions regarding the mapping between willingness-to-pay and private value, but we do our best to show the extent to which our conclusions are robust to the relaxation of those assumptions. Second, we estimate the effects of prior authorization on beneficiary health. We do this both via an aggregate analysis as well as via a case study of a specific drug class, oral anticoagulants.

6.1 Revealed Preference Approach

Our goal is to measure the total loss in consumer surplus due to beneficiaries being turned away from a restricted drug due to prior authorization. To do this, we start by estimating beneficiaries' willingness-to-pay for a drug. Returning to the model from Section 1.2, again let $V_d(\theta_{id})$ represent the incremental valuation that a beneficiary of type θ_{id} has for drug d . Let $\theta_{id} \in [0, 1]$ be an ordering such that $V_d(\theta)$ is non-increasing in θ and $V_d(0) = \max_i \Delta v_{id}$. What we want to estimate is the total consumer surplus loss for restricted drugs:

$$\Delta CS_d = - \int_{\Theta_M} V_d(\theta) d\theta$$

with Θ_M denoting the set of marginal patients who are deterred from the drug as a result of the restrictions. This expression thus represents the cumulative incremental value of the restricted drug, relative to the next

best alternative, for marginal beneficiaries who receive the drug when prior authorization is not in place, but do not receive the drug when prior authorization is in place. We estimate this by inferring beneficiaries' valuation of drugs typically under prior authorization restrictions from their demand response to price changes for those drugs.

Let $W_d(\theta_{id})$ denote the willingness to pay for drug d for θ -type beneficiaries, the maximum out-of-pocket price at which they would purchase the drug. Then, we can define the market demand curve, $D_d(P_d) = \int 1\{W_d(\theta) \geq P_d\}d\theta$, defined as the share of beneficiaries with willingness to pay at or above the out-of-pocket price P_d . If we have exogenous variation in P_d , we can use it to trace out $D_d(\cdot)$ and, with it, the distribution of W_d .

If willingness to pay reveals beneficiary valuation for a drug, $W_d(\theta) = V_d(\theta)$, $D_d(\cdot)$ will provide a picture of not just the distribution of W_d but also the distribution of V_d . Then, given knowledge of which θ -types were deterred by prior authorization, we can estimate the loss in consumer surplus. While assuming that willingness to pay is equal to consumer value is a strong assumption that rules out phenomena like behavioral hazard (Baicker et al. 2015) and liquidity constraints (Gross et al. forthcoming), such an approach is nonetheless commonly invoked in order to estimate consumer valuation (Einav et al. 2010, Lieber and Lockwood 2019).

To trace out $D_d(P_d)$, we rely on a separate natural experiment originally used by Gross et al. (forthcoming). As discussed in Section 2.1, the LIS program heavily subsidizes out-of-pocket costs for prescription drugs. Thus, when beneficiaries enter this program, they experience a large decrease in the price they pay for their prescriptions due to the cost-sharing subsidy. We leverage the transitions of 62,785 beneficiaries into the LIS program as a source of exogenous variation in drug prices and estimate the demand response to that variation in prices. In Appendix Table A16, we provide summary statistics for this population.

We observe each of these beneficiaries one year prior to their transition, the year of transition and one year post-transition. We estimate the following regression at the person-drug-year level:

$$\log(\mathbb{E}[Y_{idt}]) = \frac{\epsilon}{100} P_{dt} \times \text{NotLIS}_{it} + \alpha_i + \gamma_{dmt} + \epsilon_{it} \quad (3)$$

The interaction term $P_{dt} \times \text{NotLIS}_{it}$ is equal to the price of the drug d in year t for those not yet enrolled in the LIS program ($\text{NotLIS}_{it} = 1$) and zero for those enrolled ($\text{NotLIS}_{it} = 0$). We measure the price P_{dt} as the average copayment paid per year across all (non-LIS) beneficiaries observed using the drug in year t ; that is, we compute the 'price' as the expense of taking a drug for a year. Its coefficient captures the demand response to the change in price that occurs due to gaining the LIS cost-sharing subsidy. Because the outcome is in log terms, the coefficient estimates the price semi-elasticity of demand, ϵ , divided by 100. Note that because the prices used are fixed across beneficiaries, we only use the variation that comes from the transition into the LIS program and not any across-plan differences in copayments for a given drug.

This regression estimates a weighted average treatment effect of price changes across drugs, with each drug weighted by its variance in the outcome $P_{dt} \times \text{NotLIS}_{it}$. However, these weights will not necessarily match the implicit weights derived from the effects of prior authorization on utilization. Therefore, we weight observations by $w_{idt} = \frac{\text{Var}[\text{Auth}^{\text{Assigned}}|d, m(it)]}{\text{Var}[P_{dt} \times \text{NotLIS}_{it}]}$, such that the drug-specific weights are equal across the two regressions. With this, we can interpret ϵ as the weighted-average price semi-elasticity of demand,

weighted such that it is directly comparable to the effect of prior authorization on utilization.³⁴

For Y_{idt} , our measure of utilization, we mimic our approach in Sections 3 and 4 and use a binary indicator for whether the beneficiary ever filled the prescription. Because this outcome frequently takes on values of zero, we estimate Equation 3 using Poisson regression (Wooldridge 1999). Table 8 shows the results from this regression. The appropriately weighted semi-elasticity of demand with respect to the out-of-pocket expense is approximately 0.51 when estimated without beneficiary fixed effects, and 0.15 when estimated with beneficiary fixed effects. Prior authorization on a focal drug deters 28.9% of consumption of restricted drugs (as per Table 6). This implies that imposing prior authorization on a drug is equivalent to raising annual out-of-pocket costs for the drug by approximately \$227.³⁵ To put this in context, transitioning to the LIS program lowers the weighted average copayment for the restricted drugs by \$49 per drug fill or \$265 per drug-year, implying that prior authorization has almost the same effect on drug-specific utilization of the subsidies provided by the LIS program.

To compute the loss in consumer surplus, we need a full demand curve, not simply an elasticity. We assume that the demand curve for drug d is given by $D_d(P_d) = D_d(0)e^{\frac{\epsilon}{100}P_d}$, where $D_d(0)$ is the share of beneficiaries who consume the drug when the price is zero, and ϵ is the price semi-elasticity of demand. This is a constant semi-elasticity demand function, and is equivalent to assuming that the cumulative density function of W_d is

$$F_d(W) = (1 - D_d(0)) + D_d(0)(1 - e^{\frac{\epsilon}{100}W}) = 1 - D_d(0)e^{\frac{\epsilon}{100}W}$$

for $W \in [0, \infty)$. That is, the distribution of W has a mass point at zero with $1 - D_d(0)$ share of beneficiaries having a willingness-to-pay of zero or less, and, for the $D_d(0)$ share of beneficiaries who have willingness to pay at or above zero, their willingness to pay is exponentially distributed with scale parameter $\frac{\epsilon}{100}$. With this structure, the semi-elasticity and a measure of $D_d(0)$ are sufficient to identify the full demand curve and distribution of willingness to pay. $D_d(0)$ is defined conditional on a prior authorization regime. We assume that the willingness to pay which reveals value is the one revealed under no prior authorization restrictions. Since our data include observations of beneficiaries who do face restrictions, we cannot simply compute $D_d(0)$ from data on the LIS population. Instead, to compute it, we use our simulations in Section 4 to estimate the demand for each drug in the absence of prior authorization restrictions.

Finally, to estimate the average willingness-to-pay of the forgone consumption under prior authorization restrictions, we have to define Θ_M , the set of types corresponding to marginal beneficiaries. 28.9% of beneficiaries who would consume the drug in the absence of prior authorization are deterred from it when prior authorization applies. However, a question remains: Where are these 28.9% located on the demand curve with respect to price? Since we don't observe θ_{id} , we make two assumptions. First, we assume

³⁴There are two complications which potentially bias our estimates of β . First, the LIS transition also lowers the price of potential substitute drugs. This shifts the demand curve for d to the left, lowering the quantity of d demanded and deflating our estimates of price response. Second, the LIS transition may be contemporaneous with an income decrease. If prescription drugs are a normal good, this will shift demand to the left, causing us to further underestimate the response to prices, although this may be small if changes that trigger LIS eligibility are small, and/or if LIS enrollment is triggered by information about eligibility rather than income changes. Since estimated consumer surplus loss is inversely proportional to the demand elasticity, both effects will cause us to overestimate the consumer surplus loss.

³⁵This calculation assumes that demand for the drug is $D_d(P_d) = D_d(0)e^{\frac{\epsilon}{100}P_d}$, where $D_d(P_d)/D_d(0) = 1 - 0.289$ and $\epsilon = -0.15$

the best-case scenario: That the marginal beneficiaries screened away from the drug are those with the lowest willingness-to-pay for the drug (and thus the highest θ types of those who take the drug). Under this assumption, the total amount consumers are willing to pay for the forgone consumption is equal to

$$\Delta CS_d^{\text{Best-Case}} = - \int_{(1-0.289)D_d(0)}^{D_d(0)} D_d^{-1}(\theta) d\theta$$

This is the ideal case for prior authorization to be potentially welfare-improving. However, price and prior authorization may screen out different beneficiaries. Further, if willingness-to-pay reflects value, in Section 3.4 we find suggestive evidence that prior authorization restrictions may screen on factors likely unrelated to patient value. Thus, we consider an alternative approach. In this approach, we assume screening is random with respect to willingness-to-pay, i.e., that a random 28.9% of beneficiaries lose access. The associated total amount consumers are willing to pay for the forgone consumption can be computed as

$$\Delta CS_d^{\text{Random}} = -0.289 \int_0^{D_d(0)} D_d^{-1}(\theta) d\theta$$

i.e., willingness-to-pay for the forgone consumption is just 28.9% of the total amount all consumers are willing to pay for the drug.

We present the total amount consumers are willing to pay for the forgone consumption computed across all drugs, per beneficiary-year, in Table 9.³⁶ The two columns represent the amounts consumers are willing to pay under the two estimates of the semi-elasticity of demand. Our preferred estimate for the semi-elasticity is the one that incorporates beneficiary fixed effects. Under best-case screening, consumers are willing to pay little for the forgone consumption, only \$13 per beneficiary-year. Under random screening, consumers are willing to pay more: \$81 per beneficiary-year.³⁷

These estimates provide a benchmark by which we can evaluate the savings. In the perfect screening case, the amount consumers are willing to pay for the forgone drugs is around 15% of the net savings induced by prior authorization restrictions. In the random screening case, the amount consumers are willing to pay for the forgone consumption is in the same general range as the net savings. Thus, as long as prior authorization screens people randomly or better with respect to willingness-to-pay, consumer willingness-to-pay falls below net savings. However, if screening is worse than random, willingness-to-pay will often exceed savings.

While a comparison of consumer willingness-to-pay to net savings is a useful benchmark, what we are really interested in is a comparison of *lost consumer surplus* relative to net savings. This requires us to map from willingness-to-pay to consumer value. If we take the standard approach of assuming equivalence

³⁶Readers may notice that this requires us to aggregate across drugs. In Appendix E we show that ΔCS_d is linear in $D_d(0)$, therefore, we can replace $D_d(0)$ with $\sum_d D_d(0)$ to compute the total lost consumer surplus across all drugs. Aggregation presents a second complication: In classes with multiple restricted drugs, the composite next-best alternative to a restricted drug may include the use of another restricted drug. In this case, it may therefore be true that for some beneficiaries, the full set of restrictions will move them to their third-most-preferred option rather than their second-most-preferred. Accounting for this case would require us to estimate the joint distribution of valuations across drugs within a therapeutic class, so we assume this case away.

³⁷We also consider the possibility that screening is worse than random as in Deshpande and Li (2019). In the worst-case scenario where the forgone drug consumption comes from the consumers with highest willingness-to-pay, the total amount consumers are willing to pay for the forgone consumption would be $\Delta CS_d^{\text{Worst-Case}} = - \int_0^{0.289D_d(0)} D_d^{-1}(\theta) d\theta$, or \$181 per beneficiary-year, around two times the net savings due to prior authorization.

between willingness-to-pay and value, then the results above clearly imply large welfare gains from prior authorization in the perfect screening case and relatively modest welfare gains in the random screening case. Therefore, in these two cases, this exercise suggests that there is enough moral hazard among these beneficiaries such that the reductions in utilization are Kaldor-Hicks efficient, even after considering the administrative costs required to achieve these utilization reductions. Importantly, this is still true, albeit marginally, even when the screening is done *at random*. However, unsurprisingly, this is not true when screening is worse than random.

A caveat is that this approach takes willingness to pay seriously as a signal of beneficiary value, despite the fact that prior work has shown evidence of behavioral biases in demand for health care. While we cannot estimate a specific bias level in our setting, we can bound how large the wedge between value and willingness to pay must be to overturn our results. We assume that $W_d(\theta) = \rho V_d(\theta)$, i.e., a θ -type consumer is willing to pay at most $(100 \times \rho)\%$ of their true valuation, due to a bias or another constraint. In Appendix E, we show that the ‘debiased’ consumer surplus measure is $\Delta CS^{\text{Debiased}} = \frac{1}{\rho} \Delta CS^{\text{Estimated}}$, so the value of ρ required to make prior authorization inefficient is $\rho^* = \frac{\Delta CS^{\text{Estimated}}}{NFS}$. For the best-case and random scenarios, ρ^* is 0.15 and 0.93 respectively. That is, if screening is perfect, patients need to value their drugs at approximately 7 their WTP for prior authorization to be inefficient, whereas if screening is random, patients need to value drugs 7.5% more than their WTP. (In the worst case, prior authorization is always inefficient)

Another caveat is that the approach here focuses on patient incentives; again, there is substantial evidence that patients have limited *control* over treatment decisions, with control instead passed to their provider, their agent in medical decisions. Instead, one can think about a simple model of provider agency with altruism where providers unilaterally choose drugs for their patients trading off patient benefit against administrative costs. In Appendix E.2 we walk through this model. The fact that use of restricted drugs is reduced by 28.9% in response to provider admin costs that are \$22 implies a semi-elasticity of roughly -1.29, implying much smaller consumer surplus losses than our patient-based approach: In the best-case and random screening scenarios, the surplus losses are \$2 and \$9, respectively. Thus, in order to conclude that there are large consumer surplus losses in a world where the response is largely about provider decision-making would require providers to put extremely low weight on patient value relative to their own administrative costs.

Ultimately, we interpret the results from this revealed preference approach as suggestive evidence that the lost consumer surplus may be sufficiently low relative to the cost of care to indicate substantial moral hazard in this setting. This moral hazard motivates prior authorization restrictions as a potentially efficiency-enhancing rationing device.

6.2 Health Effects

Next, we investigate the effects of prior authorization on patient health. The American Medical Association’s public petitions on prior authorization make strong claims about potential health harms, claiming restrictions might lead to “hospitalization, disability, or even death,” although these claims are based on a survey of physicians’ opinions rather than measured outcomes.³⁸ While any revealed preference analysis may be

³⁸See <https://fixpriorauth.org/patients>.

misleading due to various behavioral and agency frictions, if there are negative effects of prior authorization on patient health, then there are obviously important losses in consumer surplus. Our primary research design from Section 3 does not permit estimation of health effects, as health is defined at the patient level, whereas our research design induces variation at the patient-drug level. To estimate health effects, we either need drug-specific measures of health, or an alternative design that induces variation at the patient level. We attempt both approaches.

6.2.1 Health Effects Part 1: Case Study of Oral Anticoagulants

A drug-specific approach to measuring health effects of prior authorization requires a setting where (1) prior authorization is sufficiently frequent (but not universal), (2) use of the drug is common, and (3) it is plausible that the restricted drug has important observable short-run clinical effects on patient health. Requirement (1) rules out common drugs like statins and beta blockers, where the prior authorization restrictions are rare. Requirement (2) rules out some of the most commonly restricted drug classes, like drugs that treat cancer, as these are not widely used. Requirement (3) rules out drugs used to treat multiple sclerosis or mental health conditions, as health effects are unlikely to be measurable in the short run.

