Rationing Medicine Through Bureaucracy: Authorization Restrictions in Medicare*

Zarek Brot† Samantha Burn‡ Timothy Layton§ Boris Vabson¶

November 9, 2024

Abstract

We study the trade-off between spending reductions and bureaucratic costs associated with managed care tools in healthcare. We leverage random assignment to Medicare Part D plans to study this trade-off for prior authorization restrictions. Prior authorization reduces a drug's utilization by 26.9%. Half of marginal beneficiaries are diverted to another related drug, while the other half are diverted to no drug. The removal of these policies would increase drug spending by $96 per beneficiary-year (3.6% of drug spending), while eliminating $10 in paperwork costs.

*We thank Frank Gaunt and Michelle Gard 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, Dan Waldinger, Anna Zink, the editors, Erzo F.P. Luttmer and John Friedman, four anonymous reviewers, and seminar participants at Columbia, Chicago Booth, Chicago Fed, Johns Hopkins, Notre Dame, NYU, UC Berkeley, UC San Diego, UVA Batten, 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 and NBER. Email: zarek@uchicago.edu

‡Imperial College London. Email: s.burn@imperial.ac.uk

§University of Virginia and NBER. Email: timothyjlayton@virginia.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 purported savings often proposed as a way to 'pay for' eligibility expansions and increases in generosity of public programs.[1]

This discussion is at odds, however, with a simple fact: 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 in healthcare 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 and reduce the use of low-value care. 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 may 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 improvements in social surplus for marginal patients who are deterred. Understanding the welfare consequences of these policies therefore requires measuring the size

[1] One 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 reforms would reduce administrative expenses by 50-60%, with this reduction having no effect on other outcomes (Pollin et al. 2018, Friedman 2019).

and composition of these marginal and inframarginal groups.

We study prior authorization empirically in Medicare Part D, the public drug insurance program for the 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. 2023). 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 common, especially among more expensive drugs: 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 enrolled in plans that have authorization restrictions on that drug against utilization for beneficiaries enrolled in 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, and it passes a large suite of balance tests.

We estimate that prior authorization restrictions reduce the (extensive margin) use of focal drugs by 26.9%, with slightly larger relative effects among non-white and older patients, and smaller relative effects on drugs in high-benefit classes. We estimate similar effects for spending (21.9%) as well as prescription days (33.8%) of the focal drug. We use the same regression specification to estimate that roughly half of utilization deterred by prior authorization is diverted to other drugs in the same therapeutic class. However, since restricted drugs tend to be much more expensive than unrestricted drugs, this results in only 13.5% of the spending reduction on the focal drug being offset by spending on substitute drugs.

Our results indicate that prior authorization restrictions clearly reduce the amount payers must spend to provide health insurance. 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. We study this trade-off empirically via a counterfactual simulation comparing drug spending and administrative costs under the status quo set of prior authorization policies from 2007-2015 to an alternative world where prior authorization restrictions were banned. To perform this counterfactual simulation while accounting for interactions between restrictions on substitute drugs, we estimate a nested logit discrete choice model of drug demand. Our simulations indicate that the use of status quo prior authorization restrictions lowered drug spending by roughly $96 per beneficiary-year, 3.6% of spending, relative to a counterfactual where prior authorization was disallowed. However, these restrictions 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 administrative costs. Instead, we calibrate per-application administrative costs and request rejection rates from prior studies and combine

them with counts of the number of restricted drugs used to estimate the size of the burden. Under our preferred calibration, we estimate that the administrative burden of the status quo prior authorization restrictions is roughly $10 per beneficiary-year, around 10% of the spending increase. While the costs of bureaucracy are non-trivial, our simulations suggest that they are second-order relative to the effects on utilization, and eliminating prior authorization would be cost-*increasing* rather than cost-decreasing. This result, however, *only* applies to the regime of restrictions that was implemented. It does not necessarily apply to *any* possible use of prior authorization. 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 appear to have been imposing 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 the value of the forgone drugs is low enough to justify such policies. Estimating consumer valuation of forgone drug consumption is a difficult exercise, given that economists have previously documented under-consumption of some drugs due to behavioral frictions (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. Our estimates of the demand response to this price change suggest that prior authorization is approximately equivalent (in terms of its effects on demand) to charging $441 more per prescription. This price exceeds the paperwork cost of prior authorization in a year by an order of magnitude.

We use this demand slope to infer patients' willingness-to-pay for forgone drugs. 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 $55.53 and $206.60 per beneficiary-year, respectively, though our estimates are smaller using alternative specifications. To the extent that this reflects consumer surplus, it says that for prior authorization restrictions (as applied in the status quo) to improve social welfare, they need to do a very good job of screening out use by those with low value. We construct various measures of the components of net social welfare, and our results vary across these approaches.

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 construct a measure of patient-level aggregate exposure to prior authorization that quantifies the share of drugs that the

patient took in the prior year that are restricted in their plan. We find inconclusive effects on health with low statistical precision. We cannot statistically reject that prior authorization generates large negative health effects, but we also cannot reject the possibility that prior authorization significantly *improves* beneficiary health.

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 substitute 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, willingness to pay is far larger.

Irrespective of whether bureaucracy in health care is an overall force for good or for bad, our results suggest that its *effects on spending on prescription drugs* are of greater orders of magnitude than the direct costs of its operation. This implies that, when considering its overall effect, 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 work exists to measure the latter; the closest to our study is Dunn et al. (2023), 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 2024), 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 Notowidigdo 2019, Homonoff and Somerville 2021, Shepard and Wagner 2023), 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. Indeed, contemporaneous work on restrictions in physical therapy suggests that the beneficial effects of screening largely come from the monitoring rather than the ordeal (Gandhi and Shi 2024).

[2] A related literature has studied the effects of bureaucratic institutions that deal with contractual incompleteness on prices and quantities in other public procurement settings (Bajari et al. 2014).

[3] There is also a small existing literature on the quantity effects of prior authorization policies for prescription drugs (Seabury et al. 2014, Sarig 2024). 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.

[4] This 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 transfer itself (Lieber and Lockwood 2019, Waldinger 2021).

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. 2022). Our work suggests that non-price rationing mechanisms may be provide a promising alternative.

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.[5] 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.[6]

Prior authorization restrictions are generally applied to discrete services.[7] 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

[5] In our sample period, forms were primarily sent via fax for legal compliance reasons. In later years, some providers were able to obtain electronic prior authorization assistance integrated into their IT apparatus.

[6] Both the forms and the sequence of events required for authorization for drugs, or for other services, are broadly similar for insurers in other U.S. insurance market segments.

[7] Since 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. (2023) and Shi (2024) for studies of such mechanisms.

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.[8] 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 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.[9]

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 cost.[10] 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.[11] In this setting, that will

[8] The 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.

[9] Note that screening need not be the only motivation for prior authorization. Insurers may also impose restrictions to discourage users of the restricted drug from enrolling in their plan (Geruso et al. 2019). They may impose restrictions on rival drugs as a reward for rebate payments (Brot-Goldberg et al. 2022, Ho and Lee 2023).

[10] We note that, for a complete welfare analysis, this "cost" should be the social cost of the drug. This would generally be the marginal cost of production of the drug, not the price of the drug paid by Medicare or some other insurer. However, given that Medicare is a tax-financed program, we could also think of the "cost" as the marginal cost of funds required to finance the price paid by Medicare. For this section, we remain agnostic about the precise definition of the cost, and we come back to this question when we attempt to analyze welfare in Section 5.

[11] We 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.

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.[12]

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 neither patients nor providers 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. 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 inframarginal patient, i.e. those who get the drug despite the presence of restrictions.[13]

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 $d\theta$):

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

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.[14]

This welfare change has three components. First, patient surplus is reduced, since the program is moving

[12] We note that this assumes that there is no uncertainty about the drugs that a given beneficiary will need over the course of the next year. With uncertainty, there would be some insurance value to leaving drugs unrestricted, as, under prior authorization restrictions, the beneficiary would have a worse outcome in the "bad" state in which they need the restricted drug versus the outcome in the bad state without prior authorization. The removal of prior authorization thus plays a small role in equating marginal utilities in the good and bad states. The assumption of no uncertainty seems reasonable here where most drug consumption is fairly persistent and predictable over time, suggesting little uncertainty regarding drug consumption and thus little insurance value. If, however, we assessed the uncertainty from "behind the veil of ignorance," insurance value would make up a larger part of welfare (Hendren 2020).

[13] In 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 4.

[14] We assume $\Theta_1 \subset \Theta_0$, i.e. that there are no 'defiers' who get the drug only when authorization restrictions are in place.

them away from their most-preferred choice to another option that they value less.[15] 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 3, 4, and 5 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 unlikely to work well when there are many inframarginal patients relative to the number of marginal patients, 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 patients, 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 (as the next alternative is likely to be

[15] Some prior authorization restrictions are imposed 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.

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.[16] Consumers choose from the plans offered in their service region (of which there are 34), 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 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.[17] 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.[18] This differentiation is substantial. Taking

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

[17] A different group of 'partial LIS' beneficiaries receive lesser subsidies, but are omitted from our analysis.

[18] Note 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.

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. (2023) 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, 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 beneficiary actively chooses a non-default option.

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 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. 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.[19]

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 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.[20] 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.[21] 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

[19] We opt for these definitions 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. Additionally, note that our definition of exclusion is given by non-inclusion, so it is possible that some drugs we designate as being excluded are covered but this coverage is not reported by the plan to CMS. Our claims data includes covered claims for drugs we designate as excluded, which may either reflect mis-classification or insurer-granted exceptions.

[20] Drug 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.

[21] The 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.

types whose coverage status we would not be able to track in formularies whatsoever.[22] 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 two groups of LIS beneficiaries who faced the automatic reassignment mechanism described in Section 2.1: (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 and (2) those whose prior plan exited the market entirely. We focus on these beneficiaries, rather than new Medicare enrollees, for two key reasons: (1) we observe a full year of post-randomization data for them, and (2) for this group we can observe pre-assignment data, providing a useful outcome for placebo tests and useful information about the set of drugs demanded by the beneficiary that we use in some analyses. We exclude beneficiaries whose reassignment is expected to be non-random based on program rules.[23] 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 differs from the full LIS population due to our sample restrictions intended to isolate those who are randomized to plans. These restrictions result in our sample being 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.[24] Our estimates will thus apply most directly to this selected population, though we expect the results to generalize due to the similarities between the sample and the general LIS population. 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 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.

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

[23] For 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.

[24] Note that, in our analysis, we retain those who qualify for Medicare due to disability rather than old age, lowering the average age of our population.

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 restrictions imposed 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. This is not explained by some insurers being more prone to

using prior authorization restrictions than others. Within drug-region-years, only 0.8% of residual variation in the use of prior authorization is explained by carriers. Even if we consider carrier-by-therapeutic class variation, this only explains 7.8% of residual variation.

We cannot definitively explain the reasons for the remaining variation, but we provide two potential explanations. First, it is plausible that the rebate concession motive (wherein manufacturers offer larger rebates in order to avoid restrictions) might generate separating equilibria where insurers pursue different strategies (e.g., one insurer caters to manufacturer A while another insurer caters to manufacturer B), leading to heterogeneity in formularies. Second, some of the variation may come from the 'path to equilibrium,' as different insurers observe the restrictions imposed by their competitors and respond. Even if we might expect a symmetric equilibrium, it may take time for insurers to reach this point, generating within-drug-market-year variation in restrictions across plans along the way.

3 The Effect of Authorization Restrictions on Drug Utilization and Spending

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

To estimate the effect of prior authorization on drug utilization, we leverage variation in prior authorization restrictions across drugs and across plans. Generally, both of these sources of variation could be correlated with potential outcomes and thus lead to biased estimates of the effects of the restrictions. First, beneficiaries can choose their plans. This could clearly lead to the composition of beneficiaries differing across plans with different types of restrictions, with the most obvious potential endogenous sorting problem coming from the possibility that beneficiaries avoid plans that restrict the drugs they intend to use. Second, drugs with prior authorization restrictions differ in many ways from unrestricted drugs, leading to differences in demand for restricted versus unrestricted drugs even in a world where all drugs are unrestricted. As we showed in Section 2.4, drugs have a higher chance of facing prior authorization when they are more expensive and less commonly used. Both of these patterns would lead us to overestimate the effect of prior authorization if we naively estimated the association between restriction and use.

We use separate strategies to deal with the potential endogeneity of each source of variation in prior authorization restrictions. To deal with beneficiary selection into plans, we use the random assignment of defaults discussed in Section 2. The beneficiary's assigned default is, by construction, orthogonal to their underlying drug preferences, removing any compositional differences of beneficiaries across plans. We restrict to *only* beneficiaries who faced randomization, due to their previous year's plan either losing benchmark status or exiting from the market. 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. Since assignment is random *within a market* (a service

region-by-year pair), we conduct our primary analysis within-market by interacting all fixed effects with market-year fixed effects. To deal with selection of drugs into restricted status, we conduct all analyses *within-drug* by including drug-by-market-by-year fixed effects. This absorbs any secular differences in use across drugs that may be correlated with the propensity to face prior authorization.

Our approach is similar in spirit to a difference-in-differences design, where the two dimensions of comparison are drug and plan. To build intuition, consider a case where there are two plans (1 and 2) and two drugs (d and d'). Drug d is never restricted in either plan 1 or in plan 2, and drug d' is restricted only in plan 2 but not in plan 1. Drug effects are identified by comparing utilization of drugs d and d' in plan 1, where neither drug is restricted. Plan effects (when included) are identified by comparing utilization of drug A among those randomly assigned to plan 1 versus plan 2. The effect of prior authorization restrictions is identified by comparing use of drug d' in plan 2 (where it is restricted) versus use of drug d' in plan 1 (where it is unrestricted) and subtracting any plan effects (i.e., the use of drug d in plan 2 versus plan 1).

Identifying the effects of prior authorization in this way requires two key assumptions. First, we need default plan assignment to be as good as random within markets so that the cross-plan comparisons of use of a given drug reflect plan effects on the use of the drug rather than differences in the composition of beneficiaries across plans. Random assignment to plans satisfies this assumption, and we verify it with balance tests. Second, we need within-drug variation in prior authorization across plans to be uncorrelated with any *other* actions taken by plans that affect use of the restricted drug. Fortunately, plans have only a few tools with which they can influence drug utilization, and the primary tool aside from prior authorization (exclusion) is observed and can be controlled for (which we do). The main threat to this assumption is that plans also differ in their (observable) restrictions on *other drugs*. For example, for two drugs d, d' that treat similar illnesses, if formularies are designed strategically, their formulary statuses will likely be correlated, and restrictions on d' will encourage use of d. We account for this by explicitly including controls for the formulary status of therapeutic substitutes. The ideal approach would be to control for the formulary status of each potential substitute, but this would require many more unique formulary arrangements than are present in the data. We instead construct a single control that measures the weighted share of all other drugs in the same therapeutic class as the focal drug that face an authorization restriction in the assigned plan, with weights equal to the substitute drug's market share in the entire sample in that year. We also include a similar control for formulary exclusion of substitute drugs. Plans have virtually no other tools to influence drug utilization other than prior authorization and exclusion, making other potential violations of this assumption unlikely, but we also test robustness of our main estimates to stronger plan-level controls, such as plan fixed effects and plan fixed effects interacted with various drug characteristics, that remove variation but allow for weaker assumptions regarding the presence of unobservable plan actions.[25]

Our final system of estimating equations is

[25] In these specifications with both plan fixed effects and controls for restrictions on potential substitute drugs, plan effects need to be identified via cross-plan comparisons of use of drugs that are unrestricted in all plans *and* that are not substitutes for drugs that are restricted in any plan. This is a high bar, but we show that our main results are still robust to the inclusion of these fixed effects. We do not include plan fixed effects in our main specification, however, as they complicate the identification of cross-drug spillovers that we study in Section 3.4. We prefer to maintain a single primary specification across all outcomes for simplicity.

$$
\begin{aligned}
Y_{idt} &= \beta_1 \text{Auth}_{idt}^{\text{Enrolled}} + \beta_2 \text{Excl}_{idt}^{\text{Enrolled}} + \kappa_{dm(it)} \qquad (1)\\
&+ \gamma_1 \text{Auth}_{j(it)dt}^{\text{Sub,Assigned}} + \gamma_2 \text{Excl}_{j(it)dt}^{\text{Sub,Assigned}} + \nu_{idt}
\end{aligned}
$$

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

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 fixed effects, and our set of controls for restrictions on substitutes. Our target parameter to be estimated is β_1, the average effect of prior authorization on drug utilization relative to unrestricted coverage. 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. For all regressions, we cluster standard errors at the assigned plan and year level.

We note that this regression is effectively equivalent to a two-step approach. In the hypothetical first step, we would, for each drug, compare utilization for those randomly assigned to a plan that restricts that drug against utilization for those assigned to a plan that covers the drug with no restrictions. In the second step, we would take the weighted average of these drug-specific differences. This produces an estimate of the average effect of prior authorization on drug utilization purged of selection bias. The implicit weights in the OLS estimation of our effects reflect how precisely each drug-specific difference is estimated (Gibbons et al. 2019). Specifically, the weights are equal to $w_{idt} = p_{dm(it)}(1 - p_{dm(it)})$, where p_{dm} is the probability that a beneficiary in market m will be assigned to a plan that restricts drug d. For instance, if a drug never faces prior authorization, we can never estimate the effect on its use; if 50% of beneficiaries face prior authorization, we can estimate drug-specific effects most precisely. The weighted average treatment effect, with these weights, is estimated more precisely than the unweighted average treatment effect.[26]

In Table 4, we report estimates of our first stage, given by Equation 2. We measure the beneficiary's enrolled plan as of December 31 of year t. The instrument is extremely strong, with first-stage F-statistics in the tens of thousands. We estimate that assignment to a plan that restricts a drug increases the probability of facing a restriction on that drug by 0.91, consistent with the fact that 91% of beneficiaries enroll in the plan that they are assigned to.[27]

We perform three sets of balance tests to verify that beneficiary formulary assignment is conditionally random. First, 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, we estimate the 'effect' of prior authorization restrictions on utilization outcomes *in the year prior to assignment*. Finally,

[26] All specifications we present in this paper use these weights, either implicitly or explicitly.

[27] In the right panel of Table 4, we address the concern that our first-stage is strong because beneficiaries are happy to comply with defaults for drugs they do not plan to take by restricting to only drugs that the beneficiary used in the prior year. Reassuringly, the associated coefficients are only slightly smaller than those estimated without this restriction.

we estimate the 'effect' of prior authorization restrictions on beneficiary characteristics (gender, race, age, and Elixhauser Comorbidity Index) which should not have any relationship with the assignment mechanism. We report the results from these tests in Appendix Tables A3, A4, and A5. Reassuringly, in all three cases we can reject even extremely small effects.

3.2 Main Estimates

With our research design established, we can now estimate the effect of prior authorization restrictions on utilization. We focus on three different measures of utilization: A binary indicator for whether the beneficiary filled the drug at least once in the year (multiplied by 100 to reflect percentage point changes), a count of the total days supply of the drug filled in the year, and total allowed spending on the drug (net of rebates).

As stated above, our approach is equivalent to an average over many within-drug comparisons. Before running regressions, we demonstrate some of these comparisons to show that the effects of prior authorization are apparent even when looking at the data in its rawest form. In Figure 5a, for a subsample of drugs, we plot the share of beneficiaries who filled the drug at least once in the year (after residualizing market fixed effects), comparing those assigned to a plan that restricted the drug versus those who faced unrestricted coverage of the drug.[28] Unsurprisingly, for each drug, when a beneficiary faces a prior authorization restriction, they are less likely to ever fill the drug. The regression estimates we discuss below can be viewed as weighted averages over many of these drug-specific comparisons.

In Table 5, we present our regression estimates. Generally, the absolute magnitude of the estimated effects is quite small; for example, prior authorization's effect on whether a beneficiary ever uses a drug is less than one percentage point. However, this small coefficient is a consequence of the fact that most beneficiaries do not take most drugs, and so most observations have zero utilization. Therefore, we additionally compute the percent change relative to the (weighted) average utilization of the drug without prior authorization.[29] Taking the baseline utilization of drugs into account, we find that the reduction in utilization induced by prior authorization restrictions is substantial: Prior authorization restrictions on a drug reduce the probability a beneficiary will use it by 26.9%, prescription days filled by 33.8%, and spending by 21.9%. These effects are quite large, and refute claims that high prior authorization approval rates mean that prior authorization has little significance for actual utilization. All of these estimates are fairly similar in magnitude, suggesting that most of the response is on the margin of taking the drug or not, rather than the margin of quantity taken conditional on any use.

In Appendix Table A7, we explore robustness of the estimates of effects of prior authorization on the

[28] To select the drugs in the subsample, we generate a score for each drug equal to the product of the average variation in prior authorization across markets (akin to the weights discussed in Section 3.1) and the share of all beneficiaries in the sample who ever fill the drug. Given that the regression will weight drugs according to the variation in prior authorization, this approximately orders the drugs according to their importance for our estimates below.

[29] Specifically, for each drug-market dyad, we first take average utilization by beneficiaries in the market assigned to plans that did not place restrictions on the drug. Then we take a weighted average over dyads, using the weights implicitly imposed by OLS estimation. As discussed in Section 3.1, the implicit weights are $w_{idt} = p_{dm(it)}(1 - p_{dm(it)})$: the probability that a beneficiary faces a restriction on drug d given their presence in market m in time t. This weighting is necessary: As the table shows, the unweighted control means are much larger than their weighted counterparts, since they, for instance, assign positive weight to frequently-used drugs that never face restrictions.

primary outcome, any use, to alternative specifications. Our main specification (which leaves out plan fixed effects) assumes that plans that are more likely to use prior authorization are not also more likely to take any other actions that affect drug utilization, in ways not captured by our controls for formulary status of therapeutic substitutes. By adding additional fixed effects, we can trade off power for weaker assumptions. By adding in plan fixed effects, we can allow for plans to differ in their propensity to impose restrictions *and* any other actions they take to influence drug use, as long as those actions only have constant effects on use across drugs (i.e., they do not differentially affect the use of specific drugs that are more likely to face restrictions). By adding plan-by-drug-type fixed effects (where we define types by therapeutic class or by decile of drug price), we can weaken this assumption further to allow plans to take actions that differentially affect types of drugs, as long as plans do not take actions that differentially affect drugs within a type. We show that our estimates are robust to these weaker assumptions, suggesting we should not be concerned about uncontrolled correlated unobservable plan actions. Our estimates are also robust to 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.

3.3 Heterogeneous Effects

Our results indicate that there are substantial average effects of prior authorization restrictions on consumption of prescription drugs. However, the stated goal of prior authorization is not to deter all types of drug consumption but instead to specifically deter *low-value* consumption. One approach to determining whether prior authorization restrictions are deterring the 'correct' drug consumption is to examine who is deterred, and what sort of consumption 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. In all heterogeneity analyses we analyze effects on utilization only using our measure of days supply filled by a beneficiary in a given year. The same qualitative patterns emerge when using other utilization measures.

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 of certain groups for reasons unrelated to the value of the drug (e.g. due to the patient's race or gender), prior authorization may deter care more strongly 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 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 6. We estimate 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.

We also measure differential effects by health status, segmenting beneficiaries by their Elixhauser Comorbidity Index score, which measures how many chronic conditions they had in the prior year. We estimate

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. 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 by drug type. These estimates are displayed in Figure 8. In some cases, prior authorization is used for safety 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 typically 'scheduled' 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. (2023).[30] Encouragingly, restrictions bind less tightly on these drugs, suggesting that prior authorization restrictions may be at least modestly well-targeted as a rationing mechanism.[31]

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.

3.4 Substitution Responses

Our results show that prior authorization reduces the use of restricted drugs. But do beneficiaries who are deterred from those drugs substitute to a different drug or go without any drug at all? To understand how prior authorization affects total spending on drugs, the answer to this question is critical: While substitution

[30] These include anticonvulsants, antidiabetic agents, antihyperlipidemic drugs, cardiac drugs, oral anticoagulants, antipsychotics, and antidepressants.

[31] We also estimate heterogeneous effects by deciles of drug price per day supply (Appendix Figure A2), by therapeutic class (Appendix Figure A3), by service region (Appendix Figure A4), and by insurance carrier (Appendix Figure A5). In all three cases, we find some heterogeneity, but not with any clear economic pattern, and we cannot statistically reject joint equality across categories. We also find roughly uniform effects when segmenting by decile of how frequently a drug is restricted (Appendix Figure A6), giving us some confidence that our treatment effect estimates—which are identified based on variation in restriction status for drugs that are sometimes, but not always, restricted—generalize to drugs that are *always* restricted. While we find much larger responses for rarely-restricted drugs, these come with large standard errors.

to no drug entails savings equal to the cost of the restricted drug, substitution to an alternative drug entails savings equal to that amount minus the cost of the alternative drug. Additionally, the intention of many applications of prior authorization is to shift demand from high-cost drugs to lower-cost alternatives; understanding substitution patterns allows us to assess its ability to do just that versus causing patients not to use any drug at all.