We focus on a class that satisfies all three requirements: Oral anticoagulants, commonly referred to as blood thinners. Anticoagulants reduce the extent of blood clotting, therefore reducing the risk of strokes (blood clots that occur in the brain), as well as other clot-driven health events such as heart attacks and pulmonary embolisms. They are typically taken over a long period of time, generally for many years. The standard oral anticoagulant that was prescribed until the 2010s was warfarin, which, by the beginning of our sample period, had existed primarily as a low-priced generic for decades, costing approximately \$0.30 per pill.

In the 2010s, however, a series of drugs called non-Vitamin K oral anticoagulants (NOACs) were approved by the FDA and introduced into usage. There were two main advantages of these new drugs over warfarin: (1) the required dose varies less across patients and over time, so there is less need for frequent monitoring of blood clotting, and (2) there are fewer potentially dangerous interactions with foods and other drugs. By 2015, NOACs represented around one-eighth of anticoagulant prescriptions, but *two-thirds* of spending (see Appendix Figures A6 and A7). Total anticoagulation spending rose substantially over this period.

For this analysis, we restrict to a subsample of individuals with a medical history of atrial fibrillation, deep vein thrombosis, or pulmonary embolism, the typical conditions treated by anticoagulants. We report summary statistics for this sample of beneficiaries in Table A17. 29.1% of beneficiary-years in the subsample fill a prescription for any oral anticoagulant, reflecting 84.9% of all oral anticoagulant use within our broader LIS sample.

As Table A17 shows, most beneficiaries are enrolled in plans which either put authorization restrictions on all available NOACs, or cover them without restriction.³⁹ To simplify our analysis, we therefore focus

³⁹Our formulary data tends to list the formulary status for newly-introduced drugs with a lag of a year. However, we know that such drugs were available to beneficiaries, since we observe them being consumed. For this analysis, we ‘backfill’ formularies: If a drug does not appear in a plan’s formulary data but does appear in claims data for some plan that year, we assign the plan its formulary status for the first year the drug appeared in *any* formulary as covered.

on comparing these two cases to measure the health effects of reduced access to NOACs broadly. For our analysis, we run regressions of the form

$$Y_{it} = \beta \text{AuthAllNOACs}_{j(it)t} + \gamma \text{OtherFormulary}_{j(it)t} + \delta_{m(it)} + \epsilon_{it}$$

where $\text{AuthAllNOACs}_{j(it)t}$ is a dummy variable indicating whether i was assigned to a plan j where *all* NOACs were restricted in year t , and $\text{OtherFormulary}_{j(it)t}$ is a dummy variable indicating whether i was assigned to a plan with any formulary other than formularies fully restricting all NOACs or formularies with no restrictions on NOACs (the omitted group). Since the omitted group is plans where *no* NOACs were restricted, β thus measures the difference between restricted access to NOACs compared to free access. All plans cover generic warfarin without restriction during this period. In this specification, we do not include plan fixed effects, as we estimate this regression for a single set of drugs, thus resulting in plan identifiers being co-linear with the treatment dummies.

We begin by replicating our utilization results for this specific drug class. To do so, we estimate the regression above with the following outcomes: total spending on anticoagulants and dummy variables for whether the beneficiary ever filled any anticoagulant, warfarin, or any NOAC during the year. We report the estimated coefficients in Table 10. Prior authorization on NOACs significantly reduced overall spending on anticoagulants, with a decrease of \$16.60 per beneficiary year or about 15%. This spending reduction was largely driven by a shift of 22.6% of beneficiaries who take NOACs in the plans that don't restrict them away from NOACs. However, we find that there was no overall effect on the probability of taking *any* anticoagulant, indicating that essentially all of those deterred by prior authorization restrictions from taking NOACs substituted to the much cheaper generic warfarin.⁴⁰

These results already suggest that large health effects are unlikely in this case, as they would have to come via the incremental health benefit of the NOACs versus warfarin rather than the absolute benefit of taking an anticoagulant. To investigate further, we estimate the effect of prior authorization restrictions on an indicator for negative anticoagulant-relevant health events during the year: Stroke, anticoagulant-related bleeding, and death. We report the results from these regressions in Table 11.⁴¹ Our point estimates are modest, with effect sizes of roughly 2% for strokes and death, and -2% for bleeding. Unfortunately, these estimates are also extremely noisy, not allowing us to reject effect sizes between 11% and -7% for strokes, between 4% and -9% for bleeding, and between 10% and -15% for death. We thus learn very little from this exercise.

For oral anticoagulants, prior authorization thus delivers large spending reductions, with those reductions coming entirely from shifting patients from expensive new branded drugs to an older, much cheaper generic substitute. There could be some health consequences of this shift, but our estimates of these consequences are too noisy to make any strong conclusions. Despite the disappointing results on health effects, we include this analysis here to highlight the difficulty of evaluating the health effects of such policies even in large samples with exogenous variation.

⁴⁰The NOAC and warfarin 'any use' coefficients do not perfectly offset each other in this case because some beneficiaries took both NOACs and Warfarin during a given year. This most likely represents beneficiaries who start on one and move to the other later in the year.

⁴¹We do not perform IV regressions here, as these regressions would produce even noisier estimates of health effects.

6.2.2 Health Effects Part 2: Aggregate Health Effects of Prior Authorization

Given that a focused case study provided inconclusive results regarding health effects of prior authorization, we next attempt a more aggregate analysis. To do so, we construct a beneficiary-level measure of exposure to prior authorization aggregated across classes. We do so by following in the spirit of [Brot-Goldberg et al. \(2021\)](#) and constructing a measure of formulary ‘fit,’ where exposure to prior authorization restrictions is measured in terms of how often it applies to the set of drugs previously taken by the beneficiary. We construct a measure of exposure to prior authorization by measuring the share of drugs that the beneficiary filled at least once in the prior year which would face prior authorization in the plan the beneficiary was subsequently assigned to, and for any other benchmark plan that they could have been randomly assigned to. We construct the same measure for formulary exclusion.

Regressing beneficiary-specific outcomes directly on this exposure measure may be confounded by the fact that this measure may be higher for beneficiaries who take niche drugs, who may also be in poorer health. We continue to follow [Brot-Goldberg et al. \(2021\)](#) by constructing a rank-based measure of exposure, where, for each beneficiary, we order each benchmark plan that the beneficiary could have been assigned to in terms of the prior authorization exposure measure, and bin the plans into quintiles of beneficiary-specific exposure. We then generate an indicator for whether the beneficiary’s assigned plan was in the quintile with the most exposure to prior authorization, and call this indicator AuthExposureQ5_i . Unlike the raw ‘fit’ measure, this measure is orthogonal to beneficiary heterogeneity by construction since the probability of assignment is always $\frac{1}{5}$ for every single beneficiary.

To estimate the effects of being assigned to a high-exposure plan on beneficiary health, we run a set of regressions of the form

$$Y_i = \beta \text{AuthExposureQ5}_i + \gamma \text{ExclExposureQ5}_i + \delta_{m(i)} + \epsilon_i$$

β thus represents the health consequences of greater exposure to prior authorization. We hold exposure to drug exclusion fixed by constructing a measure ExclExposureQ5_i of assignment to the worst plans in terms of their exclusion of previously-taken drugs.

We first show that this measure does indeed predict reductions in utilization. In the first column of Table 12, we examine the effect on total drug spending. In line with our results in Section 4, greater exposure to prior authorization does indeed lower spending. In the second column, we estimate the effects of greater prior authorization exposure on the probability that the beneficiary dies during the year (multiplied by 100 so that our measure represents whole percentage point changes). Greater prior authorization exposure is estimated to raise current-year mortality by 0.06 percentage points, roughly a 2.6% increase over baseline mortality. This is large; however, our standard errors are even larger, and we cannot rule out an 8.2% increase in mortality, nor a 3.1% decrease. We also measure utilization of non-drug medical care; as [Chandra et al. \(2010\)](#) point out, reductions in the use of valuable drugs can generate offset effects by worsening patient health. We measure total spending on inpatient hospitalizations, and total spending on all non-drug medical care. We estimate that both increase, but the standard errors are large, and we cannot reject zero or even negative effects of prior authorization on these forms of care.

In the second panel of Table 12, we again borrow from [Brot-Goldberg et al. \(2021\)](#) and restrict to a subset

of beneficiaries who face the greatest variation across plans in terms of exposure to prior authorization. We measure this in terms of the variance of the exposure measure across benchmark plans, and restrict to the top 25% of beneficiaries in the sample by this variance measure. We find that for these beneficiaries, the spending effects are predictably much larger, by a factor of three. However, the consequences for health outcomes and non-drug spending are not worse, and instead flip sign completely (e.g., we now estimate that greater exposure to prior authorization reduces mortality and lowers non-drug spending), although the estimates are even noisier than before.

Ultimately, this aggregate analysis also proves to be inconclusive regarding the health consequences of prior authorization. We cannot be overly definitive about the results from either of the two exercises: Neither set of results shows that prior authorization necessarily has *no* effect on patient health. Rather, even with substantial data and large, well-powered effects on quantities, we do not have the statistical power to pin down precise effects on health outcomes.

7 Conclusion

Our results suggest that prior authorization restrictions are a powerful tool for reducing health care costs. As highlighted by the American Medical Association and other interest groups, these restrictions do also generate substantial administrative costs. However, even under generous assumptions, these administrative costs are small relative to the reductions in drug spending achieved by these restrictions. Additionally, the administrative costs of prior authorization have decreased over time, as estimated by the CAQH.

Our results thus indicate that the first-order effect of prior authorization is *not* wasteful spending on bureaucratic sludge associated with the authorization process. Instead, the first-order effects are on drug utilization, with around one-quarter of individuals preferring restricted drugs being deterred from taking those drugs by prior authorization restrictions. However, not all of those deterred consumers opt out of taking *any* drug for the particular condition they desire to treat; indeed, around half of these marginal consumers opt for a (typically much cheaper) clinical substitute instead.

Overall, the welfare consequences of these policies, as implied by our results, are two-handed. On the one hand, our revealed preference approaches suggest that, under plausible assumptions, the value that patients put on restricted drugs falls below the cost of acquiring those drugs. Concluding the opposite requires the presence of substantial biases in decision-making; or that the marginal patients have higher valuations for drugs than the average patient. However, while our estimates of valuation fall below the cost to insurers of financially procuring drugs, they surely exceed the marginal cost of producing the drugs, making it unclear whether this consumption is inefficient. Further, we find that the quantity reductions due to prior authorization may be inefficiently targeted towards sicker patients and non-white patients, meaning that prior authorization policies may reduce equity and thus not be socially desirable. Moreover, we are not able to conclusively estimate the effects on patient health. Finally, we are unable to quantify any of the effects of prior authorization on *patient* administrative hassle, as opposed to provider hassle. Our results suggest a potential positive role for prior authorization on social welfare, although more research is needed to understand these missing pieces.

While prior authorization appears to be an effective policy as used in practice, this does not imply that

it would be an effective policy if implemented more widely. Insurers in our setting restrict drugs in the way that our model suggests may be most socially efficient, primarily targeting niche branded drugs with few inframarginal users and high prices. An expansion of authorization restrictions to other drugs, especially those with many inframarginal users and low prices, could easily be inefficient, generating substantial administrative burden for little value. Furthermore, it need not be true that the current administration of these policies is optimal, even if these policies are beneficial relative to no alternative. The use of prior authorization as a tool for information transfer may be inefficient relative to other technological solutions, such as giving payers access to patient medical records electronically ([Cutler 2020b](#)). The extent to which prior authorization serves as a tool for formal information transfer, versus serving as a device to allow physicians to signal private information, will be an important topic for future research.

Our results motivate three additional broader points. First, although managed care policies can improve social welfare, they also *raise* costs for physician and other health care providers, by increasing their paperwork burdens. These policies are Kaldor-Hicks efficient in the sense that providers *could* be transferred a portion of the savings to be made at least indifferent between being the stewards of these policies and not. In practice, however, no direct transfers are apparently made (although implicit transfers may occur through changes in reimbursement rates), and so the gains are primarily realized by payers. In our setting, these payers are largely drug insurers who have no direct contractual relationship with providers through which a transfer could occur. Finding a way to efficiently share the gains with providers is a serious political economy issue. The AMA has internally proposed developing billing codes to allow providers to bill insurers for time spent on paperwork ([Frieden 2022](#)). This could allow for some sharing of the savings. However, a per-application reimbursement could also make the signal of the value of the prescription sent by the physician's willingness to go through the prior authorization process much less informative, by reducing the net cost of an application. Thus, more creative contractual forms of sharing gains with providers may be necessary and their design is a fruitful area for future research.

Second, our results speak to the choice of rationing mechanisms within the U.S. health care system. The primary mechanism for allocating care in the U.S. is a patient-price-based market mechanism, and screening out low-value care is done on the basis of willingness of patients to pay. However, our results suggest that, to deter the same amount of care, an insurer can either charge a patient \$227 per year in copayments, or induce a provider to spend \$22 in administrative costs. In a way, bureaucratic restrictions may be a 'cheaper' way to restrict costs compared to greater cost-sharing, although the administrative costs require real effort (and therefore deadweight loss) rather than Kaldor-Hicks-neutral transfers.

Finally, our results have important implications for the broader discourse around international health care spending comparisons and U.S. health care reform. Non-price rationing in U.S. health care is primarily done formally through managed care policies, which generate administrative costs on accounting balance sheets since they are paid through administrative salaries. In contrast, queue-based rationing mechanisms, used more frequently in other OECD health care systems, also generate waste by forcing patients to wait, but these costs are not captured in formal cost accounting. More research is needed to characterize the relative costs and benefits of other sources of administrative cost burden, as well as to compare how other rationing mechanisms induce hassle costs, both those that show up in accounting data and those that do not.