We start by providing evidence of the *existence* of substitution effects and proceed to estimate the *extent* of those effects. To demonstrate the existence of substitution, Figure 5b presents results similar to those presented in Figure 5a but focusing on substitution instead of use of the focal drug. Specifically, for a subset of drugs, we compare plans that restrict the drug to those that don't. However, in this figure we plot the use of *other drugs* in the same therapeutic class as the listed drug for beneficiaries randomly assigned to plans that place prior authorization restrictions on the listed drug versus beneficiaries randomly assigned to plans that do not restrict the listed drug (residualizing market fixed effects). We see that beneficiaries who face restrictions on the listed drug (who, as shown in Figure 5a, are *less* likely to consume the listed drug) are *more* likely to consume other drugs in the same class as the listed drug, relative to beneficiaries who do not face restrictions on the listed drug. These results clearly indicate some degree of substitution due to prior authorization restrictions.

Measuring the extent of substitution is more complicated. The ideal approach to measuring substitution would be to estimate the effect of prior authorization for every focal drug-substitute drug pair. This, however, is infeasible. Estimating the direct effect of prior authorization on the focal drug only requires that we observe some plans that restrict the drug and some that do not, i.e., two unique formularies. In contrast, even if we assumed that substitution only occurs between drugs within the same class, when there are D drugs in a class, we would need to observe $D + 1$ unique formularies (in terms of their restrictions on drugs within the class) to estimate each drug pair's cross-effect. In practice, we never observe enough unique formularies in our data. Therefore, rather than estimating each cross-effect, we target estimation of the aggregate diversion ratio; that is, of the use of the focal drug deterred by prior authorization, what share is shifted to a substitute drug (and, thus, what share is shifted to taking no drug at all)?

To estimate this diversion ratio, we replicate our 2SLS estimation of Equation 1 with outcome variables that reflect the utilization of substitute drugs. Our two primary outcomes are (1) the use of any drug in the focal drug's therapeutic class *other than* the focal drug; and (2) the use of any drug in the focal drug's therapeutic class, *including* the focal drug. We report the results in Table 6. As in the prior table, interpreting the parameter estimates directly is difficult given the small magnitudes due to infrequent use of any particular drug. Instead, we can interpret proportional effects. To compute the diversion ratio, we take the ratio of the effect on other drugs to the effect on the focal drug.[32]

For both of our quantity measures, the diversion ratio is roughly one-half, implying that half of drug consumption deterred by prior authorization is made up for by consumption of a therapeutic substitute, with the other half being diverted to no drug in the class. This implies a substantial level of extensive margin

[32]One exception is when we measure the outcome of whether the beneficiary ever took a drug within the year. Computing the diversion ratio in this way leads to a deceptively small number, coming from the fact that beneficiaries sometimes take multiple drugs within a class in a given year (i.e., the effect on the focal drug and the effect on other drugs do not add up to the class-level effect). For this quantity measure, we measure diversion as 1 minus the ratio of the effect on all drugs in class relative to the effect on the focal drug.

substitution to no drug use. This may seem somewhat surprising, given that many authorization restrictions are designed to induce substitution by taking the form of "step therapy" where the intention is for patients to experiment with (typically cheaper) substitute drugs before attempting the (more-expensive) focal drugs. However, the fact that we find such high levels of substitution could also seem surprising when contrasting this result with prior results on the effects of cost-sharing on healthcare utilization, where substitution is uncommon and virtually *all* substitution is on the extensive margin (Newhouse and the Insurance Experiment Group 1993, Brot-Goldberg et al. 2017). Our explanation for these results is that extensive margin effects 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.[33] Substitution in this way (to no drug) may be undesirable unless treatment for the associated condition was otherwise of very low value.

In contrast, the estimated diversion ratio for spending is much smaller, at 13.5%; that is, 13.5% of the reduction in spending due to reduced use of the focal drug is offset by substitution to other drugs, though our regressions have relatively larger standard errors due to the high variance in spending among some substitute drugs. The fact that this diversion ratio is much smaller than those for our quantity measures reflects the fact that restricted drugs are much more expensive than unrestricted drugs. Overall, the average effect of prior authorization on a single drug is to reduce the quantity of drugs in that class used by 0.4-0.7%, while reducing spending by 1%.

In Figure 7, we explore beneficiary heterogeneity in the extent of diversion. Interestingly, those groups who see the biggest reductions in utilization due to prior authorization also are the ones who are most likely to be diverted to a substitute drug. For instance, sicker patients have slightly larger direct effects of prior authorization, but are much more likely to take another drug. For instance, of beneficiaries with 5 or more chronic conditions who stop taking a drug due to prior authorization, 78.6% of their use is diverted to a therapeutic substitute. In contrast, for those with no chronic conditions, only 37.9% of use is diverted. We also estimate significantly larger diversion for men (78.4%) than women (58.5%).

4 Spending Reductions vs. Administrative Costs

In Section 1.2 we showed that a key trade-off when assessing the welfare consequences of prior authorization is the comparison of the cost savings due to prior authorization and the associated administrative costs. Many providers believe that prior authorization's effects are small relative to the administrative costs it brings.[34] In this section, we attempt to assess this trade-off empirically.

In this analysis, we focus on a simple comparison: What is the effect of the use of prior authorization in the status quo on total resources allocated to spending on prescription drugs and administrative costs, relative to a counterfactual world in which *all* currently restricted drugs were unrestricted. This is distinct from our reduced form estimates, which focus on the effects of placing prior authorization restrictions on a

[33] In 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.

[34] For example, a vice president at the American College of Physicians called prior authorization "the No. 1 burden to physician practice and a major impediment to the patient-physician relationship and patient care." (American College of Physicians 2022)

single drug, holding all else constant. This global approach represents the policy relevant comparison, where the relevant policy is one that would ban prior authorization restrictions. Importantly, simply scaling up our reduced form estimates may fail to account for interactions between restrictions on drugs. For instance, restricting drugs A *and* B is likely to have different effects than the sum of the effects of restricting each drug individually. This is because restricting A may induce substitution to B (and vice versa). These interactions may be important for understanding the full extent of the effects of allowing prior authorization as an institution. Further, turning restrictions off for one drug at a time (the effect captured in our reduced form approach) may induce substitution to another restricted drug, thus not reducing administrative costs. As discussed earlier, we lack the variation to cleanly identify such rich substitution patterns, making this empirically difficult without imposing additional structure.

Thus, we proceed by explicitly modeling demand for drugs using a standard microfounded model of demand for restricted and unrestricted drugs (and their substitutes) based on utility maximization. We then use the estimated model to simulate removing all prior authorization restrictions and measure the effects on utilization, spending, and administrative costs. Such an approach requires us to make additional data restrictions and modeling assumptions relative to our reduced-form approach, but we show that these choices have relatively minimal effects via comparisons to our reduced-form estimates.

4.1 Estimating Drug Demand

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 in different markets to have different preferences across drugs ($\kappa_{dm(it)}$). 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 V_{i0t}, the mean utility of the outside option of getting no drug ($d = 0$), to zero.

Finally, we assume that variation in preferences across beneficiary-provider pairs within a market are governed by two random factors. The first, ϵ_{idt}, represents idiosyncratic preferences of pair i for drug d in time t. The second, ξ_{iCt}, represents the preference of i for *any* drug in class C, relative to taking no drug in the class. We assume that ϵ_{idt} is distributed standard i.i.d. Gumbel, whereas ξ_{iCt} is distributed with variance that depends on an unknown parameter $\lambda_C \in [0, 1]$, such that the demand system takes on the familiar nested logit form (Berry 1994), with a nest including all drugs within the class, and another nest including only the outside option to take no drug.

In this model, λ_C governs how much of the variation in preferences is driven by class-specific preferences relative to drug-specific preferences. For instance, if variation is driven primarily by ξ_{iCt}, then

beneficiary-provider pairs choose a drug within the class because they have a strong preference for the *class* of drugs; therefore, when a drug is restricted, beneficiaries are likely to react on the intensive margin by taking a different drug within the same class. In contrast, if variation is driven primarily by ϵ_{idt}, beneficiary-provider pairs choose drug d because they have a strong preference for the *specific drug*, and we should expect greater extensive margin responses wherein the beneficiary shifts to taking no drug at all.

Identifying β_C and λ_C relies on the same variation used in Section 3. 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. λ_C is identified from the same variation, but measuring the difference in the market shares of all other drugs in the same class. We once again use the formulary of the assigned default plan as an instrument for the formulary of the plan that the beneficiary enrolls in.

Since our model assumes a discrete choice, we assign each beneficiary to a unique choice of drug within a given class equal to the drug in that class that they filled during that year. When a beneficiary has filled prescriptions for multiple drugs, we assign them to the drug for which they received the most days supply during the year, and break ties randomly.[35] Additionally, as we are estimating many class-specific parameters in this nonlinear model, to separately identify these parameters we need to restrict our sample to classes that have sufficient variation in prior authorization restrictions. Specifically, in order to identify both parameters we need to observe at least two drugs ever being taken (one focal drug, and one substitute drug), with at least one being a drug that faces prior authorization for some plans but not other plans. For every therapeutic class and market, we measure whether both of these criteria are satisfied, and we drop any classes where this is not true for more than 10% of markets. Our final dataset includes classes making up 97.8% 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. To do so efficiently, we exploit the equivalence between the likelihood functions of the conditional logit and the Poisson generalized linear model (Guimarães et al. 2003), and estimate the model using recently-developed techniques in high-dimensional Poisson pseudo-maximum-likelihood estimation (Correia et al. 2020). We implement the control function approach of Petrin and Train (2010) to instrument for formulary status, and compute standard errors using a Bayesian bootstrap procedure. This estimation approach is described in detail in Appendix D. In that Appendix, we also discuss the intuition behind the Poisson regression approach to logit demand estimation.

4.2 Total Spending Effects of Prior Authorization Restrictions

We use our estimated model to 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. To compute drug spending in these simulations, we combine our demand model with drug prices, assuming that the 'price' of any given drug is equal to the empirical average total paid amount (net of rebate) across all filled prescriptions of that

[35] On average, 15.9% 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, 63.6% 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 89.8% of total days supply.

drug during the same year across beneficiaries in our sample filling at least one prescription of that drug during that year. We can then compare these two simulations to measure how outcomes would change were policymakers to remove all prior authorization restrictions.

In Table 7, we measure the effects of moving from simulation (2) to (1) on per capita spending and utilization for all drugs, restricted drugs only, and unrestricted drugs only.[36] Our results suggest that prior authorization policies reduced drug spending by 3.6%, or approximately $96 per beneficiary-year. This spending reduction is composed of a $112 reduction in spending on restricted drugs, and an offsetting $15.7 increase in spending on (much cheaper) unrestricted drugs. Reassuringly, the results of our estimated effects on quantities line up closely to our reduced-form estimates. Our structural approach also suggests that prior authorization reduces the use of restricted drugs by roughly one-quarter, and half of those beneficiaries substitute to another drug.

4.3 Measuring Administrative Costs

We next contrast these spending results against the administrative costs induced by the prior authorization process. Unfortunately, we do not have data on the back-and-forth process between providers and insurers. Therefore, unlike prior studies, we cannot directly estimate the costs associated with the prior authorization process from accounting data (Shi 2024) or revealed preference (Dunn et al. 2023). Instead, we calibrate relevant parameters (per-application costs and rejection rates) based on estimates from the health policy literature and combine these calibrations with our demand system estimates to estimate the total administrative costs generated by compliance with prior authorization restrictions.

To do this, we first assume that any beneficiary wishing to fill a prescription for a restricted drug must receive authorization for that drug once per year.[37] We assume that making a request incurs some constant joint cost a to the requesting physician and the insurer. Because the number of requests is unobserved, we assume that any patient we observe taking a restricted drug must have made an authorization request. However, we also recognize that there are some who made a request but were rejected and who we thus do not observe taking the restricted drug. We assume that this happens with a constant rejection rate r across all drugs and years. Given r, a, and the number of patients taking a restricted drug, we know that there were $\frac{N}{1-r}$ requests and total administrative costs are $\frac{aN}{1-r}$.[38]

We use the estimated demand system to compute 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. We calibrate the per-request administrative cost a and rejection rate r collected in prior studies. 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

[36] The use of restricted drugs is defined at the beneficiary-drug level, i.e., a drug may be 'restricted' in the status quo for some beneficiaries and 'unrestricted' for others based on the beneficiary's plan.

[37] In general, authorization is typically required once per treatment course, though beneficiaries cannot usually carry over their authorization from a previous insurer.

[38] We abstract from repeat interactions between the requesting physician and the insurer.

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.[39] Our preferred estimate is from CAQH (2013), the study covering the largest number of providers. Their estimate is $18.53 per application.[40] 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 manual processing costs of $3.95 for insurers in 2012.[41] Adding these insurer costs to the preceding estimates of provider costs 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. Unfortunately, however, none of them are easily comparable to our setting; some studies are too narrow in that they cover a single, potentially unrepresentative area of care, while others are too broad in that they include unrelated services (e.g. hospital services and physician-administered drugs), and none precisely get at the exact quantity of interest–the number of (unobserved) requests per (observed) successful fill. Nevertheless, we calibrate a range of rejection rates reported in these studies. 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%.

In Table 8, we report, for every pair of calibrated values of a and r, the estimated total administrative costs from prior authorization per beneficiary-year. Unsurprisingly, higher calibrated values of a and r increase the administrative burden and reduce the net financial savings from prior authorization. Using our preferred calibrated measure of a ($22.48 per request) and our preferred measure of r (4%, an intermediate rate) we estimate administrative costs to be $9.76 per beneficiary-year.

4.4 Net Financial Effects

We have estimated that prior authorization reduces spending by $95.88 per beneficiary-year, while, under our preferred calibration, imposing $9.76 per beneficiary-year in administrative costs. This implies that, on net, the net financial consequences of prior authorization during the years we study result in a *decrease* in total spending of $86.12 per beneficiary-year. We can compare $95.88 to the value in other cells of Table 8 to simulate financial effects under alternative calibrations. In nearly all of our calibrations, prior authorization generates net financial savings; even at implausibly extreme calibrations—per-application costs of $200, which are 10x most estimated values, and rejection rates of 15%, well above most estimates—prior authorization roughly breaks even on financial grounds.

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,

[39] Another study, 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.

[40] 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.

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

we construct the ratio of spending reductions per dollar of administrative costs. In Appendix Table A10 we provide this ratio for the set of all drugs; values above 1 imply net financial savings due to prior authorization, 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-administrative-cost ratios in Figure 9, 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, although prior authorization restrictions resulted in net financial savings for the drugs selected by insurers for these restrictions, savings are not guaranteed under every conceivable application of prior authorization restrictions. Our result is specific to the way that prior authorization is applied under the status quo. Prior authorization would not necessarily achieve similar savings for drugs not selected for restrictions. We explore this point explicitly in Appendix Tables A11 and A12 by replicating Tables 7 and 8 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. Moreover, restricting currently-unrestricted drugs would result in substantially more diversion of marginal beneficiaries deterred from those drugs (90.8% of marginal beneficiaries) to non-use of any drug. 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 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 financial savings *only if targeted appropriately*. This exercise also suggests that historically, prior authorization was targeted reasonably well across drugs.

5 Welfare Effects of Prior Authorization on Beneficiaries

Our results in Section 4 suggest that prior authorization restrictions generate net financial savings even when taking administrative costs born by inframarginal patients' providers and insurers into account. We can thus conclude that, in contrast to previous discourse about bureaucracy, the actual paperwork costs are second-order relative to reduced spending due to these policies.

As discussed in Section 1.2, a full welfare analysis would consider not only the net savings produced by prior authorization but also changes in consumer surplus due to the reduction in utilization of restricted drugs. To re-emphasize this point, we reproduce our welfare equation from Section 1.2 here:

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

We have estimated the sum of the second and third terms to be \$86.[42] Thus, the welfare consequences clearly depend on how the reduction in patient surplus due to prior authorization restrictions compares to the net savings induced by those restrictions. The welfare consequences of prior authorization thus turn on the value of the first term, the reduction in patient surplus.

For a variety of reasons, it is difficult to estimate this lost patient surplus. The two primary approaches for estimating patient surplus involve (1) inferring patient valuation from patient choices (revealed preference) and (2) inferring valuation from estimates of health effects combined with the value of a statistical life-year. In our setting, (1) is difficult due to the fact that LIS beneficiaries do not face non-zero prices for filling prescriptions, limiting our ability to assess patient willingness-to-pay for drugs. Further, when patients do face positive prices for healthcare, there is substantial evidence that demand for drugs and services may not reflect clinical or private value (Baicker et al. 2015, Brot-Goldberg et al. 2017, Chandra et al. 2021). Similarly, (2) is difficult due to the fact that the typical beneficiary-level measure of health that is available is mortality, and a marginal reduction in drug utilization, especially for the categories of drugs which are frequently restricted, may take years to have a meaningful effect on mortality rates.

Despite these difficulties, we argue that it is a useful exercise to attempt to estimate lost patient surplus using variations on these methods in order to (1) provide guidance for how one might evaluate welfare consequences of this type of "rationing via bureaucracy" and (2) provide a kind of benchmark of what might be a reasonable guess for the reduction in patient surplus due to prior authorization. We do so in this section.

5.1 Revealed Preference Approach

A common approach to estimating lost consumer surplus relies on consumer choices trading off drugs and money to estimate their willingness-to-pay (WTP) for those drugs. Under the assumption that this WTP reflects their welfare-relevant valuation of the drug, we can then compute the lost surplus as a result of prior authorization. Estimates of WTP come from first estimating the demand curve for restricted drugs, then integrating under that demand curve.

As is typical, to trace out the demand curve for drugs we require exogenous variation in the price of drugs. We rely on a separate natural experiment originally used by Gross et al. (2022) to provide this variation. 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 29,733 beneficiaries

[42] As noted in Section 1.2, this is not quite accurate for a complete welfare analysis because the second term is not the reduction in program costs but the reduction in social costs. We come back to this at the end of this section.

into the LIS program as a source of exogenous variation in drug prices and estimate the demand response to that variation in prices. The transition to the LIS provides us with two points on the demand curve for drugs: one at the observed out-of-pocket price in the absence of the LIS and one at a price of $0. We use these two points to trace out the demand curve between those two prices and then extrapolate to other prices.

We follow Gross et al. (2022) and restrict to those who transition into the LIS program between the years 2007-2015. In Appendix Table A13, we provide summary statistics for this population. We observe each of these beneficiaries two years prior to their transition, the year of transition, and two years post-transition. We estimate the following regression at the person-drug-year level:

$$Y_{idt} = \alpha_i + \gamma_{dt} + \zeta \times \mathbf{1}\{t - s >= 0\} + \epsilon_{idt} \tag{3}$$

The α_is are beneficiary fixed effects, and the γ_{dt}s are drug-by-year fixed effects. ζ is interpreted as the change in utilization due to the transition to the LIS. We estimate the regression for four outcomes (at the beneficiary-drug-year level): the out-of-pocket payment per prescription, a dummy variable for whether the person filled any prescription for the drug during the year, total days supply of medication and total spending. Our goal is to *only* use the variation in price stemming from the transition to the LIS to identify the slope of the demand curve, not the variation in drug prices across plans in the pre-LIS period. Therefore, when measuring out-of-pocket payments we reprice all drugs in the pre-LIS period to be the average out-of-pocket payment across all prescription fills for the drug during the year among beneficiaries not in the LIS.[43] The elasticity we estimate thus reflects a weighted average of drug-specific elasticities. To ensure that the elasticity is comparable to our estimates of the effects of prior authorization restrictions in Section 3, we reweight drugs in these regressions to match the weighting in those regressions, using the weights described earlier in Footnote 26. Our estimates of the response to the LIS transition will be unbiased as long as there are no other changes that affect demand that occur around the time of any particular individual's transition into the LIS program.[44]

We present the estimated coefficients from these regressions in Table 9. We estimate that, upon entering the LIS program, the average out-of-pocket payment per prescription drops by around $192.28, reflecting the high out-of-pocket costs for restricted drugs. In response to this price decrease, the probability of filling any prescription for a given drug increased by 0.047 percentage points (an increase of 18.3%), days supply filled increased by around 0.089 drug-days (a 23.7% increase), and net expenditures increased by $0.57 (a 38% increase). This implies that placing a prior authorization restriction on a drug, which reduces the probability of use by 26.9%, is equivalent to increasing the copayment from $0 per prescription to $441 per prescription.

[43] The "correct" price to use here is not immediately obvious, given the non-linear price schedule. A rational consumer would respond to the expected end-of-year price. However, there is substantial evidence that consumers also respond to 'spot' prices (Brot-Goldberg et al. 2017, Abaluck et al. 2018, Dalton et al. 2020), so we opt to use the actual spot prices paid by consumers.

[44] There are two complications which potentially bias our estimates of ζ. First, the LIS transition also lowers the price of potential substitute drugs. This shifts the demand curve for the focal drug 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 for d 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.

5.1.1 Forgone Beneficiary Value

We can then use our above estimates to infer the demand curve for drugs. We begin by assuming a single linear demand curve that encompasses all drug use. We infer the slope of the (inverse) demand curve by dividing our estimates of the LIS program's impact on out-of-pocket costs by its effect on utilization. We infer the x-intercept (quantity demanded when price is zero) as the average utilization of drugs among those in the LIS in plans with no prior authorization, as measured by the "reweighted control mean" statistic given in Table 5.

A linear demand curve implies that WTP is uniformly distributed among those who might plausibly purchase the drug, with a mass below zero for those who do not fill the drug even when it is free and unrestricted. To estimate how much consumer surplus is lost due to prior authorization, we simply need to identify the set of marginal beneficiaries — those who would consume the drug in the absence of prior authorization, but would not when it is restricted. Importantly, we need to be able to measure the distribution of WTP conditional on being in this marginal group. Put another way, we need to know where the marginal beneficiaries lie on the demand curve. However, the relationship between WTP and responsiveness to prior authorization is not identified in our data. One possibility is that the beneficiaries turned away by prior authorization do so because they have little value for the drug (or their physician has decided they have little value), and thus also have the lowest WTP. In this case, the total amount beneficiaries would be willing to pay for drugs forgone due to prior authorization would be equal to

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

where $D_d^{-1}(\theta)$ is the height of the demand curve, reflecting WTP for a consumer of type θ, and $D_d(0)$ is the x-intercept, the share of beneficiaries who consume the drug under a price of zero and no prior authorization restrictions. Δq_d here is reduction in quantity used of drug d due to prior authorization. Since we estimate that prior authorization reduces the share of beneficiaries who use the drug by 26.9%, we would integrate over the 26.9% of beneficiaries with the lowest WTP. We categorize this as the "best-case" option since, given the estimated demand curve, it is the lower bound on total WTP for the foregone drugs.

However, responsiveness to price may not be perfectly reflective of responsiveness to prior authorization. For instance, physicians play a major role in the prior authorization process, and their beliefs about the value of a given drug for a given beneficiary may differ from the patient's own values. Similarly, physicians may differ in their administrative capacity to deal with the authorization process (Gandhi and Shi 2024), and patient matching to different physicians need not be related to their WTP for specific restricted drugs. Therefore, which beneficiaries get screened out by prior authorization may be *completely orthogonal* to their WTP. We therefore also consider another scenario that reflects this "random screening" possibility. We can imagine *randomly* selecting a set of β consumers and removing their drug utilization and then assessing the WTP for that *randomly* forgone consumption. The total WTP for this random forgone consumption would then be:

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

Here, the WTP for the forgone consumption is just a share $\frac{\Delta q_d}{D_d(0)}$ of the total amount all consumers are willing to pay for the drug. For instance, since we estimate that prior authorization reduces the share of beneficiaries who use a drug by 26.9%, under this assumption we would measure that WTP of the drugs foregone due to prior authorization is just 26.9% of total WTP. After computing the total willingness to pay for each specific drug, we simply sum this measure up across drugs. In Appendix E.1, we describe in more detail the specific formulae of the demand curves and total WTP, as well as how we aggregate.[45]

We present the total amount consumers are willing to pay for the forgone consumption computed across all drugs, per beneficiary-year, in Table 10. We include a variety of alternative estimates. In all estimates, our willingness-to-pay measures are derived from estimates using whether a beneficiary ever fills the drug as the quantity outcome. Column (1) is based on linear demand curves in the method described above. As Taubinsky and Rees-Jones (2018) highlight, aggregating effects while ignoring heterogeneity in demand responses may underestimate total beneficiary value. Therefore, in Column (2), we present a version where we estimate therapeutic-class-level responses to the LIS transition and use them to construct class-specific demand curves, then aggregate up. Finally, in Column (3), we take an alternative approach where we assume that demand curves have constant semi-elasticity rather than constant slope, and estimate the semi-elasticity using Poisson regression. We describe these latter two methods in more detail in Appendix E.1. In each case, we obtain standard errors by bootstrapping.

The columns report the amounts consumers are willing to pay for the foregone consumption under the various demand curves. Our preferred estimate incorporates beneficiary fixed effects and linear demand. Under best-case screening, consumers are willing to pay $55.53 per beneficiary-year. Under random screening, consumers are willing to pay more: $206.60 per beneficiary-year.[46] Our class-specific estimates produce much smaller total WTP estimates—$18.69 in the best case and $69.54 in the random case—albeit with very little precision due to the noisy class-specific demand curve estimates. Our constant-semi-elasticity approach also produces smaller estimates, with $26.18 in the best case and $176.45 in the random case.