References

- Alpert, Abby E., Sarah E. Dykstra, and Mireille Jacobson**, “How Do Prescription Drug Monitoring Programs Reduce Opioid Prescribing? The Role of Hassle Costs versus Information,” 2020. NBER Working Paper No. 27584. 4
- American Medical Association**, “2017 AMA Prior Authorization Physician Survey,” 2017. 1, 5
- America’s Health Insurance Providers**, “Key Results of Industry Survey on Prior Authorization,” 2020. <https://www.ahip.org/wp-content/uploads/Prior-Authorization-Survey-Results.pdf>. 5
- Archer, Diane**, “Medicare Is More Efficient Than Private Insurance,” *Health Affairs Blog*, 2011. <https://www.healthaffairs.org/doi/10.1377/hblog20110920.013390/full/>. 1
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein**, “Behavioral Hazard in Health Insurance,” *Quarterly Journal of Economics*, 2015, 130, 1623–1667. 3, 4, 7, 26
- Bajari, Patrick, Stephanie Houghton, and Steven Tadelis**, “Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs,” *American Economic Review*, 2014, 104 (4), 1288–1319. 4
- Berry, Steven T.**, “Estimating Discrete-Choice Models of Product Differentiation,” *RAND Journal of Economics*, 1994, pp. 242–262. 20, A26, A27
- Birdsall, Ashley D., Ashley M. Kappenman, Bryce T. Covey, and Matthew H. Rim**, “Implementation and impact assessment of integrated electronic prior authorization in an academic health system,” *Journal of the American Pharmacists Association*, 2020, 60 (4), e93–e99. A16
- Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad**, “What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics,” *Quarterly Journal of Economics*, 2017, 132 (3), 1261–1318. 4, 22
- **and Mathijs de Vaan**, “Intermediation and Vertical Integration in the Market for Surgeons,” 2018. 19
- **, Timothy Layton, Boris Vabson, and Adelina Yanyue Wang**, “The Behavioral Foundations of Default Effects: Theory and Evidence from Medicare Part D,” 2021. NBER Working Paper No. 28331. 2, 9, 13, 15, 18, 32
- Bukstein, Don A., Greeta A. Cherayil, Adam D. Gepner, Allan T. Luskin, J. Brent Kooistra, and Reid M. Olson**, “The Economic Burden Associated with Prior Authorizations in an Allergist Office,” *Allergy & Asthma Proceedings*, 2006, 27 (2), 119–122. 23, A15
- Carlisle, Ryan P., Nicholas D. Flint, Zachary H. Hopkins, Mark J. Eliason, Kristina C. Duffin, and Aaron M. Secrest**, “Administrative Burden and Costs of Prior Authorizations in a Dermatology Department,” *JAMA Dermatology*, 2020, 156 (10), 1074–1078. 23, A15, A16
- Casalino, Lawrence P., Sean Nicholson, David N. Gans, Terry Hammons, Dante Morra, Theodore Karrison, and Wendy Levinson**, “What Does It Cost Physician Practices To Interact With Health Insurance Plans?,” *Health Affairs*, 2009, 28 (4), w533–w543. 1, 4
- Chandra, Amitabh, Evan Flack, and Ziad Obermeyer**, “The Health Costs of Cost-Sharing,” 2021. NBER Working Paper No. 28439. 3, 4

- , **Jonathan Gruber, and Robin McKnight**, “Patient Cost-Sharing and Hospitalization Offsets in the Elderly,” *American Economic Review*, 2010, 100 (1), 193–213. 32
- Chernew, Michael and Harrison Mintz**, “Administrative Expenses in the US Health Care System: Why So High?,” *Journal of the American Medical Association*, 11 2021, 326 (17), 1679–1680. 1
- Correia, Sergio, Paulo Guimarães, and Thomas Zylkin**, “Fast Poisson Estimation with High-Dimensional Fixed Effects,” *Review of Economics and Statistics*, 2020, 20 (1), 95–115. 21, A31
- Council for Affordable Quality Healthcare**, “2013 U.S. Healthcare Efficiency Index: Electronic Administrative Transaction Adoption and Savings,” 2014. 23, A15
- CoverMyMeds**, “2020 Electronic Prior Authorization Report,” 2020. <https://www.covermymeds.com/main/medication-access-report/electronic-prior-authorization/>. Retrieved January 9, 2023. 22
- Cutler, David M.**, “Reducing Administrative Costs in U.S. Health Care,” 2020. The Hamilton Project. 1
- , “Taming The Paper Tiger,” 2020. *Health Affairs Blog*, <https://www.healthaffairs.org/doi/10.1377/forefront.20200929.284683>. 34
- **and Dan P. Ly**, “The (Paper)Work of Medicine: Understanding International Medical Costs,” *Journal of Economic Perspectives*, 2011, 25 (2), 3–25. 1
- , **Elizabeth Wikler, and Peter Basch**, “Reducing Administrative Costs and Improving the Health Care System,” *New England Journal of Medicine*, 2012, pp. 1875–1878. 1, 4
- Delate, Thomas, Douglas E. Mager, Jagat Sheth, and Brenda R. Motheral**, “Clinical and Financial Outcomes Associated With a Proton Pump Inhibitor Prior-Authorization Program in a Medicaid Population,” *American Journal of Managed Care*, 2005, 11 (1), 29–36. 23, A16
- Deshpande, Manasi and Yue Li**, “Who Is Screened Out? Application Costs and the Targeting of Disability Programs,” *American Economic Journal: Economic Policy*, 2019, 11 (4), 213–248. 4, 28
- Dillender, Marcus**, “What Happens When the Insurer Can Say No? Assessing Prior Authorization as a Tool to Prevent High-Risk Prescriptions and to Lower Costs,” *Journal of Public Economics*, 2018, 165 (1), 170–200. 4
- Drum, Kevin**, “Join Me On a Dive Down the Rabbit Hole of Health Care Admin Costs,” 2019. *Mother Jones*, <https://www.motherjones.com/kevin-drum/2019/06/join-me-on-a-dive-down-the-rabbit-hole-of-health-care-admin-costs>. 1
- Dunn, Abe, Joshua D. Gottlieb, Adam Shapiro, Daniel J. Sonnenstuhl, and Pietro Tebaldi**, “A Denial a Day Keeps the Doctor Away,” 2021. NBER Working Paper No. 29010. 4, 5, 22
- , — , **and Adam Hale Shapiro**, “Administration Above Administrators: The Changing Technology of Health Care Management,” *AEA Papers and Proceedings*, 2020, 110, 274–278. 1, 4
- Einav, Liran, Amy Finkelstein, and Mark Cullen**, “Estimating Welfare in Insurance Markets Using Variation in Prices,” *The Quarterly Journal of Economics*, 2010, 125 (3), 877–921. 26
- Eliason, Paul, Riley League, Jetson Leder-Luis, Ryan C. McDevitt, and James W. Roberts**, “Ambulance Taxis: The Impact of Regulation and Litigation on Health Care Fraud,” 2021. NBER Working Paper No. 29491. 4

- Finkelstein, Amy and Matthew J. Notowidigdo**, “Take-up and Targeting: Experimental Evidence from SNAP,” *The Quarterly Journal of Economics*, 2019, 134 (3), 1505–1556. 4, 6
- Frakt, Austin**, “Is Medicare for All the Answer to Sky-High Administrative Costs?,” *The New York Times*, 2018. 1
- Frieden, Joyce**, “Should Docs Get Paid for Time Spent on Prior Authorizations?,” 2022. *MedPage Today*, <https://www.medpagetoday.com/meetingcoverage/ama/99189>. 34
- Friedman, Gerald**, “Yes, We Can Have Improved Medicare for All,” 2019. Hopbrook Institute Working Paper No. 2. 1
- Geruso, Michael, Timothy Layton, and Daniel Prinz**, “Screening in Contract Design: Evidence from the ACA Marketplaces,” *American Economic Journal: Economic Policy*, 2019, 11, 64–107. 14
- Gibbons, Charles E., Juan Carlos Suárez Serrato, and Michael B. Urbancic**, “Broken or Fixed Effects?,” *Journal of Econometric Methods*, 2019, 8 (1). 16
- Glied, Sherry**, “Managed Care,” in “Handbook of Health Economics,” Vol. 1, Elsevier, 2000, pp. 707–753. 5
- Goldsmith-Pinkham, Paul, Peter Hull, and Michal Kolesár**, “Contamination Bias in Linear Regressions,” 2022. NBER Working Paper No. 30108. 16
- Gottlieb, Joshua D., Adam Hale Shapiro, and Abe Dunn**, “The Complexity Of Billing And Paying For Physician Care,” *Health Affairs*, 2018, 37 (4), 619–626. 4
- Gross, Tal, Timothy Layton, and Daniel Prinz**, “The Liquidity Sensitivity of Healthcare Consumption: Evidence from Social Security Payments,” *American Economic Review: Insights*, forthcoming. 4, 26
- Guimarães, Paulo, Octávio Figueiredo, and Douglas Woodward**, “A Tractable Approach to the Firm Location Decision Problem,” *Review of Economics and Statistics*, 2003, 85 (1), 201–204. 20, A26
- Himmelstein, David U., Terry Campbell, and Steffie Woolhandler**, “Health Care Administrative Costs in the United States and Canada, 2017,” *Annals of Internal Medicine*, 2020, 172 (2), 134–142. 1
- Homonoff, Tatiana and Jason Somerville**, “Program Recertification Costs: Evidence from SNAP,” *American Economic Journal: Economic Policy*, 2021, 13 (4), 271–298. 4
- Kakani, Pragya, Michael Chernew, and Amitabh Chandra**, “Rebates in the Pharmaceutical Industry: Evidence from Medicines Sold in Retail Pharmacies in the U.S.,” *Journal of Health Politics, Policy and Law*, 2022, 47, 1004–1079. 11
- LaPensee, Kenneth**, “Analysis of a Prescription Drug Prior Authorization Program in a Medicaid Health Maintenance Organization,” *Journal of Managed Care Pharmacy*, 2003, 9 (1), 36–44. A16
- Lee, Victoria, Todd Berland, Glenn Jacobowitz, Caron Rockman, Mikel Sadek, Michael Barfield, Neal Cayne, and Thomas S. Maldonado**, “Prior authorization as a utilization management tool for elective superficial venous procedures results in high administrative cost and low efficacy in reducing utilization,” *Journal of Vascular Surgery: Venous and Lymphatic Disorders*, 2020, 8 (3), 383–389. A16
- Lieber, Ethan M.J. and Lee M. Lockwood**, “Targeting with In-Kind Transfers: Evidence from Medicaid Home Care,” *American Economic Review*, 2019, 109 (4), 1461–1485. 4, 26


- Newhouse, Joseph P. and the Insurance Experiment Group**, *Free for All?: Lessons from the Rand Health Insurance Experiment*, Harvard University Press, 1993. [22](#)
- Nichols, Albert L. and Richard J. Zeckhauser**, “Targeting Transfers through Restrictions on Recipients,” *American Economic Review*, 1982, 72 (2), 372–377. [1](#), [5](#), [7](#)
- Pauly, Mark V.**, “The Economics of Moral Hazard: Comment,” *American Economic Review*, 1968, 58 (3), 531–537. [1](#), [4](#), [6](#)
- Petrin, Amil and Kenneth Train**, “A Control Function Approach to Endogeneity in Consumer Choice Models,” *Journal of Marketing Research*, 2010, 47 (1), 3–13. [21](#), [A29](#)
- Pollin, Robert, James Heintz, Peter Arno, Jeannette Wicks-Lim, and Michael Ash**, “Economic Analysis of Medicare for All,” November 2018. [1](#)
- Raper, James L., James H. Willig, Hui-Yi Lin, Jeroan J. Allison, M. Bennet Broner, Michael J. Mugavero, and Michael S. Saag**, “Uncompensated Medical Provider Costs Associated with Prior Authorization for Prescription Medications in an HIV Clinic,” *Clinical Infectious Diseases*, 2010, 51 (6), 718–724. [23](#), [A15](#), [A16](#)
- Sarig, Oren**, “Pharmaceutical Demand Response to Utilization Management,” 2020. [4](#)
- Schwartz, Aaron L., Troyen A. Brennan, Dorothea J. Verbrugge, and Joseph P. Newhouse**, “Measuring the Scope of Prior Authorization Policies: Applying Private Insurer Rules to Medicare Part B,” *JAMA Health Forum*, 2021, 2 (5), e210859–e210859. [A16](#)
- Seabury, Seth A., Darius N. Lakdawalla, Deborah Walter, John Hayes, Thomas Gustafson, Anshu Shrestha, and Dana P. Goldman**, “Patient Outcomes and Cost Effects of Medicaid Formulary Restrictions on Antidepressants,” 2014, 17 (2), 153–168. [4](#)
- Shao, Jun and Dongsheng Tu**, *The Jackknife and Bootstrap*, Springer Science + Business Media, 1995. [A31](#)
- Shepard, Mark and Myles Wagner**, “Reducing Ordeals through Automatic Enrollment: Evidence from a Subsidized Health Insurance Exchange,” 2022. NBER Working Paper No. 30781. [4](#)
- Shi, Maggie**, “Monitoring for Waste: Evidence from Medicare Audits,” 2022. [4](#), [5](#), [22](#)
- Train, Kenneth**, *Discrete Choice Methods with Simulation*, Cambridge University Press, 2009. [A27](#)
- U.S. Department of Health and Human Services Office of Inspector General**, “Medicare Advantage Appeal Outcomes and Audit Findings Raise Concerns About Service and Payment Denials,” 2018. Report OEI-09-16-00410. [A16](#)
- Waldinger, Daniel**, “Targeting In-Kind Transfers Through Market Design: A Revealed Preference Analysis of Public Housing Allocation,” *American Economic Review*, 2021, 111 (8), 2660–2696. [4](#)
- Wallace, Zachary S., Tyler Harkness, Xiaoqing Fu, John H. Stone, Hyon K. Choi, and Rochelle P. Walensky**, “Treatment Delays Associated With Prior Authorization for Infusible Medications: A Cohort Study,” *Arthritis Care & Research*, 2020, 72 (11), 1543–1549. [A16](#)
- Wooldridge, Jeffrey M.**, “Distribution-Free Estimation of Some Nonlinear Panel Data Models,” *Journal of Econometrics*, 1999, 90 (1), 77–97. [27](#)

Woolhandler, Steffie, Terry Campbell, and David U. Himmelstein, “Costs of Health Care Administration in the United States and Canada,” *New England Journal of Medicine*, 2003, 349 (8), 768–775. 1

Zeckhauser, Richard, “Medical Insurance: A Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives,” *Journal of Economic Theory*, 1970, 2 (1), 10–26. 4


8 Figures

Figure 1: Use of Prior Authorization in Our Sample




Notes: This figure plots a time series of the share of prescriptions filled among beneficiaries in our sample that required prior authorization. The blue dotted line plots the share of all filled prescriptions requiring prior authorization. The solid red line weights those prescriptions by their list price, such that it measures the share of total gross spending that required prior authorization. The dashed red line weights those prescriptions by their net price (list price net of rebate), such that it measures the share of total net spending that required prior authorization.

Figure 2: Prior Authorization Restrictions by Drug Price




Notes: This figure shows the relationship between average expenditure (net price) per day supply of a drug and the share of plans that put prior authorization restrictions on that drug. Each observation is a drug-year pair. Drugs with fewer than 20 prescriptions in a year within our sample are excluded. List price expenditure for a drug is calculated from the Medicare part D claims for beneficiaries in our sample, and deflated by average rebate for that drug from SSR Health data.

Figure 3: Prior Authorization Restrictions by Extent of Use




Notes: This figure displays the relationship between the number of users of a drug in a given year and the share of plans that put prior authorization restrictions on that drug. Each underlying observation is a single drug-year pair. Drugs with fewer than 20 prescriptions in a year are excluded.

Figure 4: Distribution of Drug-Level Frequency of Prior Authorization




Notes: This figure displays the distribution of the fraction of plans within a service region that require prior authorization for the drug in a given year, weighted by number of enrollees in the plan. Each underlying observation is a single drug-region-year pair, $N = 75,875$. Market-years where no plan requires prior authorization on a drug (74.2% of drug-region-years) or all plans require prior authorization on a drug (2.6% of drug-region-years) are excluded.

Figure 5: Heterogeneous Effects of Prior Authorization on Utilization by Beneficiary Characteristics




Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug, for subsamples of beneficiaries. Effects are presented in terms of the percent change due to prior authorization relative to a control mean, reweighted as described in Section 3.3.

Figure 6: Heterogeneous Effects of Prior Authorization on Utilization by Drug Characteristics



Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug, for subsamples of drugs. Effects are presented in the percent change due to prior authorization relative to a control mean, reweighted as described in Section 3.3.

Figure 7: Ratio of Drug Cost Reduction to Administrative Cost Burden by Class



Notes: This figure reports, for each therapeutic class, estimates of the amount of spending reduced due to status quo prior authorization policies per dollar of administrative costs induced, under the calibration $a = \$22.48$ and $r = 0.04$. This is reported for the top 30 therapeutic classes by total spending. Black brackets represent 95% confidence intervals. The red line is at \$1, at which the savings from reduced spending are exactly equal to the administrative costs. Negative values mean that prior authorization is estimated to lead to increases in spending. Blue bars indicate that the estimated savings-to-administration ratio is beyond the axes presented; for the four indicated classes, this ratio is above 200.

9 Tables

Table 1: Beneficiary Summary Statistics

	Analytic Sample	Broad LIS Population
Avg. Age	62.9	65.4
% Female	58.2	60.7
% White	60.6	64.3
Avg. Elixhauser Index	4.0	4.4
% Enrolled in Assigned Plan	91.1	
Share With Any Drug Use	91.5	91.4
Avg. # Unique Drugs Taken	10.8	11.1
Avg. # Unique Drugs Taken with Authorization Restrictions	0.2	0.2
Avg. Drug Spending	\$3,396	\$3,294
Avg. Non-Drug Medical Spending	\$11,286	\$10,034
Beneficiary-years	1,102,328	19,003,526

Notes: This table presents summary statistics for LIS beneficiaries. Observations are at the beneficiary-year level. The first column includes our primary sample, as described in Section 2.3. The second column includes all LIS beneficiaries who are observed in the data enrolling in Medicare Parts A, B, and D for all 12 months of the year. For the first column, spending outcomes are measured in the year before reassignment. In the second column, all outcomes are measured in the year of observation.