These estimates provide a benchmark by which we can evaluate the savings. In the preferred estimate of the perfect screening case (column 1), the amount consumers are willing to pay for the forgone drugs is around 64% of the net financial savings induced by prior authorization restrictions. In the random screening case, the amount consumers are willing to pay for the forgone consumption is a little over double those savings. Thus, as long as prior authorization screens beneficiaries well with respect to willingness-to-pay, losses in consumer willingness-to-pay will be below net 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

[45] Aggregating across drugs and classes presents a 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, which we cannot do. We therefore assume this case away, though the sign of the bias introduced is unknown.

[46] We also consider the possibility that screening is worse than random, as has been the empirical case in other settings (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 $\Delta CS_d^{\text{Worst-Case}} = -\int_0^{0.289 D_d(0)} D_d^{-1}(\theta) d\theta$, or $357.68 per beneficiary-year, around two times the net savings due to prior authorization.

between willingness-to-pay and value, then the results above imply some welfare gains from prior authorization in the perfect screening case and welfare losses in the random screening case. Willingness-to-pay may not reflect a beneficiary's true value of consuming a drug. Beneficiaries may, for instance, underestimate the benefit from a drug (Baicker et al. 2015). While we cannot estimate the extent of bias in our setting, we can bound how large the wedge between value and willingness-to-pay must be to overturn our above conclusions about Kaldor-Hicks efficiency. We assume that there is a scalar wedge ρ such that a beneficiary's willingness to pay is a share ρ of their true welfare-relevant value. This is effectively equivalent to making the welfare-relevant demand curve (in the sense of Bernheim and Rangel (2009)) steeper by a factor of $\frac{1}{\rho}$ for $\rho \in [0, 1]$. The value of ρ required to equate the consumer surplus reductions and the net financial benefits NFB of prior authorization is $\rho^* = \frac{\Delta CS^{\text{Revealed}}}{NFS}$, where $\Delta CS^{\text{Revealed}}$ is our pure revealed preference measure of consumer surplus. For the best-case-screening scenario, ρ^* is 0.64. That is, if we are in the best case of who is screened out, patients need to value their drugs at approximately 1.5 times their WTP for the consumer surplus effects of prior authorization to equal its effects on spending.

A complete welfare analysis also requires us to map from net *program* savings to net *social* savings. When prior authorization reduces utilization, it reduces program costs by reducing the amount of money spent by payers (the insurer and the government). However, its effects on *social* costs come through two channels: (1) the social marginal cost of public funds needed to finance the insurance program; and (2) the marginal cost of producing the forgone drugs. Our data does not allow us to evaluate (2) here, so we treat these costs as if they are zero and consider our calibration here an excessively pessimistic measure of effects on social welfare. For (1), we use typical calibrations of the marginal cost of financing a dollar of public spending ($0.30-$0.50), and multiply them by the change in spending induced by prior authorization. This results in *social* savings from reduced drug spending of $28-$48. When compared to administrative costs of $10, this results in net social financial savings of $18-$38. Under our primary estimates, there is no case in which prior authorization improves social welfare by this measure, though there are cases where this is true under our alternative estimates. To the extent that the drugs deterred by prior authorization are non-trivially costly, however, prior authorization may yet improve social welfare.

Another caveat is that the approach here focuses on patient incentives. There is substantial evidence that patients have limited control over treatment decisions, with control instead passed to their provider, their agent in medical decisions. While our model focuses on beneficiaries as consumers, one can imagine an alternative model where providers make decisions on behalf of their patients, trading off their altruistic intentions to maximize patient benefit against their own administrative costs. In Appendix E.2 we walk through such a model. Above, we showed that prior authorization is equivalent to charging patients $280 per prescription. In contrast, we know that the provider administrative costs of prior authorization are approximately $22 and induce the same response, implying that providers have much more elastic demand than beneficiaries. Since they are more elastic, the revealed surplus from the drug under their decision frame is much smaller. If we assumed that they equally weighted patient benefit and administrative costs, it would imply consumer surplus losses of $1.21 in the best-case screening case and $4.52 in the random screening case. Thus, in order to conclude that there are large welfare harms in a world where the response to prior authorization primarily reflects provider decision-making would require us to believe that providers put extremely low weight on patient value relative to their own administrative costs. Indeed, it makes sense

that prior authorization, a policy that heavily relies on imposing costs on physicians in their capacity as agents for patients, will be inefficient if physicians are poor agents.

Explicitly interpreting these estimated quantities as reflecting the welfare consequences of prior authorization requires strong assumptions. This is reflected in the wide range of estimates we obtain under different approaches. Ultimately, we interpret these results as suggestive evidence that the lost consumer surplus due to forgone drugs is of a similar order of magnitude relative to the cost of procurement. This motivates prior authorization restrictions as a *potentially* efficiency-enhancing rationing device.

5.2 Health Effects

Next, we investigate the effects of prior authorization on patient health. Our primary research design from Section 3 does not permit estimation of effects on health, as health is defined at the patient level, whereas our research design leverages variation in prior authorization restrictions at the patient-drug level. We thus modify our research design to accommodate person-level outcomes.

To do so, we construct a beneficiary-level measure of exposure to prior authorization aggregated across drugs. We follow in the spirit of Brot-Goldberg et al. (2023) and construct 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 the exposure of the beneficiary to prior authorization restrictions in a given plan by calculating the share of drugs that the beneficiary filled at least once in the prior year that would face prior authorization restrictions in that plan (as a share of covered drugs).

As in our earlier sections, since plan *enrollment* may be influenced by the beneficiary's demand for specific drugs, we use the 'fit' of the *assigned* default plan as an instrument for the 'fit' of the *enrolled* plan. However, here we face a second problem: some beneficiaries are more likely to face greater exposure to prior authorization than others simply because they take more (or different) drugs. Their drug regime is likely related to their underlying health, creating an additional spurious correlation between health outcomes and exposure (even when randomly assigned). We take the approach of Borusyak and Hull (2023) to addressing this problem. For each beneficiary, we compute their 'fit' in all benchmark plans in their choice set. We then use this to construct their ex-ante *expected* 'fit' measure given potential random assignment (i.e., the average fit across plans in the choice set). Borusyak and Hull (2023) show that subtracting this expected measure from our instrument removes the relevant omitted variable bias in this setting while retaining useful cross-sectional variation.

We therefore run a set of regressions of the form:

$$
\begin{aligned}
Y_{it} &= \beta_1 \text{AuthExposure}_{it}^{Enrolled} + \beta_2 \text{ExclExposure}_{it}^{Enrolled} + \delta_{m(it)} + \epsilon_{it} \\
\begin{bmatrix} \text{AuthExposure}_{idt}^{\text{Enrolled}} \\ \text{ExclExposure}_{idt}^{\text{Enrolled}} \end{bmatrix} &= \gamma_1 \text{AuthExposure}_{it}^{Assigned} + \gamma_2 \text{ExclExposure}_{it}^{Assigned} + \delta_{m(it)} + u_{it}
\end{aligned}
$$

where $\text{AuthExposure}_{it}^{Enrolled}$ and $\text{AuthExposure}_{it}^{Assigned}$ are the corrected measures of beneficiary i's ag-

gregate exposure to prior authorization due to assignment or enrollment in year t. β_1 thus (approximately) represents the health consequences of moving a beneficiary from facing no prior authorization restrictions on drugs they took in the prior year to facing restrictions on all previously-used drugs. We include ExclExposure_i to hold exposure to drug exclusion fixed.

We first show that this measure does indeed predict reductions in utilization. In the first column of Table 11, we examine the effect on total drug spending. In line with our results in Section 3.4, 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). Moving from being exposed to no prior authorization to having all previously-taken drugs restricted is estimated to *reduce* current-year mortality by 0.07 percentage points, a 3.1% decrease from baseline. However, this is noisily estimated; our standard errors are even larger, and we cannot rule out an 13.8% increase in mortality, nor a 20.1% 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 inpatient spending is increased by prior authorization, while all non-drug spending decreases. Again, however, the standard errors are large, and we cannot reject a wide range of effects.

Ultimately, we can conclude little about the consequences of prior authorization for beneficiary health. We cannot be overly definitive about the results from this exercise: We do *not* show 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. Our estimates include substantial negative effects on beneficiary health. If those represented the true effect of prior authorization, it would substantially influence our interpretation of the effects of prior authorization on social welfare (for the worse).

6 Conclusion

Our results suggest that prior authorization restrictions are a powerful tool for reducing healthcare costs. As highlighted by the American Medical Association and other interest groups, these restrictions 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 consequence of prior authorization is *not* wasteful spending on bureaucratic sludge associated with the authorization process, but, rather, the effect on utilization.

The welfare consequences remain ambiguous. The revealed preference approach we take suggests that, under some assumptions, willingness-to-pay for forgone drugs is less than the cost of acquiring those drugs. However, we would hesitate to strongly conclude that prior authorization is necessarily efficient. First, moving from willingness-to-pay to actual consumer benefits is difficult. Further, the "correct" social cost of the forgone drugs is not immediately obvious. Second, we are not able to conclusively estimate the effects on patient health, and our confidence intervals include positive effects on patient mortality that would likely

outweigh any financial savings. Third, we are unable to quantify any of the effects of prior authorization on *patient* administrative hassle, as opposed to provider hassle. Finally, we only consider partial equilibrium consequences and only evaluate welfare for patients and insurers. A broader consideration would require considering manufacturer welfare, as well as the general equilibrium effects on drug prices and insurance enrollment. Our results suggest a potential positive role for prior authorization on social welfare, although more research is needed to understand these missing pieces. But at a minimum, our results indicate that prior authorization restrictions are not unambiguously welfare-decreasing, as much of the rhetoric regarding these tools suggests.

Moreover, our results do *not* imply that prior authorization would be an effective policy if implemented more widely. Insurers in our setting appear to be, on average, restricting drugs that our model suggests are relatively more appropriate to restrict, 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 reduction in spending. Moreover, the current administration of these policies may not be optimal. 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, even when prior authorization raises social welfare, it does so by raising paperwork burdens for health care providers. These policies may be 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 perfect agents for these policies and not. In practice, however, there is *no* direct transfer, 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.

Second, our results speak to the choice of rationing mechanisms within the U.S. healthcare 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 $441 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 healthcare spending comparisons and U.S. healthcare 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

Abaluck, Jason, Jonathan Gruber, and Ashley Swanson, "Prescription Drug Use under Medicare Part D: A Linear Model of Nonlinear Budget Sets," *Journal of Public Economics*, 2018, *164*, 106–138. 28

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 College of Physicians, "ACP Urges Senate to Pass Legislation Aimed at Improving the Medicare Advantage Prior Authorization Process," 2022. https://www.acponline.org/advocacy/acp-advocate/archive/october-7-2022/acp-urges-senate-to-pass-legislation-aimed-at-improving-the-medicare-advantage-p 21

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. 6

Archer, Diane, "Medicare Is More Efficient Than Private Insurance," *Health Affairs Blog*, 2011. https://www.healthaffairs.org/do/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, 5, 8, 27, 31

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

Bernheim, B. Douglas and Antonio Rangel, "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," *Quarterly Journal of Economics*, 2009, *124* (1), 51–104. 31

Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 1994, pp. 242–262. 22, A22, A23

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. A15

Borusyak, Kirill and Peter Hull, "Nonrandom Exposure to Exogenous Shocks," *Econometrica*, 2023, *91* (6), 2155–2185. 32

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. 5, 21, 27, 28

_ **and Mathijs de Vaan**, "Intermediation and Vertical Integration in the Market for Surgeons," 2018. 22

_ , **Catherine Che, and Benjamin R. Handel**, "Pharmacy Benefit Managers and Vertical Relationships in Drug Supply: State of Current Research," 2022. NBER Working Paper No. 29959. 6

Brot-Goldberg, Zarek, Timothy Layton, Boris Vabson, and Adelina Yanyue Wang, "The Behavioral Foundations of Default Effects: Theory and Evidence from Medicare Part D," *American Economic Review*, 2023, *113* (10), 2718–58. 2, 10, 19, 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. 25, A14

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. 25, A14, A15

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, 5, 27

_ , **Jonathan Gruber, and Robin McKnight**, "Patient Cost-Sharing and Hospitalization Offsets in the Elderly," *American Economic Review*, 2010, *100* (1), 193–213. 33

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. 23, A27

Council for Affordable Quality Healthcare, "2013 U.S. Healthcare Efficiency Index: Electronic Administrative Transaction Adoption and Savings," 2014. 25, A14

CoverMyMeds, "2020 Electronic Prior Authorization Report," 2020. https://www.covermymeds.com/main/medication-access-report/electronic-prior-authorization/. Retrieved January 9, 2023. 21

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/do/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

Dalton, Christina M., Gautam Gowrisankaran, and Robert J. Town, "Salience, Myopia, and Complex Dynamic Incentives: Evidence from Medicare Part D," *Review of Economic Studies*, 2020, *87* (2), 822–869. 28

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. 25, A15

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, 30

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 Hale Shapiro, Daniel J. Sonnenstuhl, and Pietro Tebaldi, "A Denial a Day Keeps the Doctor Away*," *Quarterly Journal of Economics*, 06 2023, p. qjad035. 4, 5, 24

_ , _ , and _ , "Administration Above Administrators: The Changing Technology of Health Care Management," *AEA Papers and Proceedings*, 2020, *110*, 274–278. 1, 4

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

Gandhi, Ashvin and Maggie Shi, "The Limits of Spending Limits: Evidence from Medicare Physical Therapy," 2024. 4, 29

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. 6

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. 18

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*, 2022, *4* (2), 175–90. 5, 27, 28

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. 23, A22

Hendren, Nathaniel, "Measuring Ex Ante Welfare in Insurance Markets," *Review of Economic Studies*, 2020, *88* (3). 7

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

Ho, Kate and Robin S. Lee, "Contracting over Rebates: Formulary Design and Pharmaceutical Spending," 2023. 6

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. A15

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. A15

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

Newhouse, Joseph P. and the Insurance Experiment Group, *Free for All?: Lessons from the Rand Health Insurance Experiment*, Harvard University Press, 1993. 21

Nichols, Albert L. and Richard J. Zeckhauser, "Targeting Transfers through Restrictions on Recipients," *American Economic Review*, 1982, *72* (2), 372–377. 1, 6, 8

Pauly, Mark V., "The Economics of Moral Hazard: Comment," *American Economic Review*, 1968, *58* (3), 531–537. 5

Petrin, Amil and Kenneth Train, "A Control Function Approach to Endogeneity in Consumer Choice Models," *Journal of Marketing Research*, 2010, *47* (1), 3–13. 23, A25

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. 25, A14, A15

Sarig, Oren, "Pharmaceutical Demand Response to Utilization Management," *Journal of Health Economics*, 2024, *93*, 102830. 4, A14

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. A15

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," *Forum for Health Economics and Policy*, 2014, *17* (2), 153–168. 4

Shao, Jun and Dongsheng Tu, *The Jackknife and Bootstrap*, Springer Science + Business Media, 1995. A27

Shepard, Mark and Myles Wagner, "Do Ordeals Work for Selection Markets? Evidence from Health Insurance Auto-Enrollment," 2023. NBER Working Paper No. 30781. 4

Shi, Maggie, "Monitoring for Waste: Evidence from Medicare Audits," *Quarterly Journal of Economics*, 2024, *139* (2), 993–1049. 4, 5, 24

Taubinsky, Dmitry and Alex Rees-Jones, "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment," *Review of Economic Studies*, 2018, *85* (4), 2462–2496. 30

Train, Kenneth, *Discrete Choice Methods with Simulation*, Cambridge University Press, 2009. A23

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. A15

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. A15

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. 5

7 Figures

Figure 1: Use of Prior Authorization in Our Sample

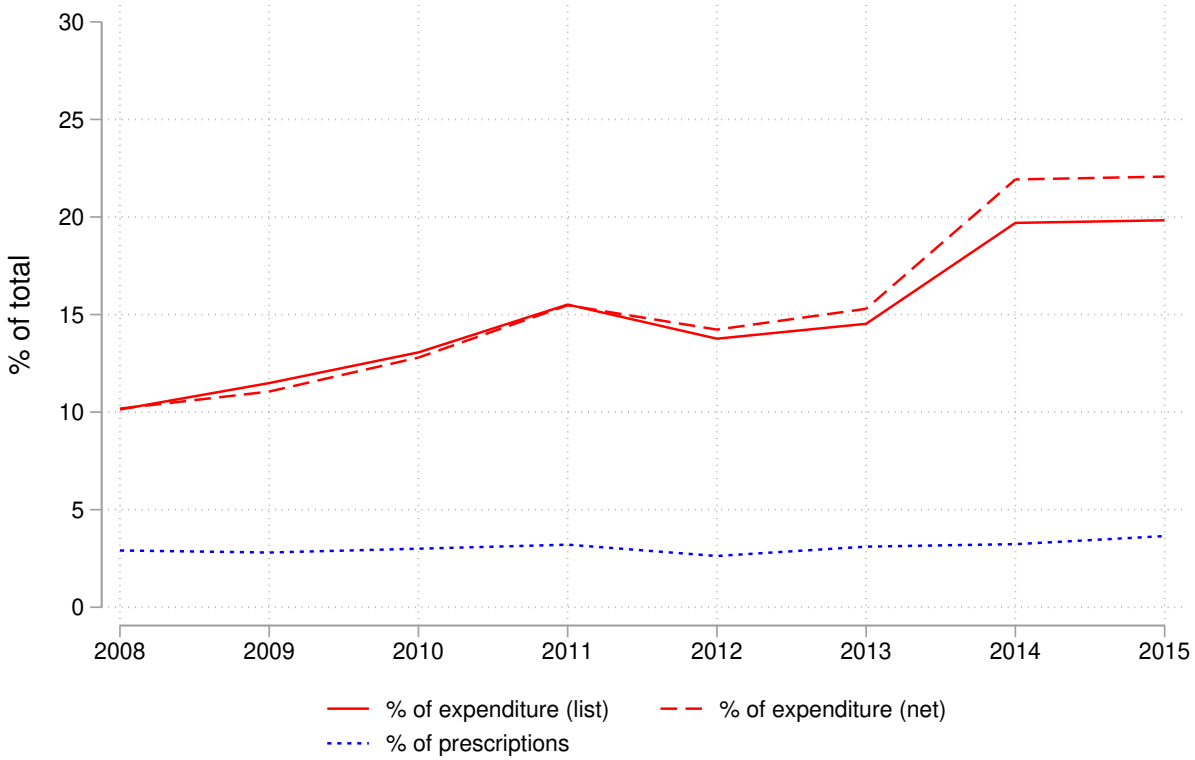

Notes: This figure plots a time series of the share of prescriptions filled among beneficiaries in our sample that required prior authorization. The blue dotted line plots the share of all filled prescriptions requiring prior authorization. The solid red line weights those prescriptions by their list 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

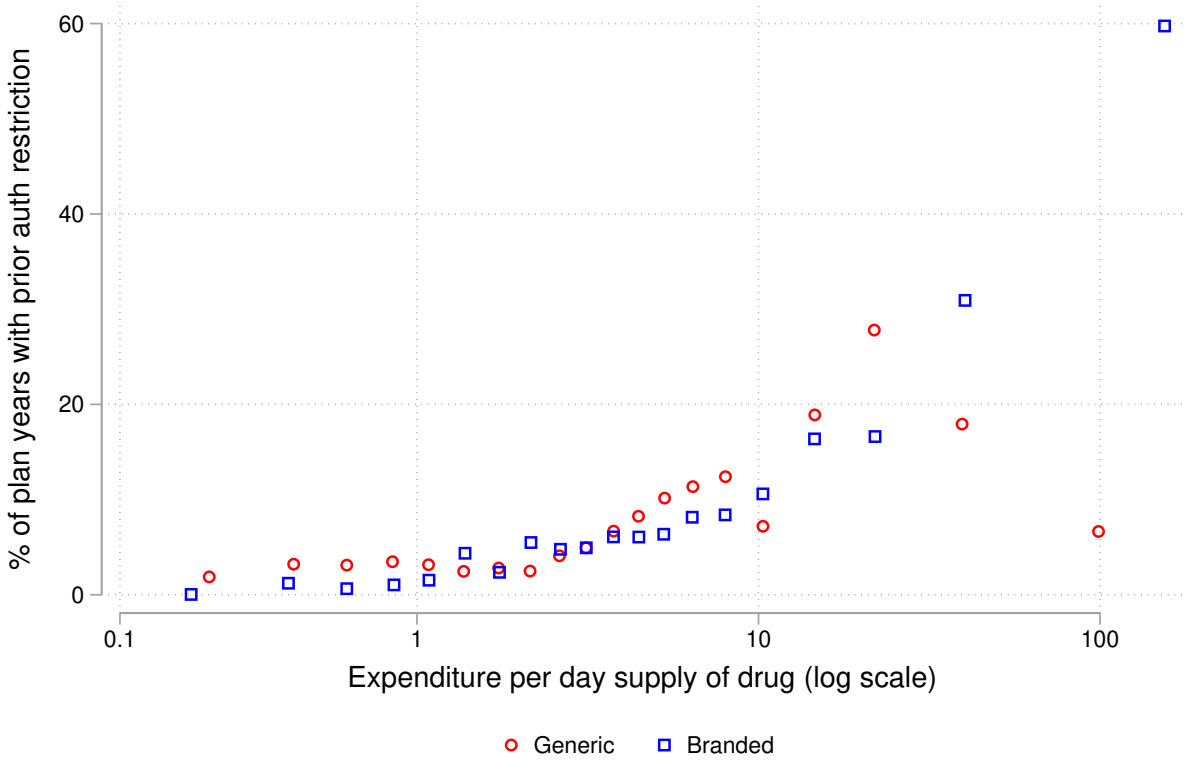

Notes: This figure shows the relationship between average expenditure (net price) per day supply of a drug and the share of plans that put prior authorization restrictions on that drug. Each observation is a drug-year pair. Drugs with fewer than 20 prescriptions in a year within our sample are excluded. 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

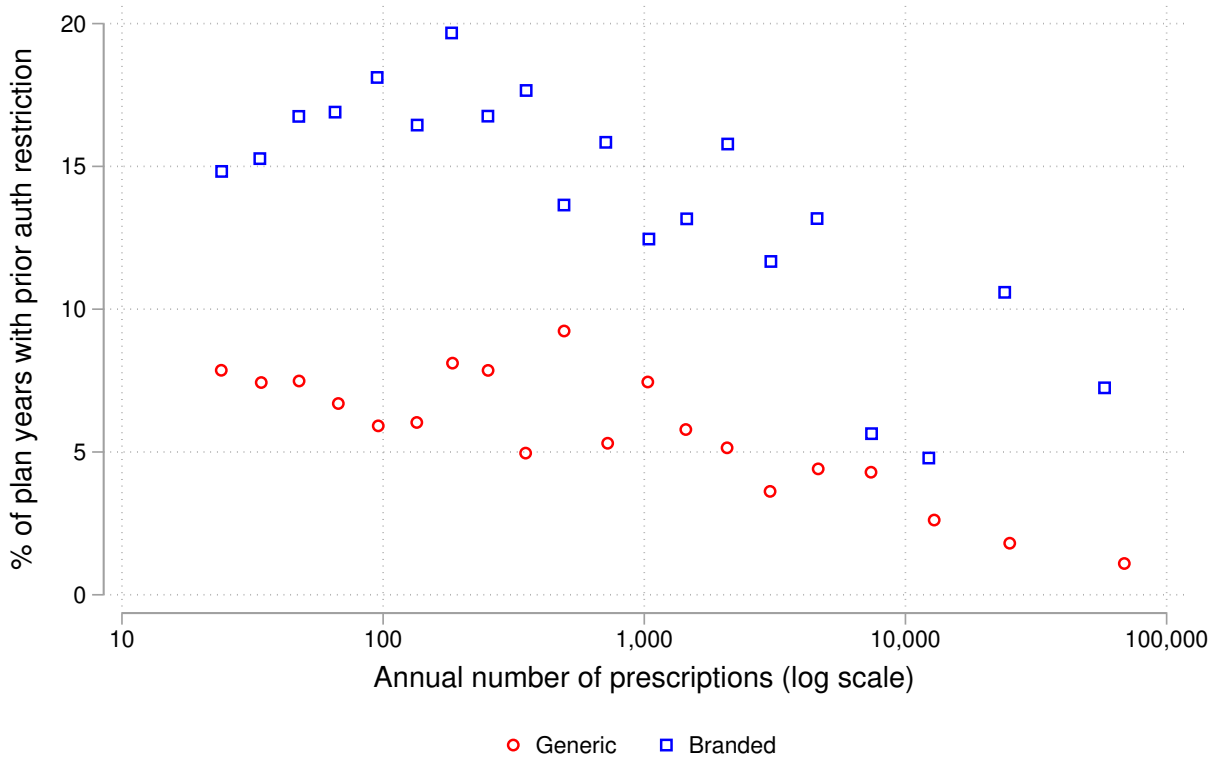

Notes: This figure displays the relationship between the number of users of a drug in a given year and the share of plans that put prior authorization restrictions on that drug. Each underlying observation is a single drug-year pair. Drugs with fewer than 20 prescriptions in a year are excluded.