Table 2: Plan Summary Statistics

	Assigned plans	Enrolled plans
Mean beneficiaries per plan	803.5	138.9
Mean % of drugs under prior authorization	12.0	12.3
Standard deviation	(4.3)	(4.5)
10th percentile	5.7	5.8
Median	12.7	12.7
90th percentile	16.2	17.4
Mean % of drugs excluded	28.0	24.1
Standard deviation	(28.0)	(24.1)
10th percentile	15.7	5.0
Median	28.9	26.1
90th percentile	39.3	39.3
Plan-years	1,386	8,015

Notes: This table presents summary statistics for benchmark plans that were qualified to receive LIS beneficiaries through the default auto-assignment mechanism. Observations are at the plan-year level. The first column includes all benchmark plans that qualified to receive beneficiaries in our sample through the auto-assignment default mechanism. The second column includes all plans that beneficiaries in our sample enrolled in.

Table 3: Drug Summary Statistics

	All drugs	Drug type		Generic
		Branded without generic bioequivalent	Branded with generic bioequivalent	
Number of drug-years	12,605	4,457	3,443	4,705
Number of unique drugs	2,005	847	609	737
% of plan-years under prior authorization	12.6	23.3	5.8	7.4
% of plan-years excluded	29.1	27.3	57.2	10.2
% of beneficiaries with any use	0.8	0.3	0.2	1.7
List price per day supply	\$27.1	\$63.2	\$16.1	\$6.0
Net price per day supply	\$23.9	\$55.8	\$13.4	\$5.9
Net spending per enrolled beneficiary	\$2.0	\$3.5	\$0.9	\$1.5

Notes: This table presents summary statistics for drugs that were featured on a formulary for at least one benchmark plan in our sample during the period 2008-2015. A ‘drug’ is defined as a combination of active-ingredient and whether the product is branded/generic. Products containing different doses of the same active ingredient and with different modes of administration are all counted as the same drug.

Table 4: First Stage Regressions

	Full Sample		Existing Users	
	Auth ^{Enrolled}	Excluded ^{Enrolled}	Auth ^{Enrolled}	Excluded ^{Enrolled}
Auth ^{Assigned}	0.908 (0.002)	-0.001 (0.000)	0.873 (0.005)	0.001 (0.001)
Excluded ^{Assigned}	0.000 (0.000)	0.905 (0.003)	0.003 (0.001)	0.849 (0.006)
F-statistic	76,297	61,845	18,382	9,539
Number of drug-beneficiary-years	1,723,975,571		10,220,638	
Number of beneficiary-years	1,113,594		1,000,779	
Number of market-years	210		210	
Average plans per market-year	6.6		6.6	
Number of drug-years	12,554		12,554	

Notes: This table presents coefficient estimates from the first stage regressions of indicators for whether the plan a beneficiary enrolled in during a given year placed prior authorization restrictions on or excluded a drug in that year, on indicators for whether the plan the beneficiary was assigned to placed prior authorization restrictions on or excluded that drug. In Columns (1) and (3), the outcome is whether the plan of enrollment restricted the drug in that year. In Columns (2) and (4), the outcome is exclusion rather than restriction. Each underlying observation is a beneficiary-drug-year tuple. Columns (1) and (2) include our entire sample, with all possible beneficiary-drug-year tuples. Columns (3) and (4) restrict to beneficiary-drug-year tuples where the beneficiary filled a prescription for the drug at least once during the year before reassignment. Standard errors are clustered at the assigned plan and year level.

Table 5: Estimates of the Effect of Prior Authorization Status on Drug Utilization

	(1)	(2)	(3)	(4)	(5)	(6)
Auth ^{Enrolled}	-1.169	-0.136	-0.098	-0.099	-0.101	-0.108
	(0.012)	(0.005)	(0.004)	(0.003)	(0.003)	(0.004)
Auth ^{Sub}						0.049
						(0.0036)
PA % Effect	-290.0	-33.7	-24.3	-24.5	-25.1	-26.8
Control Mean				1.299		
Rewighted Control Mean				0.403		
Drug FEs		X				
Drug-year FEs			X			
Drug-market-year FEs				X	X	X
Plan-market-year FEs					X	X
Substitution Controls						X
Number of drug × beneficiary-years			1,723,975,571			
Number of market years			210			
Average plans per market-year			6.6			
Average beneficiaries per plan			51			
Average drugs per year			1569.2			

Notes: This table presents coefficient estimates from instrumental variable regressions of a beneficiary's utilization of a drug in a given year on an indicator for whether the drug was put under prior authorization restrictions in the plan that beneficiary was enrolled in that year. Prior authorization in the plan in which the beneficiary is enrolled in is instrumented for by prior authorization restriction and exclusion status in the plan to which the beneficiary was randomly assigned. Each underlying observation is a beneficiary-drug-year tuple. Regressions include plan-market-year and drug-market-year fixed effects. Prior authorization of substitute drugs is mean prior authorization status of all other drugs within the class, where drugs are weighted by their average expenditure across all plans in the sample. Standard errors are clustered at the assigned plan and year level. Columns represent regressions with different sets of controls.

Table 6: Spending and Utilization Effects of Status Quo Relative to Ban on Prior Authorization Restrictions

	Total	Restricted Drugs	Unrestricted Drugs	No Drug
Change in	-3.57%	-21.8%	+0.72%	-
Spending	-95.88	-111.57	+15.69	-
Per Capita				
Change in	-0.65%	-28.9%	+0.58%	+0.06%
# Users	-0.065	-0.120	+0.056	+0.065
Per Capita				
Diversion	-	-100%	46.2%	53.8%

Notes: This table presents results from an exercise where we simulate switching beneficiaries from facing no authorization restrictions to facing the status quo formulary restrictions. The first two panels detail the change in spending and utilization of all drug, restricted drugs (those drug-plan-region-year observations where an authorization restriction was in place in the status quo), unrestricted drugs, and no drug. In those panels, the upper row gives the percent change in these quantities, while the lower row presents the absolute change per beneficiary-year. The final panel details the share of beneficiaries moving away from restricted drugs to either unrestricted drugs or no drug.

Table 7: Per Capita Administrative Burden of Authorization Restrictions

		Request Rejection Rate				
		0%	1.5%	4%	7.5%	15%
Paperwork Cost	\$11.62	\$4.84	\$4.92	\$5.04	\$5.24	\$5.70
	\$18.19	\$7.58	\$7.70	\$7.90	\$8.20	\$8.92
	\$21.72	\$9.05	\$9.19	\$9.43	\$9.79	\$10.65
	\$22.48	\$9.37	\$9.51	\$9.76	\$10.13	\$11.02
	\$31.30	\$13.04	\$13.24	\$13.59	\$14.10	\$15.35
	\$50	\$20.84	\$21.16	\$21.71	\$22.53	\$24.52
	\$100	\$41.68	\$42.31	\$43.41	\$45.06	\$49.03
	\$200	\$83.35	\$84.62	\$86.83	\$90.11	\$98.06

Notes: This table reports estimates of the administrative costs of administering the historical prior authorization restriction regimes implemented in Medicare Part D per beneficiary-year. Each cell represents the estimate under a calibrated set of values for the application cost a and rejection rate r . Spending reductions from prior authorization are estimated at \$96, and so values below that indicate that prior authorization generates net financial savings, while values below it indicate net financial losses.

Table 8: Effect of Out-of-Pocket Price Increases on Utilization for Beneficiaries Transitioning into the LIS Program

	(1)	(2)
Out-of-pocket price $\times$ Not LIS	-0.0051 (0.0011)	-0.0015 (0.0017)
Mean	0.035	
Rewighted mean	0.002	
Drug-market-year FEs	X	X
Beneficiary FEs		X
Number of drug $\times$ beneficiary years	110,498,105	
Number of market years	211	
Average drugs per year	1,054	

Notes: This table reports coefficient estimates from regressions of a beneficiary's number of prescriptions filled for a given drug in a given year on out-of-pocket price per year. These regressions use our sample who transition into the LIS program and leverage the transition as a shock to out-of-pocket prices. Standard errors are clustered at the beneficiary level.

Table 9: Revealed Preference Estimates of Consumer Surplus Loss

	(1)	(2)
Best-Case Screening	3.83	13.05
Random Screening	23.80	80.91
Beneficiary FEs		X
Net Financial Savings from Prior Authorization	86.58	

Notes: This table provides estimates of the loss in consumer surplus, in dollars per beneficiary-year, due to the presence of prior authorization restrictions moving beneficiaries away from their most-preferred drugs. The estimates are derived from estimates of the elasticity of drug use with respect to out-of-pocket price from the columns of Table 8, as well as from estimates of how prior authorization changes drug use, given in Table 6. The two columns represent the consumer surplus measures derived from the two columns in Table 8, respectively. The three rows represent different assumptions about the extent to which beneficiary value for a drug is related to their propensity to switch drugs in response to prior authorization. In the 'best case,' marginal beneficiaries who switch are those using the original drug who value it the least. In the 'random' case, marginal beneficiaries have an average value for the drug relative to others using it. In the 'worst' case, marginal beneficiaries have the highest value for the drug. The "net financial savings" listed come from the difference between our estimate of spending reductions in Table 6 and our preferred estimate of the average administrative cost of prior authorization given in bold in Table 7.

Table 10: Effects of Prior Authorization Restrictions on Utilization for Oral Anticoagulants

	Spending			Any prescription		
	All	NOACs	Warfarin	All	NOACs	Warfarin
All NOACs PA	-16.6 (6.41)	-18.3 (6.61)	1.7 (0.70)	-0.0003 (0.0033)	-0.0097 (0.0028)	0.0069 (0.0035)
Other restrictions	-16.1 (5.51)	-12.4 (5.14)	-3.7 (2.51)	-0.0011 (0.0086)	-0.0100 (0.0035)	0.0058 (0.0085)
Control Mean	111.585	77.433	34.152	0.291	0.043	0.260
Beneficiary-years	134,182					
Market-years	160					

Notes: This table presents coefficient estimates from a set of regressions of a beneficiary's utilization of oral anticoagulants in a given year on an indicator for whether the beneficiary's assigned plan put prior authorization restrictions on all non-Vitamin K oral anticoagulants that year. Regressions include market fixed effects. Standard errors are clustered at the beneficiary level.

Table 11: Effects of Oral Anticoagulant Prior Authorization Restrictions on Health Outcomes

	Died	Stroke	Bleed
All NOACs under prior auth	-0.00032 (0.00090)	0.00067 (0.00143)	-0.00124 (0.00172)
Other restrictions on NOACs	-0.00152 (0.00189)	0.00569 (0.00275)	0.00480 (0.00497)
Control Mean	0.014	0.032	0.054
N (beneficiary-years)	134,182		
N (market-years)	160		

Notes: This table presents coefficient estimates from a set of regressions of a beneficiary's health outcomes in a given year on an indicator for whether the beneficiary's assigned plan put prior authorization restrictions on all non-Vitamin K oral anticoagulants that year. Regressions include market fixed effects. Standard errors are clustered at the beneficiary level.

Table 12: Effects of Aggregate Prior Authorization Restriction Exposure on Utilization and Health Outcomes

All Beneficiaries				
	Spending	% Died in Year	Inpatient Spending	Non-Drug Medical Spending
AuthExposureQ5	-64.825 (29.508)	0.058 (0.065)	68.446 (87.435)	78.948 (120.543)
Control Mean	4,210.39	2.253	12,196.71	1,579.28
N (beneficiary-years)	609,316			

Top 25% of Beneficiaries by Spread in Fit				
	Spending	% Died in Year	Inpatient Spending	Non-Drug Medical Spending
AuthExposureQ5	-192.499 (43.018)	-0.051 (0.115)	-162.382 (145.638)	-309.17 (201.224)
Control Mean	3,909.49	2.151	5,429.84	11,090.72
N (beneficiary-years)	152,385			


Notes: This table presents coefficient estimates from a set of regressions of a beneficiary's utilization and health outcomes in a given year on an indicator for whether their assigned plan was in the bottom quintile of benchmark plans in terms of putting authorization restrictions on their previously-taken drugs. Regressions include market fixed effects and a control for exclusion exposure.

For Online Publication

Appendix for:
**Rationing Medicine Through Bureaucracy:
Authorization Restrictions in Medicare**


A Additional Figures

Appendix Figure A1: Distribution of Number of Benchmark Plans in Region-Year




Notes: This set of figures plots the distribution in the number of benchmark plans across the pairs of Part D service region-years. The top figure presents this distribution weighting all Part D service region-year pairs equally, while the bottom weights Part D service region-year pairs by the number of beneficiaries in our sample enrolled under each.

Appendix Figure A2: Heterogeneous Effects of Prior Authorization on Utilization by Drug-Year Price Deciles




Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug. We run separate regressions on groups of drug-year pairs, where pairs are grouped into decile based on their price per day supply. Effects are presented in the percent change due to prior authorization relative to a control mean, reweighted as described in 3.3.


Appendix Figure A3: Heterogeneous Effects of Prior Authorization on Utilization by Drug-Year Restriction Rate Deciles




Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug. We run separate regressions on groups of drug-year pairs, where pairs are grouped into decile based on the share of plans in that year that put the drug under a prior authorization restriction. Effects are presented in the percent change due to prior authorization relative to a control mean, reweighted as described in 3.3.

Appendix Figure A4: Heterogeneous Effects of Prior Authorization on Utilization by Class


Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug. We run separate regressions on each drug therapeutic class. We report results only for the top 30 classes by total spending. Effects are presented in the percent change due to prior authorization relative to a control mean, reweighted as described in 3.3.

Appendix Figure A5: Distribution of Nesting Parameter Across Classes

Notes: This figure presents a histogram of estimates of λ , the nesting parameter in our nested logit demand system. Each underlying observation in this figure is an estimate for a specific therapeutic class.



Appendix Figure A6: This figure presents the share of patients filling each of these oral anticoagulants at least once during the year, across time.



Appendix Figure A7: This figure presents the per-patient yearly spending on each of these oral anticoagulants, across time.

B Additional Tables

Appendix Table A1: Prior Authorization Frequency for Top Drug Classes by Medicare Part D Spending

	Spending per beneficiary year (USD)	% spending with prior auth	% fills with prior auth
Biological Response Modifiers	94	69.6	68.1
Immunosuppressants	65	66.3	54.7
Antineoplastic Agents	99	57.7	13.9
Adrenals & Comb	86	3.0	11.6
CNS Agents, Misc	94	17.6	6.9
Cardiac Drugs	88	12.4	5.9
Antidiabetic Agents, Misc	110	15.0	5.7
Estrogens & Comb	25	1.2	5.4
Bone Resorption Inhibitors	22	9.0	4.8
Misc Therapeutic Agents	58	15.0	4.0
Tranq/Antipsychotic	185	6.9	3.6
Sympathomimetic Agents	27	2.1	3.4
Antidepressants	93	7.7	3.3
Gastrointestinal Drug, Misc	132	2.8	3.2
Anticoagulants	47	14.5	2.8
Muscle Relaxants	36	1.9	2.3
Antivirals	120	14.6	2.1
NSAIDs	37	10.0	1.6
Anticonvulsants, Misc	60	4.4	1.6
Vasodilating Agents	27	44.6	1.5
Parasympathomimetic	42	3.2	1.5
Antiplatelet Agents	70	0.6	1.4
Antihyperlipidemic Drugs	212	2.7	1.1
Cardiac, Calcium Channel	49	1.5	1.0
Antidiabetic Agents, Insulins	158	0.6	0.9
Opiate Agonists	92	3.5	0.7
Eye/Ear/Nose/Throat Misc	44	1.2	0.6
Cardiac, Beta Blockers	45	0.5	0.5
Antiinflam Agents EENT	29	0.1	0.2
Anticholinergic	47	0.1	0.2

Notes: This table reports, for a set of therapeutic classes, the total spending per beneficiary-year, the share of spending where the drug being filled required a prior authorization restriction, and the share of prescription drug fills where the drug being filled required a prior authorization restriction. All statistics are limited to beneficiaries in our sample.