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

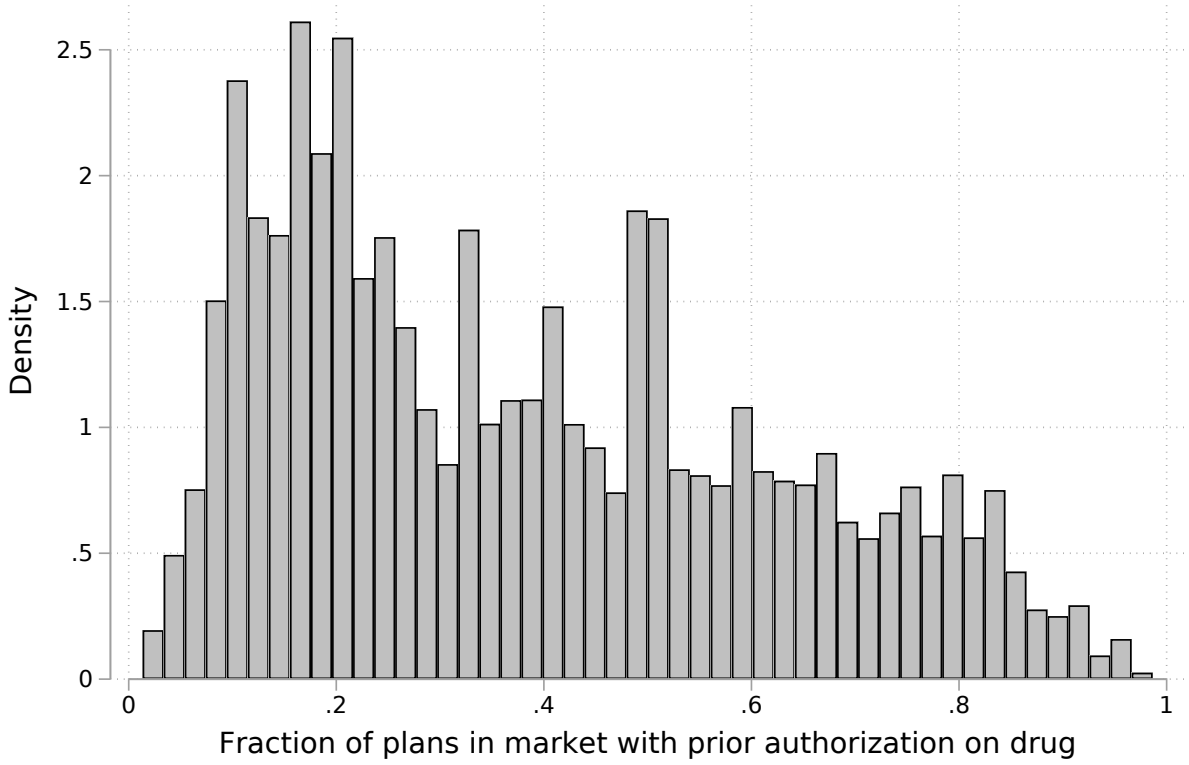

Notes: This figure displays the distribution of the fraction of plans within a service region that require prior authorization for the drug in a given year, weighted by number of enrollees in the plan. Each underlying observation is a single drug-region-year pair, $N = 75,875$. Market-years where no 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: Drug-Level Effects of Prior Authorization

(a) Effects on Use of Focal Drug

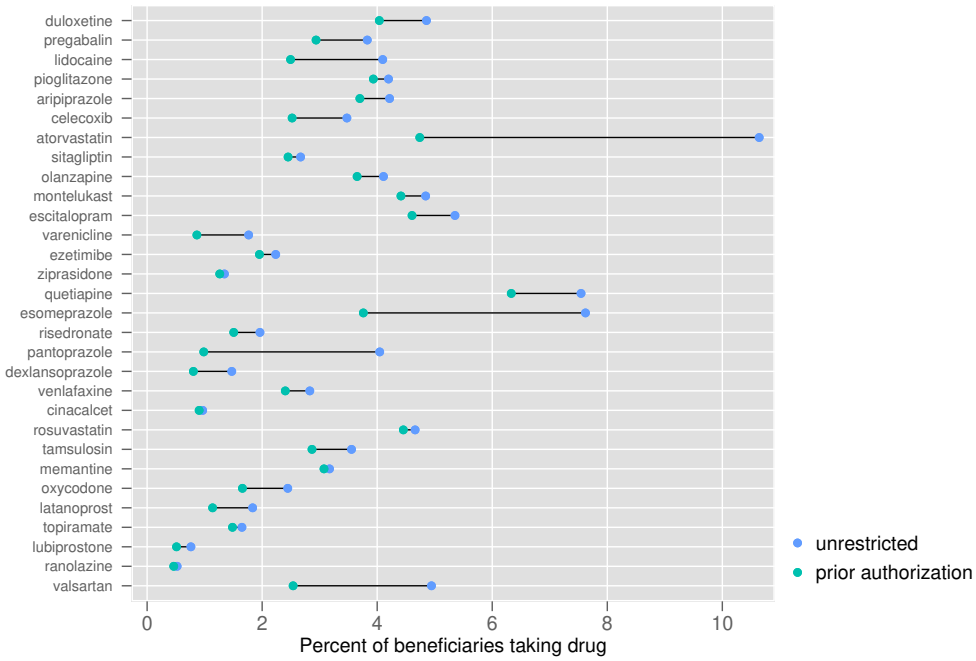

(b) Effects on Use of Other Drugs in Class

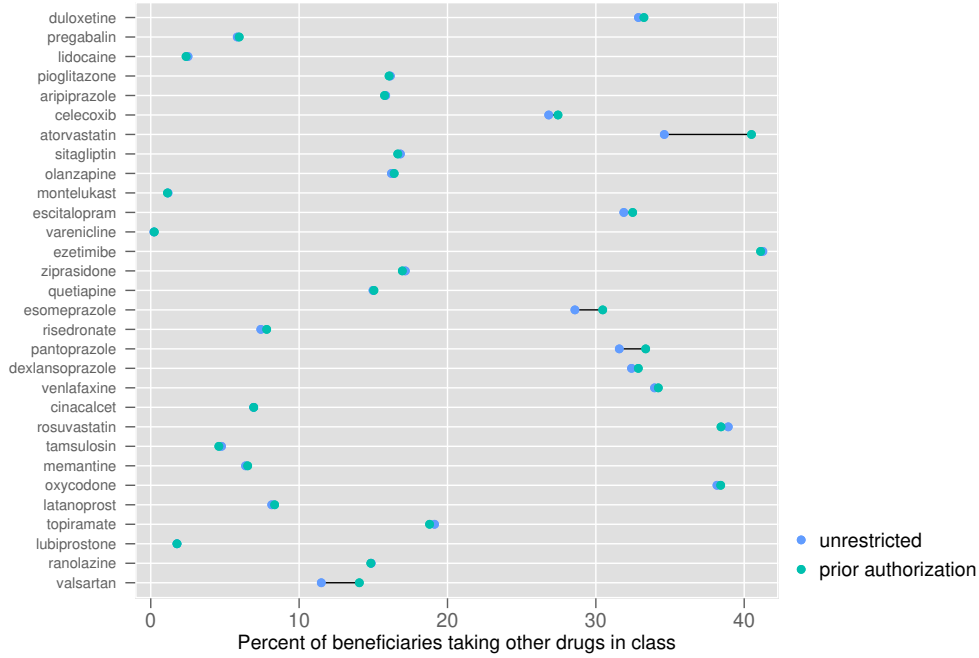

Notes: Panel (a) shows the percent of beneficiaries using a drug in plans where that drug is subject to prior authorization and plans where it is unrestricted, controlling for market-year. Panel (b) shows the percent of beneficiaries taking any other drug in the same therapeutic class as a given drug in plans where that drug is subject to prior authorization and plans where it is unrestricted, controlling for market-year. The set of drugs plotted are the branded drugs that have the highest min(% of plans with PA, % unrestricted) $\times$ number of users, summed across all market-years.

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

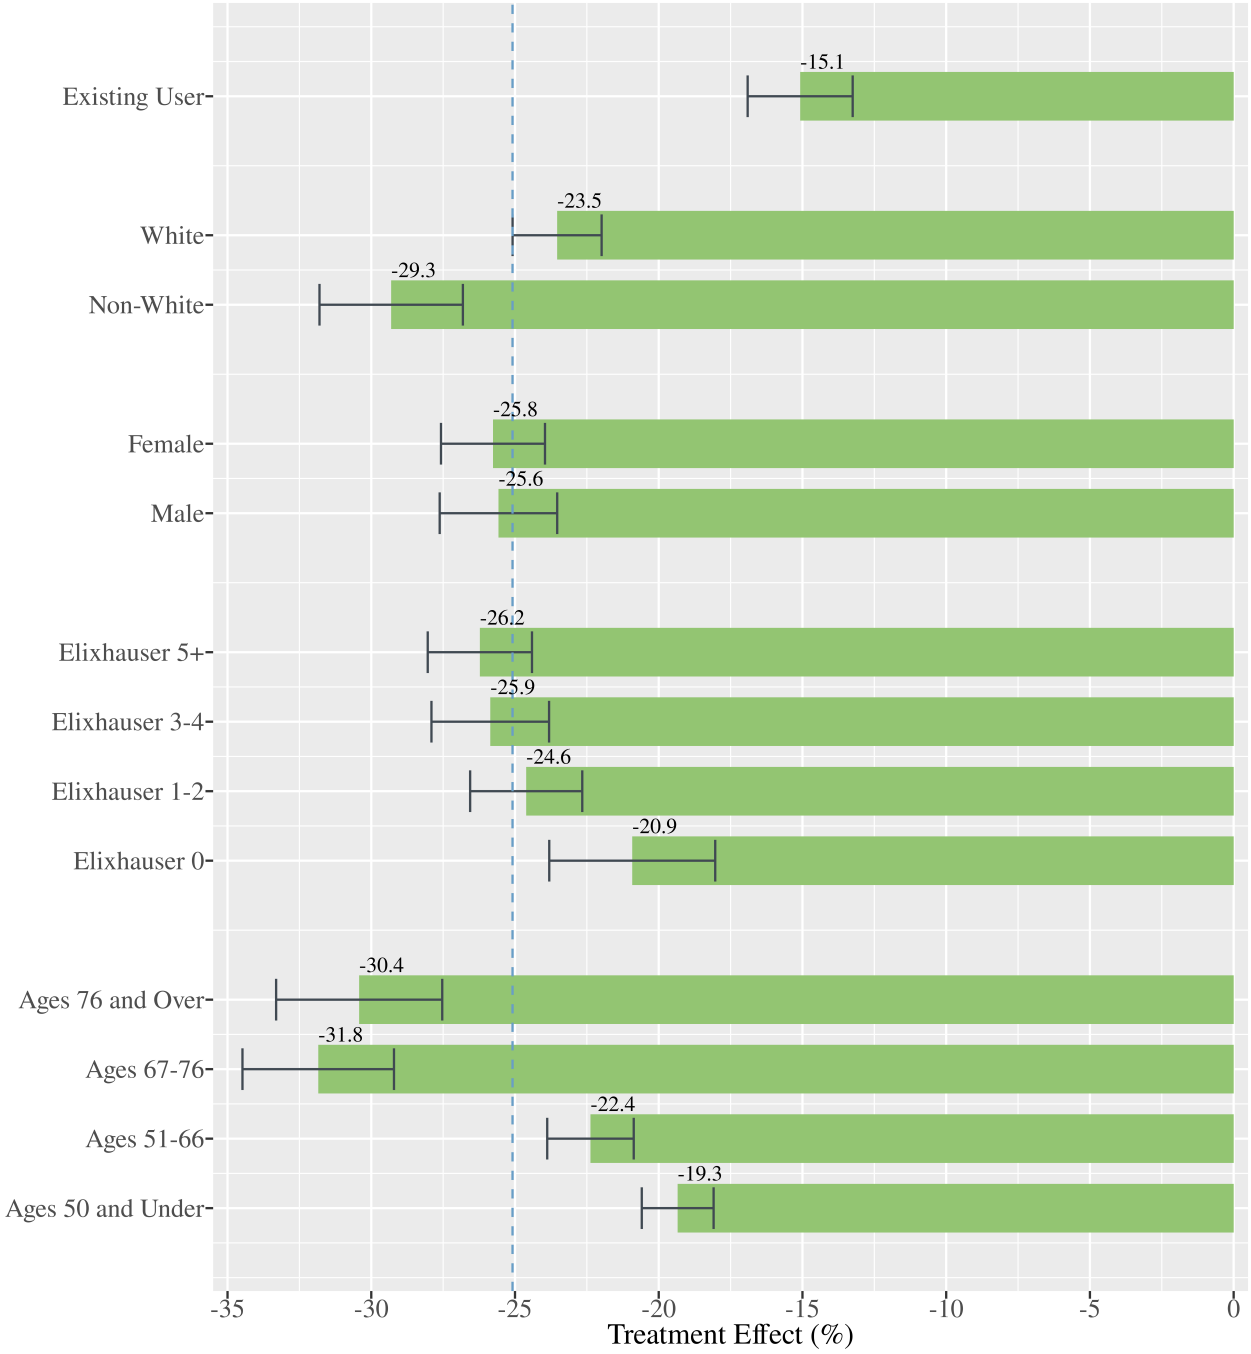

Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug, for subsamples of beneficiaries. Effects are presented in terms of the percent change due to prior authorization relative to a control mean, reweighted as described in Section 3.2.