Appendix Table A2: First Stage Regressions with Further Specifications

Auth ^{Enrolled}						
Auth ^{Assigned}	0.950 (0.002)	0.913 (0.003)	0.908 (0.003)	0.908 (0.002)	0.908 (0.002)	0.908 (0.002)
Excluded ^{Assigned}	0.002 (0.000)	-0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Excluded ^{Enrolled}						
Auth ^{Assigned}	0.003 (0.000)	0.002 (0.000)	-0.001 (0.000)	-0.001 (0.000)	-0.001 (0.000)	-0.001 (0.000)
Excluded ^{Assigned}	0.950 (0.002)	0.918 (0.003)	0.905 (0.003)	0.905 (0.003)	0.905 (0.003)	0.905 (0.003)
Drug FEs	X					
Drug-year FEs	X					
Drug-market-year FEs	X					
Plan-market-year FEs	X					
Substitution Controls	X					
Number of drug-beneficiary-years	1,723,975,571					
Number of beneficiary-years	1,113,594					
Number of market-years	210					
Average plans per market-year	6.6					
Number of drug-years	12,554					

Notes: This table presents coefficient estimates from the ‘first stage’ regressions of indicators for whether the plan a beneficiary enrolled in during a given year placed prior authorization restrictions on or excluded a drug in that year, on indicators for whether the plan the beneficiary was assigned to placed prior authorization restrictions on or excluded that drug. In the upper panel, the outcome is whether the plan of enrollment restricted the drug in that year. In the lower panel, the outcome is exclusion rather than restriction. Each underlying observation is a beneficiary-drug-year tuple. Standard errors are clustered at the assigned plan and year level. Columns represent regressions with different sets of controls.

Appendix Table A3: Placebo Test: Formulary Status in Prior Year

	Auth ^{Enrolled} _{t-1}	Excluded ^{Enrolled} _{t-1}
Auth ^{Assigned} _t	-0.001 (0.002)	-0.001 (0.001)
Excluded ^{Assigned} _t	-0.001 (0.001)	-0.003 (0.001)
F-statistic	1	4
Number of drug-beneficiary-years	1,510,671,381	
Number of beneficiary-years	1,037,159	
Number of market-years	210	
Average plans per market-year	6.6	
Number of drug-years	11,906	

Notes: This table presents estimates from a set of ‘placebo’ versions of our first-stage regressions, where we regress indicators for whether the plan a beneficiary enrolled in during a given year placed prior authorization restrictions on or excluded a drug in the year before reassignment on indicators for whether the plan the beneficiary was assigned to placed prior authorization restrictions on or excluded that drug in a given year following reassignment. Each underlying observation is a beneficiary-drug-year tuple. Standard errors are clustered at the plan and year level. In columns (1), the outcome is whether the plan of enrollment restricted the drug in that year. In columns (2) the outcome is exclusion rather than restriction.

Appendix Table A4: Placebo Test: Utilization in Prior Year

	Spending	# Fills	# Days Supply	% Ever Filled
Auth ^{Assigned}	-0.011 (0.032)	0.000 (0.000)	0.003 (0.005)	0.002 (0.005)
Rewighted Control Mean	2.651	0.135	0.403	0.307

Notes: This table presents estimates from a set of ‘placebo’ utilization regressions, where we regress a beneficiary’s utilization of a drug in the year before reassignment on an indicator for whether the drug was put under prior authorization restrictions in the plan that beneficiary was assigned to in a given year following reassignment. Each underlying observation is a beneficiary-drug-year tuple. Standard errors are clustered at the plan and year level.

Appendix Table A5: Placebo Test: Demographics

	Female	White	Age	Elixhauser Index $_{t-1}$
Auth ^{Assigned}	-0.000 (0.000)	-0.000 (0.000)	0.002 (0.001)	0.001 (0.000)
Excluded ^{Assigned}	-0.000 (0.000)	0.000 (0.000)	-0.006 (0.001)	-0.002 (0.000)
Control Mean	0.583	0.614	62.6	3.56
Number of drug-plan-years	2,149,673			

Notes: This table presents estimates from a set of ‘placebo’ utilization regressions, where we regress indicators for a beneficiary being in certain demographic groups on an indicator for whether the drug was put under prior authorization restrictions in the plan that beneficiary was assigned to in a given year following reassignment. Each underlying observation is a beneficiary-drug-year tuple. Standard errors are clustered at the plan and year level.

Appendix Table A6: Estimates of the Effect of Prior Authorization Restrictions on Utilization: Additional Specifications

	(7)	(8)	(9)
Auth ^{Assigned}	-0.107 (0.004)	-0.090 (0.003)	-0.108 (0.005)
Auth ^{Sub}	0.049 (0.0043)	0.391 (0.0186)	0.049 (0.0065)
PA % Effect	-26.7	-22.4	-26.7
Control Mean		1.305	
Rewighted Control Mean	0.403		0.395
Drug-market-year FEs	X	X	X
Plan-market-year FEs			X
Substitution Controls	X	X	X
Plan-by-cost FEs	X		
Plan-by-class FEs		X	
Number of drug × beneficiary-years	1,723,975,571		1,237,515,645
Number of market years		210	
Average plans per market-year		6.6	
Average beneficiaries per plan		803	
Average drugs per year	1569.2		1460.2

Notes: This table presents coefficient estimates from regressions of a beneficiary's utilization of a drug in a given year on an indicator for whether the drug was put under prior authorization restrictions in the plan that beneficiary was assigned to in that year. Each underlying observation is a beneficiary-drug-year tuple. Regressions include plan-market-year and drug-market-year fixed effects. Prior authorization of substitute drugs is mean prior authorization status of all other drugs within the class, where drugs are weighted by their average expenditure across all plans in the sample. Standard errors are clustered at the assigned plan and year level. Columns represent regressions with different sets of controls, except for the final column, which represents a version of the main regression specification that drops all observations where the drug in question was excluded. This table presents specifications not otherwise presented in Table 5.

Appendix Table A7: Estimates of the Effect of Prior Authorization Restrictions on Alternative Utilization Outcomes

	Days Supply	Spending	Fills
Auth ^{Assigned}	-0.156 (0.057)	-0.767 (0.155)	-0.005 (0.000)
Auth ^{Sub}	0.057 (0.0065)	0.155 (0.0285)	0.002 (0.0002)
PA % Effect	-30.9	-21.2	-28.5
Control Mean	1.529	3.555	0.051
Reweighted Control Mean	0.504	3.613	0.017
R ²	0.969	0.831	0.966
Number of drug × beneficiary-years	1,732,564,415		
Number of market years	210		
Average plans per market-year	6.6		
Average beneficiaries per plan	807		

Notes: This table presents coefficient estimates from regressions of a beneficiary's utilization of a drug in a given year on an indicator for whether the drug was put under prior authorization restrictions in the plan that beneficiary was assigned to in that year. Each underlying observation is a beneficiary-drug-year tuple. Regressions include plan-market-year and drug-market-year fixed effects. Prior authorization of substitute drugs is mean prior authorization status of all other drugs within the class, where drugs are weighted by their average expenditure across all plans in the sample. Standard errors are clustered at the assigned plan and year level. Columns represent different outcomes. The three outcomes are, in order, total allowed net spending on the drug in the year, number of prescription fills for the drug in the year, and total days supply for the drug in the year.

Appendix Table A8: Estimates of Prior Authorization Per-Application Administrative Costs

Study	Setting	Method	Estimate
Bukstein et al. (2006)	Single allergist clinic,	Staff time at hourly wages, ^a mean	\$17.77
Raper et al. (2010)	Single HIV clinic,	Staff time at hourly wages, plus materials costs, mean	\$14.24
		Staff time at opportunity costs, ^b plus materials costs, mean	\$27.35
CAQH (2013)	Many surveyed practices	Staff time at estimated global rates, mean	
		...for manual filing ^c	\$18.53
		...for electronic filing	\$5.20
Carlisle et al. (2020)	Single dermatology clinic	Staff time at hourly wages, median	\$7.67

Notes: This table presents estimates from the literature on the per-application administrative costs associated with drugs restricted under prior authorization. All studies are in U.S. settings unless otherwise noted.

^a In this method, the researchers convert employees' salaries to hourly wage equivalents, then price their time using those hourly equivalents.

^b In this method, the researchers calculate the revenue the practice would have received if the nurse involved took the time spent on the prior authorization request and instead billed insurers for the time-equivalent number of 30-minute visits for established patients (CPT code 99213) at standard Medicare rates at the time. In their manuscript, [Raper et al. \(2010\)](#) incorrectly add their wage-equivalent and opportunity cost estimates together, which is incorrect since it double-counts the nurse's time.

^c [CAQH \(2013\)](#) distinguish between the costs of filing manually (i.e., with a fax machine or phone) or electronically (through the internet). Few prior authorization requests during our period were electronic, so we only use the manual costs in our calibration exercise.

Appendix Table A9: Estimates of Prior Authorization Request Rejection Rates

Study	Setting	Services	Estimate
LaPensee (2003)	One Medicaid MCO	All drugs	4.4%
		Non-formulary drugs	3.7%
		Formulary drugs	7.1%
Delate et al. (2005)	Medicaid	Proton-pump inhibitors	4.9%
Raper et al. (2010)	Single HIV clinic	All drugs	33%
Initial application ^a			
U.S. OIG (2018)	All Medicare Advantage MCOs	All services and drugs	4.1%
Birdsall et al. (2020)	Academic health system	All drugs	
Initial application			15%
Final application			7.4%
Carlisle et al. (2020) ^a	Single dermatology clinic	Biologics	21.1%
Initial application		Other drugs	41.8%
Lee et al. (2020) ^a	Division of Vascular Surgery New York University Hospital, 2017	Lower-extremity venous procedures	6.1%
Wallace et al. (2020)	Single rheumatology clinic	Infusable drugs	
Initial application			21%
Final application			4%
Schwartz et al. (2021)	Large private insurer	Hosp. services and drugs	4.2%
AthenaHealth ^b	Physician clients	All drugs	1.5%

Notes: This table presents estimates from the literature on the rejection rates associated with requests made for services and drugs restricted under prior authorization.

^a This study does not report interpretable final application approval rates.

^b <https://www.athenahealth.com/prior-authorization>. Last accessed on 07/13/22.

Appendix Table A10: Simulation of Moving From No Prior Authorization Restrictions to Status Quo (with Standard Errors)

	Total	Restricted Drugs	Unrestricted Drugs	No Drug
Change in	-3.57%	-21.8%	+0.72%	-
Spending	(0.84)	(4.25)	(0.05)	
Per Capita	-95.88	-111.57	+15.69	-
	(23.92)	(23.78)	(1.00)	
Change in	-0.65%	-28.9%	+0.58%	+0.06%
# Users	(0.13)	(3.17)	(0.02)	(0.01)
Per Capita	-0.065	-0.120	+0.056	+0.065
	(0.013)	(0.013)	(0.002)	(0.013)
Diversion	-	-100%	46.2%	53.8%
			(7.48)	(7.48)

Notes: This table replicates 6 but adds standard errors to the relevant estimates, given in parentheses under their respective estimate.

Appendix Table A11: Administrative Costs From Authorization Restrictions (with Standard Errors)

		Request Rejection Rate				
		0%	1.5%	4%	7.5%	15%
Paperwork Cost	\$11.62	4.84	4.92	5.04	5.24	5.70
		(0.05)	(0.05)	(0.05)	(0.05)	(0.06)
	\$18.19	7.58	7.70	7.90	8.20	8.92
		(0.07)	(0.07)	(0.08)	(0.08)	(0.09)
	\$21.72	9.05	9.19	9.43	9.79	10.65
		(0.09)	(0.09)	(0.09)	(0.09)	(0.10)
	\$22.48	9.37	9.51	9.76	10.13	11.02
		(0.09)	(0.09)	(0.09)	(0.10)	(0.11)
	\$31.30	13.04	13.24	13.59	14.10	15.35
		(0.13)	(0.13)	(0.13)	(0.14)	(0.15)
	\$50	20.84	21.16	21.71	22.53	24.52
		(0.20)	(0.21)	(0.21)	(0.22)	(0.24)
	\$100	41.68	42.31	43.41	45.06	49.03
		(0.40)	(0.41)	(0.42)	(0.44)	(0.48)
	\$200	83.35	84.62	86.83	90.11	98.06
		(0.81)	(0.82)	(0.84)	(0.87)	(0.95)

Notes: This table replicates 7 but adds standard errors to the relevant estimates, given in parentheses under their respective estimate.

Appendix Table A12: Savings per Administrative Dollar From Authorization Restrictions (with Standard Errors)

		Request Rejection Rate				
		0%	1.5%	4%	7.5%	15%
Paperwork Cost	\$11.62	19.80 (4.83)	19.50 (4.76)	19.01 (4.64)	18.31 (4.47)	16.83 (4.11)
	\$18.19	12.65 (3.09)	12.46 (3.04)	12.14 (2.96)	11.70 (2.86)	10.75 (2.62)
	\$21.72	10.59 (2.58)	10.43 (2.55)	10.17 (2.48)	9.80 (2.39)	9.00 (2.20)
	\$22.48	10.23 (2.50)	10.08 (2.46)	9.82 (2.40)	9.47 (2.31)	8.70 (2.12)
	\$31.30	7.35 (1.79)	7.24 (1.77)	7.06 (1.72)	6.80 (1.66)	6.25 (1.52)
	\$50	4.60 (1.12)	4.53 (1.11)	4.42 (1.08)	4.26 (1.04)	3.91 (0.95)
	\$100	2.30 (0.56)	2.27 (0.55)	2.21 (0.54)	2.13 (0.52)	1.96 (0.48)
	\$200	1.15 (0.28)	1.13 (0.28)	1.10 (0.27)	1.06 (0.26)	0.98 (0.24)

Notes: This table reports estimated ratios of the spending reductions induced by the historical prior authorization restriction regimes implemented in Medicare Part D relative to the costs of paperwork. Each cell represents the estimate under a calibrated set of values for the application cost a and rejection rate r . Values above 1 indicate that prior authorization generates net financial savings, while values below it indicate net financial costs. Parenthetical terms denote bootstrap standard errors for their associated estimate.

Appendix Table A13: Spending and Utilization Effects from Applying Authorization Restrictions to Currently-Unrestricted Drugs

	Total	Unrestricted Drugs	PA/Ex Drugs	No Drug
Change in	-7.02%	-11.91%	+0.14%	-
Spending	-181.55	-249.79	+68.25	-
Per Capita				
Change in	-11.71%	-13.52%	+26.28%	+1.05%
# Users	-1.16	-1.28	+0.12	+1.16
Per Capita				
Diversion	-	-100.0%	9.2%	90.8%

Notes: This table presents results from an exercise where we simulate switching beneficiaries from facing the status quo formulary restrictions to facing prior authorization restrictions on all previously-unrestricted drugs. The first two panels detail the change in spending and utilization of all drug, restricted drugs (those drug-plan-region-year observations where an authorization restriction was in place in the status quo), unrestricted drugs, and no drug. In those panels, the upper row gives the percent change in these quantities, while the lower row presents the absolute change per beneficiary-year. The final panel details the share of beneficiaries moving away from restricted drugs to either unrestricted drugs or no drug .

Appendix Table A14: Per Capita Administrative Burden of Authorization Restrictions from Applying Authorization Restrictions to Currently-Unrestricted Drugs

		Request Rejection Rate				
		0%	1.5%	4%	7.5%	15%
Paperwork Cost	\$11.62	\$110.04	\$111.72	\$114.63	\$118.96	\$129.46
	\$18.19	\$172.26	\$174.88	\$179.44	\$186.22	\$202.66
	\$21.72	\$205.69	\$208.82	\$214.26	\$222.36	\$241.98
	\$22.48	\$212.88	\$216.13	\$221.75	\$230.14	\$250.45
	\$31.30	\$296.41	\$300.92	\$308.76	\$320.44	\$348.72
	\$50	\$473.50	\$480.71	\$493.22	\$511.89	\$557.05
	\$100	\$946.99	\$961.41	\$986.45	\$1023.77	\$1114.11
	\$200	\$1893.98	\$1922.83	\$1972.90	\$2047.55	\$2228.22

Notes: This table reports estimates of the increase in administrative costs from a simulation of switching beneficiaries from facing the status quo formulary restrictions to facing prior authorization restrictions on all previously-unrestricted drugs. Each cell represents the estimate under a calibrated set of values for the application cost a and rejection rate r .

Appendix Table A15: Spending Reductions per Administrative Dollar from Applying Authorization Restrictions to Currently-Unrestricted Drugs

		Request Rejection Rate				
		0%	1.5%	4%	7.5%	15%
Paperwork Cost	\$11.62	1.65	1.63	1.58	1.53	1.40
	\$18.19	1.05	1.04	1.01	0.97	0.90
	\$21.72	0.88	0.87	0.85	0.82	0.75
	\$22.48	0.85	0.84	0.82	0.79	0.72
	\$31.30	0.61	0.60	0.59	0.57	0.52
	\$50	0.38	0.38	0.37	0.35	0.33
	\$100	0.19	0.19	0.18	0.18	0.16
	\$200	0.10	0.09	0.09	0.09	0.08

Notes: This table reports estimates of the ratio of reductions in drug spending to the increase in administrative costs from a simulation of switching beneficiaries from facing the status quo formulary restrictions to facing prior authorization restrictions on all previously-unrestricted drugs. Each cell represents the estimate under a calibrated set of values for the application cost a and rejection rate r . Ratios above 1 indicate net financial savings, while ratios below 1 indicate net financial losses.

Appendix Table A16: Summary Statistics for LIS Transition Sample

	Analytic Sample
Avg. Age	70.4
Share Female	64.7
Share White	72.1
Avg. Elixhauser Index	3.36
Share With Any Drug Use	93.0
Avg. # Unique Drugs Taken	10.0
Avg. # Unique Drugs Taken with Authorization Restrictions	0.1
Avg. Drug Spending	\$2,418
Avg. Non-Drug Medical Spending	\$4,978
Beneficiary-year observations	956,460

Notes: This table provides summary statistics for the sample of beneficiaries who transition into the LIS program during our sample window. This is the primary sample used in Section 6.1.

Appendix Table A17: Summary Statistics for Oral Anticoagulant Sample

Mean Age	71.0
Share of female beneficiaries	63.5%
Share of white beneficiaries	61.8%
Mean CHADS ² Vasc ² score	5.61
Share of anticoagulant users	29.1%
Share of NOAC users	4.33%
Share of Warfarin users	26.0%
Share of beneficiary-years facing prior auth on all NOACs	17.1%
Share of beneficiary-years facing prior auth on no NOACs	79.3%
Beneficiary-years	134,182

Notes: This table provides summary statistics for the sample of beneficiaries with prior diagnosis of atrial fibrillation, deep vein thrombosis, and/or pulmonary embolism. This is the primary sample used in Section 6.2.1.

C Prior Authorization Form Examples



<https://providers.amerigroup.com>

Novel Oral Anticoagulants Prior Authorization of Benefits Form

CONTAINS CONFIDENTIAL PATIENT INFORMATION

Complete form in its entirety and fax to: Prior Authorization of Benefits Center at 1-844-512-9004.

Provider Help Desk: 1-800-454-3730

1. Patient information		2. Physician information	
Patient name: _____		Prescribing physician: _____	
Patient ID #: _____		Physician address: _____	
Patient DOB: _____		Physician phone #: _____	
Date of Rx: _____		Physician fax #: _____	
Patient phone #: _____		Physician specialty: _____	
Patient email address: _____		Physician DEA: _____	
		Physician NPI #: _____	
		Physician email address: _____	
3. Medication	4. Strength	5. Directions	6. Quantity per 30 days
_____	_____	_____	Specify: _____
7. Diagnosis: _____			
8. Approval criteria: (Check all boxes that apply. Note: Any areas not filled out are considered not applicable to your patient and may affect the outcome of this request.)			
Prior authorization (PA) is not required for preferred novel oral anticoagulants (NOACs). PA is required for nonpreferred NOACs. Requests for doses outside of the manufacturer recommended dose will not be considered. Payment will be considered for FDA approved or compendia indications under the following conditions: 1. Patient does not have a mechanical heart valve. 2. Patient does not have active bleeding. 3. For a diagnosis of atrial fibrillation or stroke prevention, patient has the presence of at least 1 additional risk factor for stroke, with a CHA₂DS₂-VASc score ≥ 1. 4. A recent creatinine clearance (CrCl) is provided. 5. A recent Child-Pugh score is provided. 6. Patient's current body weight is provided. 7. Patient has documentation of a trial and therapy failure at a therapeutic dose with at least two preferred NOACs. 8. For requests for edoxaban, documentation patient has had 5 to 10 days of initial therapy with a parenteral anticoagulant (low molecular weight heparin or unfractionated heparin). The required trials may be overridden when documented evidence is provided that the use of these agents would be medically contraindicated.			
Preferred (no PA required if within established quantity limits)		Nonpreferred	
☐ Eliquis ☐ Xarelto		☐ Savaysa	
☐ Pradaxa			



OptumRx has partnered with CoverMyMeds to receive prior authorization requests, saving you time and often delivering real-time determinations. Visit go.covermymeds.com/OptumRx to begin using this free service. Please note: All information below is required to process this request. Mon-Fri: 5am to 10pm Pacific / Sat: 6am to 3pm Pacific

Zetia® (ezetimibe) Prior Authorization Request Form

DO NOT COPY FOR FUTURE USE. FORMS ARE UPDATED FREQUENTLY AND MAY BE BARCODED

Member Information (required)			Provider Information (required)		
Member Name:			Provider Name:		
Insurance ID#:			NPI#:		Specialty:
Date of Birth:			Office Phone:		
Street Address:			Office Fax:		
City:	State:	Zip:	Office Street Address:		
Phone:			City:	State:	Zip:
Medication Information (required)					
Medication Name:			Strength:		Dosage Form:
☐ Check if requesting brand			Directions for Use:		
☐ Check if request is for continuation of therapy					
Clinical Information (required)					
Select the diagnosis below: ☐ Homozygous Familial Hypercholesterolemia (HoFH) ☐ Homozygous Sitosterolemia ☐ Primary Hypercholesterolemia ☐ Other diagnosis: _____ ICD-10 Code(s): _____					
Clinical information: Has the patient's diagnosis been confirmed? ☐ Yes ☐ No					
Select the medications the patient has a failure, contraindication, or intolerance to: ☐ Ezetimibe-simvastatin ☐ Lovastatin ☐ Simvastatin ☐ Other statin or statin combination product. Please specify all: _____					
Quantity limit requests: What is the quantity requested per DAY? _____ What is the reason for exceeding the plan limitations? ☐ Titration or loading dose purposes ☐ Patient is on a dose-alternating schedule (e.g., one tablet in the morning and two tablets at night, one to two tablets at bedtime) ☐ Requested strength/dose is not commercially available ☐ Other: _____					
Are there any other comments, diagnoses, symptoms, medications tried or failed, and/or any other information the physician feels is important to this review? _____ _____ _____					

Please note:

This request may be denied unless all required information is received.
 For urgent or expedited requests please call 1-800-711-4555.
 This form may be used for non-urgent requests and faxed to 1-800-527-0531.

This document and others if attached contain information that is privileged, confidential and/or may contain protected health information (PHI). The Provider named above is required to safeguard PHI by applicable law. The information in this document is for the sole use of OptumRx. Proper consent to disclose PHI between these parties has been obtained. If you received this document by mistake, please know that sharing, copying, distributing or using information in this document is against the law. **If you are not the intended recipient, please notify the sender immediately.**
 Office use only: Zetia-ezetimibe_Comm_2018Mar-W

ANTIPSYCHOTICS **PRIOR AUTHORIZATION FORM** (form effective 1/5/21)



Keystone First

PERFORMRxSM
Next Generation Pharmacy Benefits

Fax to PerformRxSM at **1-215-937-5018**, or to speak to a representative call **1-800-588-6767**.

PRIOR AUTHORIZATION REQUEST INFORMATION					
☐ New request ☐ Renewal request		Total pages:	Office contact/phone:	LTC facility contact/phone:	
PATIENT INFORMATION					
Patient name:			Patient ID#:		DOB:
Street address:			Apt #:	City/state/zip:	
PRESCRIBER INFORMATION					
Prescriber name:					
Specialty:			NPI:	State license #:	
Street address:			Suite #:	City/state/zip:	
Phone:			Fax:		
MEDICATION REQUESTED					
Preferred Agents ☐ Abilify Maintena ☐ fluphenazine elixir ☐ haloperidol tablet ☐ Invega Sustenna ☐ Perseris ER injection ☐ risperidone tablet ☐ aripiprazole tablet ☐ fluphenazine oral concentrate ☐ haloperidol decanoate inj ☐ Invega Trinza ☐ quetiapine tablet ☐ trifluoperazine tablet ☐ Aristada ER injection ☐ fluphenazine tablet ☐ haloperidol lactate inj. ☐ loxapine capsule ☐ quetiapine ER tablet ☐ ziprasidone capsule ☐ Aristada Initio injection ☐ fluphenazine decan. inj. ☐ haloperidol lactate oral concentrate ☐ olanzapine tablet ☐ Risperdal Consta ☐ Zyprexa Relprevv ☐ clozapine tablet ☐ Haldol injection					
Non-Preferred Agents ☐ Abilify Mycite ☐ chlorpromazine tablet ☐ Geodon injection ☐ olanzapine inj/ODT ☐ Saphris SL tablet ☐ Versacloz suspension ☐ Abilify tablet ☐ clozapine ODT ☐ Haldol decanoate inj. ☐ olanzapine/fluoxetine cap ☐ Secuado patch ☐ Vraylar capsule ☐ Adasuve inhalation ☐ Clozaril tablet ☐ Invega ER tablet ☐ paliperidone ER tab ☐ Seroquel tablet ☐ Zyprexa tablet/injection ☐ amitriflyline/perphenazine ☐ Fanapt tablet ☐ Latuda tablet ☐ pimoizide tablet ☐ Seroquel XR tablet ☐ Zyprexa Zydys ☐ aripiprazole ODT ☐ Fazaclon dispersible tablet ☐ molindone tablet ☐ Rexulti tablet ☐ Symbyax capsule ☐ other: ☐ aripiprazole solution ☐ fluphenazine HCl injection ☐ Nuplazid capsule ☐ Risperdal solution/tablet ☐ thioridazine tablet ☐ Caplyta capsules ☐ Geodon capsule ☐ Nuplazid tablet ☐ risperidone ODT ☐ thiothixene capsule					
Strength:	Dosage form:	Directions:	Quantity:	Refills:	
Diagnosis:			Diagnosis code (required):		
PHARMACY INFORMATION (Prescriber to identify the pharmacy that is to dispense the medication):					
Deliver to: ☐ Patient's Home ☐ Physician's Office ☐ Patient's Preferred Pharmacy Name:					
Pharmacy Phone #:			Pharmacy Fax #:		
☐ I acknowledge that the patient agrees with the pharmacy chosen for delivery of this medication.					
REQUEST FOR A NON-PREFERRED AGENT					
1. Has the patient taken the requested non-preferred antipsychotic in the past 90 days? ☐ Yes – Submit documentation. ☐ No			2. Has the patient tried and failed the preferred medications (listed above)? ☐ Yes – List medications tried: ☐ No		
3. Does the patient have a contraindication or intolerance to the preferred medications? ☐ Yes – Submit documentation of contraindication/intolerance. ☐ No			4. For oral Invega/paliperidone ER requests, does the patient have active liver disease with elevated LFTs or is the patient at risk for active liver disease? ☐ Yes – Submit documentation and lab values. ☐ No		
REQUEST FOR A PATIENT LESS THAN 18 YEARS OF AGE					
5. Is this request for a dose increase of a previously approved medication? ☐ Yes – Submit recent chart documentation supporting the increased dose. ☐ No					
6. Is the requested agent prescribed by, or in consultation with, one of the following physician specialists? ☐ Yes ☐ No Submit documentation of consultation, if applicable. ☐ child development pediatrician ☐ child & adolescent psychiatrist ☐ general psychiatrist (only if patient is ≥ 14 years of age) ☐ pediatric neurologist					
7. Does the patient have severe behavioral problems related to a psychotic or neuro-developmental disorder? ☐ Yes – Submit medical record documentation. ☐ No					
8. Has the patient tried non-drug therapies? ☐ Yes – Submit medical record documentation. ☐ No					
9. Has the patient had the following baseline and/or follow-up monitoring? Check all that apply. ☐ BMI (or weight/height) ☐ blood pressure ☐ fasting glucose level ☐ fasting lipid panel ☐ presence of extrapyramidal symptoms (EPS) using the Abnormal Involuntary Movement Scale (AIMS) Submit documentation of all monitoring/test results.					
REQUEST FOR A LOW-DOSE ORAL ANTIPSYCHOTIC FOR A PATIENT 18 YEARS OF AGE OR OLDER					
10. What is the TOTAL daily dose of the requested medication? _____mg/day Submit documentation of complete medication regimen.					
11. Is the low dose prescribed as part of a plan to titrate up to a therapeutic dose? ☐ Yes – Submit documentation of titration plan. ☐ No					
REQUEST FOR THERAPEUTIC DUPLICATION OF AN ATYPICAL OR TYPICAL ANTIPSYCHOTIC					
12. Does the patient have a medical reason for concomitant use of the requested medications? ☐ Yes – Submit documentation with justification. ☐ No					
13. Is this request for a drug that is being titrated to, or tapered from, a drug in the same class? ☐ Yes – List medication. ☐ No					
PLEASE FAX COMPLETED FORM WITH REQUIRED CLINICAL DOCUMENTATION					
Prescriber signature:					Date:

Confidentiality Notice: The documents accompanying this telecopy may contain confidential information belonging to the sender. The information is intended only for the use of the individual named above. If you are not the intended recipient, you are hereby notified that any disclosure, copying, distribution or taking of any telecopy is strictly prohibited.

D Demand Estimation: Additional Details

To overcome computational hurdles, we estimate our nested demand system in Section 4 with a Poisson pseudo-maximum-likelihood estimation approach. This appendix describes 1) the justification for doing so and the estimation routine; and 3) the data processing required to make the data ready for estimation.

D.1 Estimating Nested Logit Demand Systems with Poisson Regression

We build on the equivalence of the likelihood functions of conditional multinomial logit estimation and Poisson regression. Readers interested in a deeper dive are encouraged to read [Guimarães et al. \(2003\)](#) and the references contained therein. That paper derives the equivalence between the two. We will instead briefly walk through the intuition.

Consider a conditional logit demand system for individuals i choosing a single good d from a choice set D . Individuals choose a good to maximize utility $u_{id} = \beta X_{id} + \epsilon_{id}$ for observed X_{id} . If ϵ is i.i.d. standard Gumbel distributed, then the probability that i chooses d is

$$P_{id} = \frac{\exp(\beta X_{id})}{\sum_{k \in D} \exp(\beta X_{ik})}$$

The sample analogue is c_{id} , the choice indicator vector which is 1 if i chose d and 0 otherwise. Typical estimation involves noting that, with conditional logit demand, $E[c_{id}] = P_{id}$, and rewriting this as a maximum likelihood problem. However, note that if we assert this equality and take logs of both sides, we have

$$\log(E[c_{id}]) = \beta X_{id} - \underbrace{\log \left[\sum_{k \in D} \exp(\beta X_{ik}) \right]}_{\alpha_i} \quad (4)$$

with the term α_i as a quantity that is constant across all goods within an individual. This is equivalent to the typical Poisson regression formulation, and therefore the coefficient β on X_{id} from an individual-level conditional logit can be estimated with an individual-product-level Poisson regression that includes X_{id} and individual-level fixed effects.

Further, imagine that instead of individual-level choices, we observe group-level market shares s_{gd} for a group of individuals g where $X_{id} = X_{i'd} = X_{gd}$ for all d and for all $i, i' \in g$. Note that $E[s_{gd}] = P_{gd}$, and so a group-level Poisson regression as formulated above will equivalently estimate β .

The classic alternative to this is the approach of [Berry \(1994\)](#). He notes that if one takes Equation 4 and difference out the expression for a reference good 0, one gets

$$\log(E[s_{gd}]) - \log(E[s_{g0}]) = \beta(X_{gd} - X_{g0})$$

and if one assumes that the Law of Large Numbers applies, then the observed shares $\hat{s}_{gd}$ are approximately equal to their expectations, $E[s_{gd}]$, and the econometrician can run a regression of the log share difference between the focal good and the reference good ($\log(\hat{s}_{gd}) - \log(\hat{s}_{g0})$) on the difference in characteristics

between them (and since the reference good is often an outside good with all characteristics set to zero, the regressors can simply be the characteristics of the focal good). These approaches are analogous. Berry's approach differences out the α_i from Equation 4.

The difficulty with this approach arises in finite samples, in two ways. First, the Berry approach will be biased in finite samples where $\hat{s}_{gd} \not\approx E[s_{gd}]$ and thus Jensen's inequality ensures that $E[\log(\hat{s}_{gd})] \not\approx \log(E[s_{gd}])$; the bias will be larger when this approximation is poorer: in smaller samples and/or when groups are smaller. Second, and more importantly in our application, in finite samples, as $P_{gd} \rightarrow 0$ for a good j , the probability of observing market shares of zero for that good becomes nontrivial. Indeed, in our setting, 98.7% of beneficiary-drug pairs have zero usage. In that case, $\log(\hat{s}_{gd})$ is undefined. In contrast, the Poisson regression approach is not biased in finite samples and can accept market share observations of zero.⁴²

In Section 4, we want to estimate a nested logit model rather than a conditional logit model, with a single nest incorporating all drug options, excluding the option of taking no drug. As a reminder, the utility function for the nested logit is:

$$u_{idt} = \underbrace{\beta_C \text{Auth}_{idt} + \delta_C \text{Excl}_{idt} + \kappa_{dm(it)}}_{V_{idt}} + \xi_{iC} \mathbf{1}\{d \neq 0\} + \lambda \epsilon_{idt}$$

where ϵ_{idt} and $\xi_{iC} \mathbf{1}\{d \neq 0\} + \lambda \epsilon_{idt}$ are Gumbel distributed and the choice probabilities are

$$P_{idt} = \underbrace{\frac{\exp \frac{V_{idt}}{\lambda_C}}{\sum_{k \in C} \exp \frac{V_{ikt}}{\lambda_C}}}_{P_{id|d \neq 0}} \times \underbrace{\frac{\left(\sum_{k \in C} \exp \frac{V_{ikt}}{\lambda_C} \right)^{\lambda_C}}{1 + \left(\sum_{k \in C} \exp \frac{V_{ikt}}{\lambda_C} \right)^{\lambda_C}}}_{P_{i(d \neq 0)}}$$

for inside goods with V_{idt} as the mean utility of good d for individual i in time t , and

$$P_{i0t} = \frac{1}{1 + \left(\sum_{k \in C} \exp \frac{V_{ikt}}{\lambda_C} \right)^{\lambda_C}}$$

for the outside good.

Berry (1994) shows that the nested logit demand system can be estimated via log-linear OLS by including $\log(s_{gd}/s_{g(d \neq 0)})$ as an additional regressor, with its estimated coefficient being equal to $1 - \lambda_C$. However, in settings where s_{gd} is zero, this regressor will be undefined. Therefore, we cannot use this approach. Instead, we estimate this model using a two-step approach: First, we estimate all of the mean utility parameters using the drug choice; then, we estimate λ_C using the choice of whether to consume a drug at all or not.⁴³

Specifically, we note that the nested logit utility can be divided by λ_C to get

⁴² Additionally, the Berry approach cannot be used on individual-level data, since the outcome variable will take on the value of zero for non-chosen goods.

⁴³ Train (2009) notes that this form of estimation is consistent but inefficient, since the across-nest choice is not incorporated into the estimation of the within-nest choice. In our case, since the across-nest choice only incorporates one additional alternative, which inherently cannot face prior authorization or exclusion, the two-step approach is unlikely to cause significant efficiency loss.

$$\underbrace{\frac{u_{idt}}{\lambda_C}}_{\tilde{u}_{idt}} = \underbrace{\frac{\beta_C}{\lambda_C}}_{\tilde{\beta}_C} \text{Auth}_{idt} + \underbrace{\frac{\delta_C}{\lambda_C}}_{\tilde{\delta}_C} \text{Excl}_{idt} + \underbrace{\frac{\kappa_{dm(it)}}{\lambda_C}}_{\tilde{\kappa}_{dm(it)}} + \frac{\xi_{iC} \mathbf{1}\{d \neq 0\}}{\lambda_C} + \epsilon_{idt}$$

Additionally, if we define a reference inside good, good 1, we can rewrite the above as

$$\tilde{u}_{idt} = \tilde{\beta}_C \text{Auth}_{idt} + \tilde{\delta}_C \text{Excl}_{idt} + \underbrace{(\tilde{\kappa}_{dm(it)} - \tilde{\kappa}_{1m(it)})}_{\tilde{\Delta\kappa}_{dm(it)}} + \tilde{\kappa}_{1m(it)} + \epsilon_{idt}$$

Since $\tilde{u}$ is a monotonic transformation of u , maximizing u is equivalent to maximizing $\tilde{u}$; additionally, since ϵ is standard Gumbel, then the probability of choosing d conditional on choosing an inside good (and conditional on a draw of ξ_{iC}) is

$$P_{id|d \neq 0} = \frac{\exp \frac{V_{idt}}{\lambda_C}}{\sum_{k \in C} \exp \frac{V_{ikt}}{\lambda_C}} = \frac{\exp \left(\tilde{V}_{idt} + \tilde{\kappa}_{1m(it)} + \frac{\xi_{iC} \mathbf{1}\{d \neq 0\}}{\lambda_C} \right)}{\sum_{k \neq 0} \exp \left(\tilde{V}_{ikt} + \tilde{\kappa}_{1m(it)} + \frac{\xi_{iC} \mathbf{1}\{d \neq 0\}}{\lambda_C} \right)} = \frac{\exp(\tilde{V}_{idt})}{\sum_{k \neq 0} \exp(\tilde{V}_{ikt})}$$

with $\tilde{V}_{idt} = \tilde{\beta}_C \text{Auth}_{idt} + \tilde{\delta}_C \text{Excl}_{idt} + \tilde{\Delta\kappa}_{dm(it)}$, and the third equality coming from the fact that $\tilde{\kappa}_{1m(it)}$ and $\frac{\xi_{iC} \mathbf{1}\{d \neq 0\}}{\lambda_C}$ are common to all inside goods and thus have no effect on choice probabilities.

The key factor here is that within a nest, the choice probabilities are standard logit and so can be treated as such. Moreover, since all of the remaining regressors are defined at the group level, we can estimate the group-drug-year-level Poisson regression:

$$\log(E[s_{gdt}]) = \tilde{\beta}_C \text{Auth}_{gdt} + \tilde{\delta}_C \text{Excl}_{gdt} + \tilde{\Delta\kappa}_{dm(it)} + \alpha_{gt}$$

where we regress group-drug-year-level market shares on dummies for prior authorization and exclusion, with drug-market and group-year fixed effects. This gives us estimates, $\hat{\beta}$, $\hat{\delta}$, and $\hat{\Delta\kappa}$, with α_{gt} as nuisance parameters.

We then have two remaining unknown parameters: λ_C and $\tilde{\kappa}_{1m(it)}$. Noting again that $\frac{V_{idt}}{\lambda} = \tilde{V}_{idt} + \tilde{\kappa}_{1m(it)}$, the probability of a member of g choosing any drug (compared to no drug) is

$$P_{g(d \neq 0)} = \frac{\left(\sum_{k \in C} \exp \frac{V_{gkt}}{\lambda_C} \right)^{\lambda_C}}{1 + \left(\sum_{k \in C} \exp \frac{V_{gkt}}{\lambda_C} \right)^{\lambda_C}} = \frac{\left(\sum_{k \neq 0} \exp(\tilde{V}_{gkt} + \tilde{\kappa}_{1m(gt)}) \right)^{\lambda_C}}{1 + \left(\sum_{k \neq 0} \exp(\tilde{V}_{gkt} + \tilde{\kappa}_{1m(gt)}) \right)^{\lambda_C}}$$

Taking the log of both sides, we see that

$$\log P_{i(d \neq 0)t} = \kappa_{1m(gt)} + \lambda_C \hat{\mathcal{V}}_{gt} + \omega_{gt}$$

with $\hat{\mathcal{V}}_{gt} = \log \left(\sum_{k \neq 0} \exp(\tilde{V}_{gkt}) \right)$, the inclusive value of the inside goods, and a group fixed effect $\omega_{gt} = -\log \left(1 + \left(\sum_{k \neq 0} \exp(\tilde{V}_{gkt} + \tilde{\kappa}_{1m(gt)}) \right)^{\lambda_C} \right)$. Additionally, the choice probability of the outside good (no drug) is

$$\log P_{g(d \neq 0)} = \omega_g$$

Therefore, we can estimate $\kappa_{1m(gt)}$ and λ_C by running a Poisson regression at the group-option-year level, with options being either taking any drug or taking no drug; with the outcomes as group market shares, and the regressors being a market-level intercept for the ‘any drug’ option, the inclusive value $\mathcal{V}$ interacted with an indicator for the ‘any drug’ option, and group-class-year fixed effects. Once we have done this, all relevant parameters have been estimated.⁴⁴

D.1.1 Instrumental Variable Estimation

Our approach requires us to instrument for the prior authorization and exclusion status of a drug in the plan the beneficiary was *enrolled in* with the same from the plan they were *assigned to*. Instrumental variables approaches are tricky in nonlinear estimation. We use the control function approach of [Petrin and Train \(2010\)](#). This is further complicated by the fact that we estimate our model in two stages, both of which require a control function at each stage.

To estimate the inner nest choice (i.e., the choice of drug conditional on choosing any drug), we first run the regression:

$$\begin{bmatrix} \text{Auth}_{idt}^{\text{Enrolled}} \\ \text{Excl}_{idt}^{\text{Enrolled}} \end{bmatrix} = \gamma^1 \begin{bmatrix} \text{Auth}_{idt}^{\text{Assigned}} \\ \text{Excl}_{idt}^{\text{Assigned}} \end{bmatrix} + \vec{K}_{dm(it)} + \vec{u}_{idt}^1$$

i.e., a linear regression of dummies for formulary status in the enrolled plan on the same dummies in the assigned plan, plus drug-market fixed effects. We can then recover the estimated residuals,

$$\hat{u}_{idt}^1 = \begin{bmatrix} \text{Auth}_{idt}^{\text{Enrolled}} \\ \text{Excl}_{idt}^{\text{Enrolled}} \end{bmatrix} - \left(\hat{\gamma}^1 \begin{bmatrix} \text{Auth}_{idt}^{\text{Assigned}} \\ \text{Excl}_{idt}^{\text{Assigned}} \end{bmatrix} + \hat{K}_{dm(it)}^1 \right)$$

and include them as a control in the Poisson regression on drug choice market shares.

For the outer choice model (the choice of drug or no drug), we must also account for endogeneity: specifically, the endogeneity of the inclusive value $\mathcal{V}$, which governs the inclusive value of the formulary the beneficiary faces. To account for this, we run the following regression:

$$\mathcal{V}_{it}^{\text{Enrolled}} = \gamma^2 \mathcal{V}_{it}^{\text{Assigned}} + K_{1m(it)} + u_{it}^2$$

the linear regression of the inclusive value estimated for the plan of enrollment on the inclusive value of the plan of assignment (only having an effect for the ‘any drug’ choice), with a market-level fixed effect.

We can then construct the estimated residuals from this regression,

$$\hat{u}_{it}^2 = \mathcal{V}_{it}^{\text{Enrolled}} - \left(\hat{\gamma}^2 \mathcal{V}_{it}^{\text{Assigned}} + \hat{K}_{1m(it)} \right)$$

and use those as controls in the Poisson regression on the shares that choose any drug.

⁴⁴While we only estimated versions of β , γ , and $\Delta\kappa$ that were normalized by λ , the normalized parameters are sufficient to compute counterfactual simulations. They can be retransformed back into their non-normalized forms if need be.

The control function approach allows us to control for the extent of deviation of beneficiaries away from their assigned formulary. The coefficient on the residuals from the ‘first stage’ in the choice model capture the extent to which beneficiaries who endogenously select into plans with different coverage than their default do so because they prefer specific drugs that they are deviating to fill more easily.

One feature of this approach is that the largest group of beneficiaries that can be constructed with modeled homogeneity within the group is at the enrolled-plan-by-assigned-plan level; therefore, this is the group g that we use.

D.1.2 Estimation Routine

To summarize, our procedure is, for each therapeutic class:

1. Restrict to only inside good options (i.e., exclude beneficiaries in a plan who took no drug in the class), and construct a dataset of group-year drug choice shares for drugs within the class, where groups are enrolled-plan-by-assigned-plan pairs.
2. Run the ‘inner choice first stage’ linear regression of dummies for formulary status in the enrolled plan on dummies for formulary status in the assigned plan and drug-market fixed effects:

$$\begin{bmatrix} \text{Auth}_{gdt}^{\text{Enrolled}} \\ \text{Excl}_{gdt}^{\text{Enrolled}} \end{bmatrix} = \tilde{\gamma}_C^1 \begin{bmatrix} \text{Auth}_{gdt}^{\text{Assigned}} \\ \text{Excl}_{gdt}^{\text{Assigned}} \end{bmatrix} + \vec{K}_{dm(gt)} + \vec{u}_{gdt}^1$$

to estimate the group-by-drug-by-year residuals $\hat{u}_{gdt}^1$.

3. Run the ‘inner choice second stage’ Poisson regression of group-year drug choice shares on dummies for the prior authorization and exclusion status of the drug in the *enrolled* plan, drug-market fixed effects, plan-year fixed effects, and the estimated residuals from above:

$$\log(E[s_{gdt}]) = \beta_C \text{Auth}_{gdt}^{\text{Enrolled}} + \delta_C \text{Excl}_{gdt}^{\text{Enrolled}} + \Delta \kappa_{dm(gt)} + \alpha_{gt} + \zeta_C^1 \hat{u}_{gdt}^1$$

4. Take the estimated parameters β_C , δ_C , and $\Delta \kappa_{dm(gt)}$, and use them to construct the inclusive values $\mathcal{V}$ for all plans in every year.
5. Construct a dataset with two observations for each plan-year, one containing the share of beneficiaries taking any drug in the class, the other containing the share of beneficiaries taking no drug in the class.
6. Run the ‘outer choice first stage’ linear regression of the inclusive value for the plan the beneficiary enrolled in on the inclusive value for the plan they were assigned to, plus a market fixed effect interacted with a dummy indicating the ‘any drug’ choice:

$$\mathcal{V}_{gt}^{\text{Enrolled}} = \gamma^2 \mathcal{V}_{gt}^{\text{Assigned}} + K_{1m(gt)} + u_{gt}^2$$

to estimate the group-by-choice-by-year residuals $\hat{u}_{gt}^2$.

7. Run the ‘outer choice second stage’ Poisson regression of group-year choice shares (drug or no drug) on the inclusive value of the enrolled plan, a market fixed effect, and the residuals estimated in the prior step, all interacted with a dummy indicating the ‘any drug’ choice, as well as a group-year fixed effect:

$$\log(E[s_{gDt}]) = [\lambda_C \mathcal{V}_{gt}^{\text{Enrolled}} + \kappa_{1m(gt)} + \zeta_C^2 \hat{u}_{gt}^2] \times \mathbf{1}\{D = 1\} + \omega_{gt}$$

where $D = 0$ reflects “no drug” and $D = 1$ reflects “any drug.”

This approach makes clear how λ_C is identified, and how it reflects the extent of intensive vs. extensive margin substitution. The components of $\mathcal{V}_{gt}$ are identical across groups g within a region and year *except* for the formularies they face; the demand parameters are otherwise identical. λ_C is identified from the extent to which plans with more stringent formularies characterized by greater use of prior authorization and exclusion (and thus lower inclusive values) have less use of any drug. When λ_C is close to zero, only intensive margin substitution matters: When beneficiaries are deterred from one drug, they will substitute to another, leaving the share of beneficiaries taking any drug constant. In contrast, when λ_C is close to one, beneficiaries will substitute to other options proportionally, and thus most beneficiaries who are deterred from a drug will move to no drug.

To estimate the Poisson regressions, we use the Poisson pseudo-maximum-likelihood estimation method developed by [Correia et al. \(2020\)](#) that allows for fast estimation of Poisson regression models with high-dimensional fixed effects. For ease of computation, we estimate this model separately for each therapeutic class.

D.2 Standard Errors

Since our estimation procedure has multiple steps, and we want our standard errors to incorporate the variation in estimators that can come from noise in any particular step, the ideal is to bootstrap the entire procedure described above. However, our estimation procedure relies on many fixed effects which are sparsely estimated, i.e., the number of observations pinning down the fixed effect is quite small. This is especially true with many of our drug-market-year fixed effects. If we cannot observe any individual taking the drug in that market-year, we will be forced to estimate the fixed effect at $-\infty$. With a standard bootstrap, the odds of this occurring for any given drug-market-year are nontrivial. This will cause our confidence intervals to necessarily be too large for some estimators, driven by computational issues rather than true variation.

Instead, we use the Bayesian bootstrap ([Shao and Tu 1995](#)).⁴⁵ Instead of resampling units with replacement, we instead, for each unit, draw random weights at each bootstrap run, and re-estimate the model with these weights applied. The distribution of parameter estimates from each run serves as our estimated sampling distribution of the parameter. That work suggests that an appropriate weight for each individual can be drawn from the exponential distribution with scale parameter 1. To speed up computation, we draw this at the group-by-drug-by-year level rather than the individual-by-year level, which we can do since the sum

⁴⁵We thank Peter Hull for alerting us to the Bayesian bootstrap’s suitability for this purpose.

of exponentially-distributed random variables has a Gamma distribution. If n individuals from group g in year t were observed, the appropriate weight is $w_{gt} \sim \text{Gamma}(n, 1)$.⁴⁶

For each therapeutic-class-specific drug demand estimation routine, we replace the Poisson pseudo-maximum-likelihood method with a **weighted** pseudo-maximum-likelihood estimator, using the drawn weights. We use 500 bootstrap runs, and preserve the weights within a run across classes, so that within a given bootstrap run, the same weights are being used to compute therapeutic-class-level market shares and spending and thus correctly aggregate across classes. Standard errors for a parameter (or function of a set of parameters) are estimated as the standard deviation of that parameter over the 500 estimated bootstrap runs.

D.3 Data Processing for Demand Estimation

Since we take a discrete choice approach to modeling drug demand, estimating such a model requires data formatted as a discrete choice. However, since our analysis is at the level of a year, this is naturally often violated: A patient may take multiple drugs in a given year, especially to satisfy step therapy requirements. In the first column of Appendix Table A18, we report the share of beneficiaries who took multiple drugs in a given year within a class (conditional on taking at least one drug). Across classes, this averages to 15.1% of beneficiaries, but ranges from 0% to 51.8%. To transform this into an appropriate dataset, we pick, for each beneficiary-year, the modal drug within the class they took that year (as defined by the drug consumed with the most days supply, breaking ties randomly), and assign that as their ‘chosen’ drug for the year. Column two of Appendix Table A18 reports, for each class of the top 30 by gross spending, the share of days supply that the assigned drug made up across beneficiary-year pairs who filled multiple drugs within a class for a year. Appendix Figure A8 plots the distribution of this multiple-drug-user share across classes. On average, across all classes, the assigned drug made up 63.9% of days supply for these beneficiaries and 90% for all beneficiaries. Appendix Figures A9 and A10 plot the distributions of these values across classes.

The identification of all of our demand parameters requires that any market (region-year) in a particular class must have at least one drug that faces prior authorization in at least one (but not all) plans in that market; otherwise, β cannot be identified from behavior in that market. Additionally, in a similar vein, it must be true that at least two drugs are ever taken; if not, β is not identified separately from λ , since both will influence inside drug vs. no drug choice.


In the third and fourth columns of Appendix Table A18, we list the share of markets that violate at least one of the two above requirements (both as a share of market-years and weighted by beneficiary counts) for the top 30 therapeutic classes by spending. In Appendix Figures A11 and A12, we plot the distribution of the unweighted and weighted shares. A sizable number of classes have very high shares of markets that do not contribute to identification. In testing, these classes tended to be ones where β was estimated with the wrong sign (i.e., we estimated that, for that class, prior authorization *increased* use of a focal drug), and ones where λ was estimated at values well outside the $[0, 1]$ interval that we would expect it to lie on. We therefore decide to only use classes where no more than 10% of markets violate at least one of the two requirements.

⁴⁶Note that the expected value of w_{gt} is n , which is the expected number of times one would draw a member from the group in a standard bootstrap approach.


Appendix Table A18: Class Level Nested Logit Summary Statistics for Top 30 Classes by Part D Spending

Class	Unweighted Market Survival	Weighted Market Survival	Share of Focal Days Supply, Multiple Drug Users	Share of Multiple Drug Users
Antihyperlipidemic Drugs, NEC	94.3%	98.7%	59.5%	23.9%
Psychother, Tranq/Antipsychotic	95.2%	99.0%	60.4%	31.9%
Antidiabetic Agents, Insulins	95.7%	99.1%	60.8%	51.8%
Gastrointestinal Drug Misc, NEC	97.1%	99.3%	72.8%	22.3%
Antivirals, NEC	89.5%	98.4%	36.6%	25.5%
Antidiabetic Agents, Misc	96.7%	99.3%	57.5%	25.5%
Antineoplastic Agents, NEC	92.9%	99.1%	67.3%	6.6%
Biological Response Modifiers	79.5%	96.7%	72.5%	4.1%
CNS Agents, Misc.	97.1%	99.3%	61.8%	10.5%
Psychother, Antidepressants	97.1%	99.2%	58.8%	32.6%
Adrenals & Comb, NEC	95.2%	99.0%	78.9%	25.9%
Analg/Antipyr, Opiate Agonists	94.3%	99.1%	66.3%	24.0%
Cardiac Drugs, NEC	96.2%	99.3%	73.0%	20.2%
Antiplatelet Agents, NEC	81.4%	94.8%	63.2%	12.5%
Immunosuppressants, NEC	91.0%	98.5%	61.9%	15.4%
Misc Therapeutic Agents, NEC	95.2%	99.1%	59.8%	24.8%
Anticonvulsants, Misc	91.4%	98.8%	57.0%	19.0%
Cardiac, Calcium Channel	93.8%	98.6%	71.6%	10.5%
Coag/Anticoag, Anticoagulants	88.6%	95.6%	85.5%	15.7%
Cardiac, Beta Blockers	90.0%	98.6%	71.9%	7.2%
Parasympathomimetic, NEC	84.3%	95.6%	72.2%	7.0%
Eye/Ear/Nose/Throat Misc, NEC	90.5%	98.7%	53.7%	35.0%
Analg/Antipyr, Nonstr/Antiinflm	96.2%	99.2%	71.1%	22.2%
Muscle Rel, Smooth-Genitour NEC	95.7%	99.2%	74.7%	15.3%
Antiinflam Agents EENT, NEC	96.7%	99.3%	67.2%	18.7%
Sympathomimetic Agents, NEC	92.9%	99.0%	73.4%	8.7%
Estrogens & Comb, NEC	90.0%	97.8%	74.4%	7.9%
Vasodilating Agents, NEC	71.0%	93.0%	79.1%	18.5%
Phosphorus Removing Agents, NEC	73.3%	94.1%	68.8%	7.9%
Cardiac, ACE Inhibitors	70.0%	90.5%	72.3%	5.3%


Notes: For each class listed, this table displays the share of markets (region-year pairs) that have (1) at least one drug in the class that is under a prior authorization restriction in between 0 and 100% of plans; and (2) where at least two drugs in the class are filled. The first and second columns give this statistic, the second weighted by beneficiary count within our sample. The fourth column lists the share of beneficiary-years who fill at least two drugs within the class in a given year, out of those who fill at least one drug. The third column lists the share of days supply made up by the most-used drug in the class, for this subpopulation of beneficiaries. Table is sorted by total Part D spending within our sample.

Appendix Figure A8: Share of Drug-Users Whom Take Multiple Drugs


Notes: For each therapeutic class, we measure, out of the set of beneficiary-year pairs where the beneficiary took at least one drug within the class in that year, how many beneficiary-year pairs were ones in which the beneficiary took at least two drugs in the class. This figure plots the distribution of that statistic across classes.

Appendix Figure A9: Focal Drug Days Supply Share for Multi-Drug Users


Notes: For each therapeutic class, we measure, out of the set of beneficiary-year pairs where the beneficiary took at least two drugs within the class in that year, what share of days supply in that class were accounted for by the focal (most-used) drug. This figure plots the distribution of that statistic across classes.

Appendix Figure A10: Focal Drug Days Supply Share for All Drug Users

Notes: For each therapeutic class, we measure, out of the set of beneficiary-year pairs where the beneficiary took at least one drug within the class in that year, what share of days supply in that class were accounted for by the focal (most-used) drug. This figure plots the distribution of that statistic across classes.

Appendix Figure A11: Unweighted Market-Year Survival After Logit Restrictions

Notes: This figure plots the distribution, across classes, of the share of markets (region-year pairs) that have (1) at least one drug in the class that is under a prior authorization restriction in between 0 and 100% of plans; and (2) where at least two drugs in the class are filled.

Appendix Figure A12: Beneficiary-Weighted Market-Year Survival After Logit Restrictions

Notes: This figure plots the distribution, across classes, of the share of markets (region-year pairs) that have (1) at least one drug in the class that is under a prior authorization restriction in between 0 and 100% of plans; and (2) where at least two drugs in the class are filled. In this figure, markets are weighted by the number of beneficiaries represented in our sample.

E Revealed Preference Analysis: Additional Details

E.1 Deriving Consumer Surplus Loss

In Section 6.1, we use demand curves of the form $D(p) = D(0)e^{\frac{\epsilon}{100}p}$. In this Appendix, we briefly derive a closed-form expression for consumer surplus.

We note that consumer surplus over the range Θ is

$$CS = \int_{\Theta} V_d(q) dq = \int_{\Theta} D^{-1}(q) dq$$

with the second equality due to the fact that we have assumed that $V_d(\theta) = W_d(\theta)$, and $W_d(\theta) = D^{-1}(\theta)$.

Note that since $D(p) = D(0)e^{\frac{\epsilon}{100}p}$, then $D^{-1}(q) = \frac{100}{\epsilon} \log\left(\frac{q}{D(0)}\right)$. Given this, we can also note that

$$\int \frac{100}{\epsilon} \log\left(\frac{q}{D(0)}\right) dq = \frac{100}{\epsilon} \left[q \log\left(\frac{q}{D(0)}\right) - q \right]$$

In the section, we generally take integrals over regions of the form $[aD(0), bD(0)]$. For such a region, the integral is therefore

$$\begin{aligned} \int_{aD(0)}^{bD(0)} \frac{100}{\epsilon} \log\left(\frac{q}{D(0)}\right) dq &= \frac{100}{\epsilon} \left[q \log\left(\frac{q}{D(0)}\right) - q \right]_{aD(0)}^{bD(0)} \\ &= \frac{100D(0)}{\epsilon} [b \log(b) - a \log(a) - (b - a)] \end{aligned}$$

noting that, while this antiderivative is undefined at $q = 0$, $\lim_{q \rightarrow 0} q \log(q) - q = 0$. Note that this measure of consumer surplus is linear in $D(0)$, as well as being linear in the reciprocal of ϵ , the semi-elasticity of demand.

At the end of that section, we relax the assumption that willingness-to-pay is equal to value, and replace it with $W_d(\theta_{id}) = \rho V_d(\theta_{id})$ for $\rho \in (0, 1]$, which is equivalent to $\frac{W_d(\theta_{id})}{\rho} = V_d(\theta_{id})$. Consumer surplus is now

$$CS^{Debiased} = \int_{\Theta} V_d(q) dq = \int_{\Theta} \frac{W_d(q)}{\rho} dq = \frac{1}{\rho} \int_{\Theta} D^{-1}(q) dq = \frac{1}{\rho} CS$$

i.e., debiased consumer surplus is linear in the ‘rational’ consumer surplus measure. Note that if we want to find the ρ such that net welfare is zero, we need

$$\begin{aligned} NFS + \Delta CS^{Debiased} &= 0 \\ NFS + \frac{1}{\rho} \Delta CS &= 0 \\ \frac{-\Delta CS}{NFS} &= \rho \end{aligned}$$

Noting that ΔCS is negative so the term on the left will be positive.

E.2 Provider-Based Revealed Preference Approach

Here we detail an alternative approach to measuring revealed preference through provider actions. In this approach we assume that decisions about which prescription drug the patient will consume are made entirely by the provider.

Consider a provider deciding whether to prescribe restricted drug d to patient i . Assume providers care about their own costs, but also put altruistic weight on the patient's preferences, such that provider utility is

$$u_{id} = \rho v_{id} - a$$

where ρ is the weight the provider places on patient preferences and a is the administrative cost of fulfilling a prior authorization request, where applicable. The provider will prescribe drug d if $\rho \Delta v_{id} = \rho V_d(\theta_i) \geq a$, resulting in a demand curve $D(a)$ that depends on administrative costs, with $D(a) = \int 1\{W_d(\theta_i) \geq a\} d\theta$, with $W_d(\theta_i) = \rho V_d(\theta_i)$, the willingness-to-do-paperwork (an analogue to willingness-to-pay).

If, as in Section 6.1, $W_d(\theta_i)$ is drawn from a zero-inflated exponential distribution with scale parameter $\frac{\epsilon}{100}$ and a mass at zero of $1 - D(0)$, then, as in the prior section, this structure gives rise to a demand curve that depends on administrative costs, $D_d(a) = D(0)e^{\frac{\epsilon}{100}a}$ for ϵ , the semi-elasticity of drug demand with respect to administrative costs. Under this structure, the demand curve for drugs once again reveals patient valuations for the drug; although, in this case, it specifically reveals how physicians value patient value for the drug. To simplify, we begin by assuming that physicians are perfectly altruistic in that they weight their patient's preferences equal to their own, i.e., $\rho = 1$.

To estimate the administrative cost semi-elasticity, we simply use the demand response to prior authorization restrictions that we observe in Sections 3 and 4. In response to prior authorization, providers prescribe restricted drugs 28.9% less. Our baseline calibration of provider-facing cost is \$22. These two numbers imply that the administrative cost semi-elasticity of prescription is $\epsilon = \frac{28.9}{22.48} = 1.29$. By this calibration, providers are several times more elastic to administrative costs relative to patients' elasticity to out-of-pocket prices.

As established in the prior section, the implied consumer surplus loss is inversely proportional to the elasticity of demand. Therefore, the loss estimated from this approach will be smaller than the loss estimated from the beneficiary-centered approach. We once again compute consumer loss under two screening scenarios: the best-case and the random case. Under those three assumptions, the consumer surplus loss is \$2 and \$9 respectively.

These measures assume $\rho = 1$; however, we have no guarantee that physicians act in the best interests of their patients per se. It might be that physicians weight their own costs to a relatively greater extent than the value for their patients. We do not have a specific estimate of ρ . Instead, we can once again find the values of ρ that would make prior authorization restrictions generate utilitarian welfare losses on net. ρ has a stronger economic interpretation in this case: When ρ is low, providers care little about their patients' welfare, and therefore policymakers should not enforce screening mechanisms that make them responsible for allocating drugs to patients. For the best-case and random case scenarios, the maximum ρ to make prior authorization inefficient is 0.02 and 0.11 respectively.