Figure 7: Heterogeneous Effects of Prior Authorization on Utilization of Other Drugs by Beneficiary Characteristics

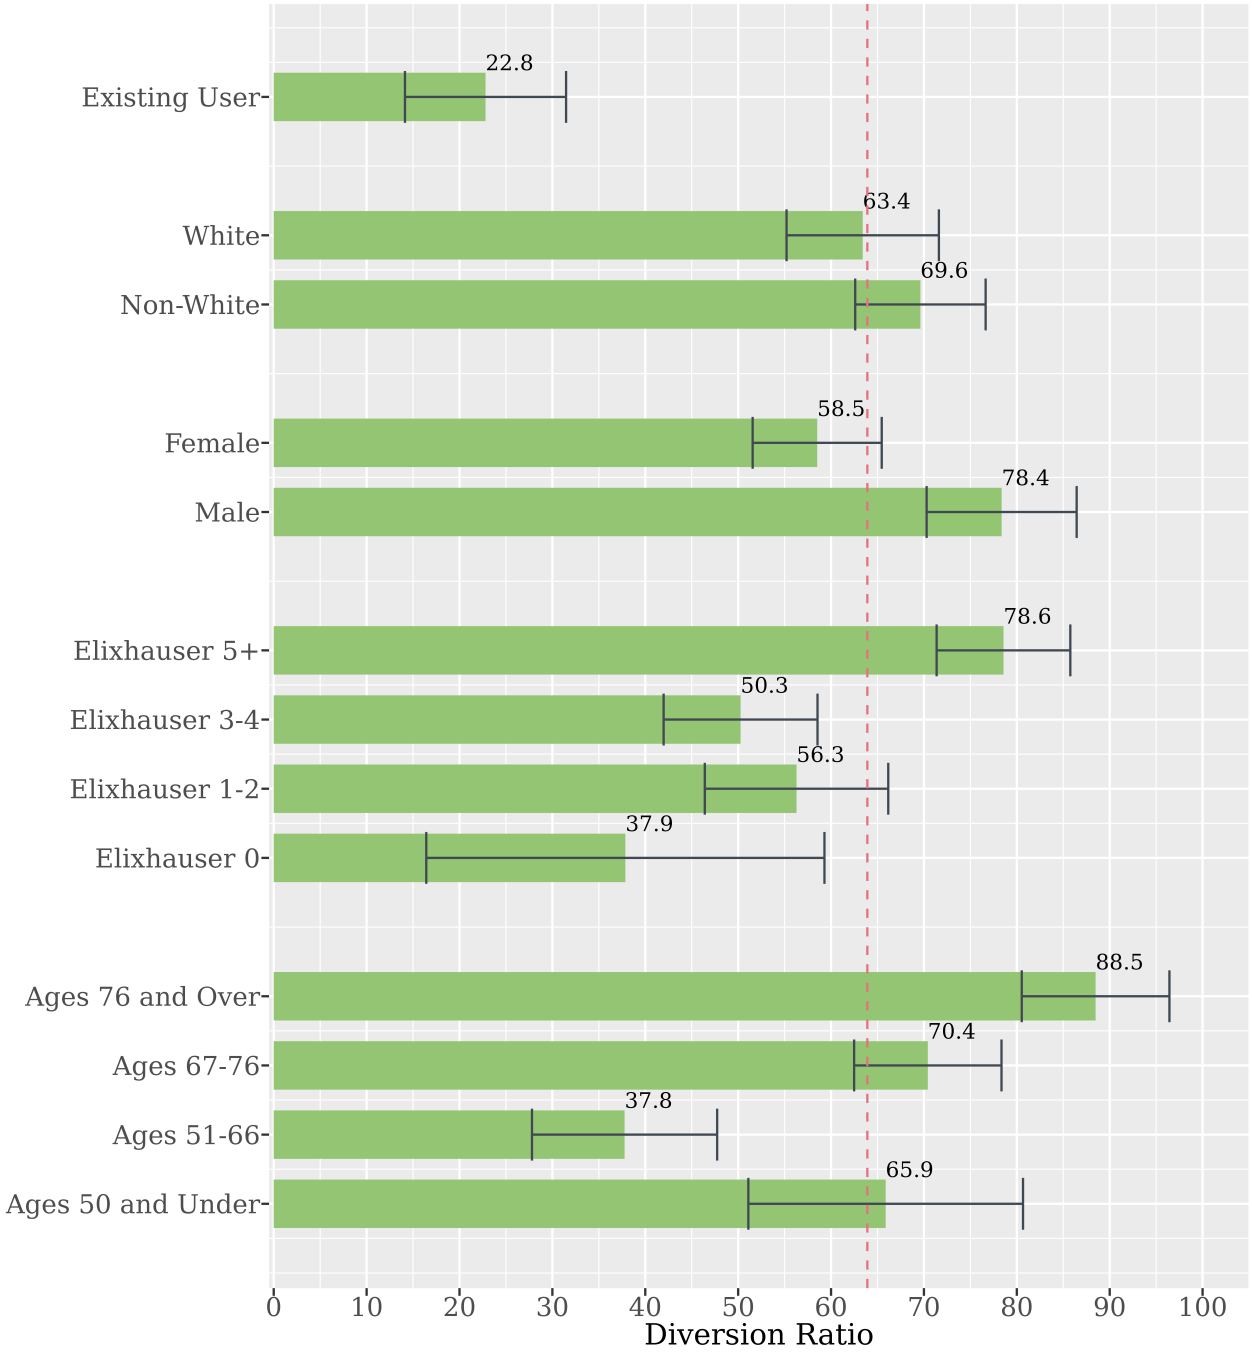

Notes: This figure presents estimates of diversion to other drugs due to prior authorization, estimated separately for subsamples of beneficiaries. The diversion ratio plotted here is the ratio of the effect of prior authorization on the use of of other drug in the same class as the focal drug, divided by its effect on the focal drug.

Figure 8: Heterogeneous Effects of Prior Authorization on Utilization by Drug Characteristics

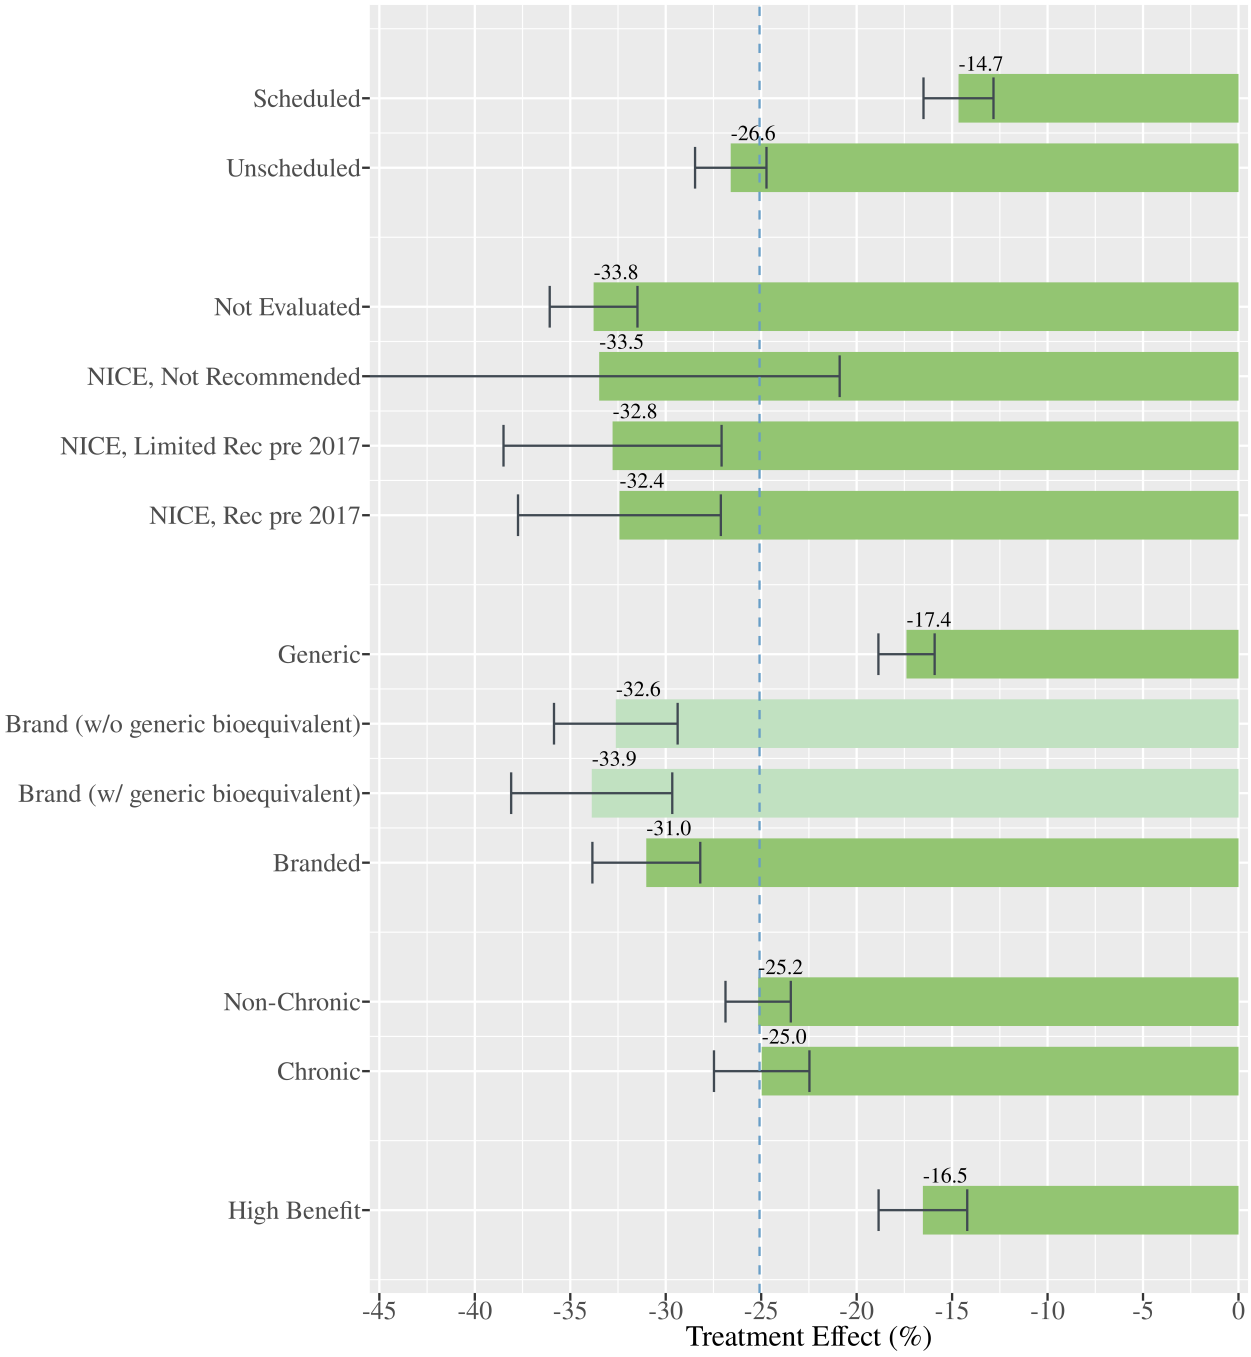

Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug, for subsamples of drugs. Effects are presented in the percent change due to prior authorization relative to a control mean, reweighted as described in Section 3.2.

Figure 9: Ratio of Drug Cost Reduction to Administrative Cost Burden by Class

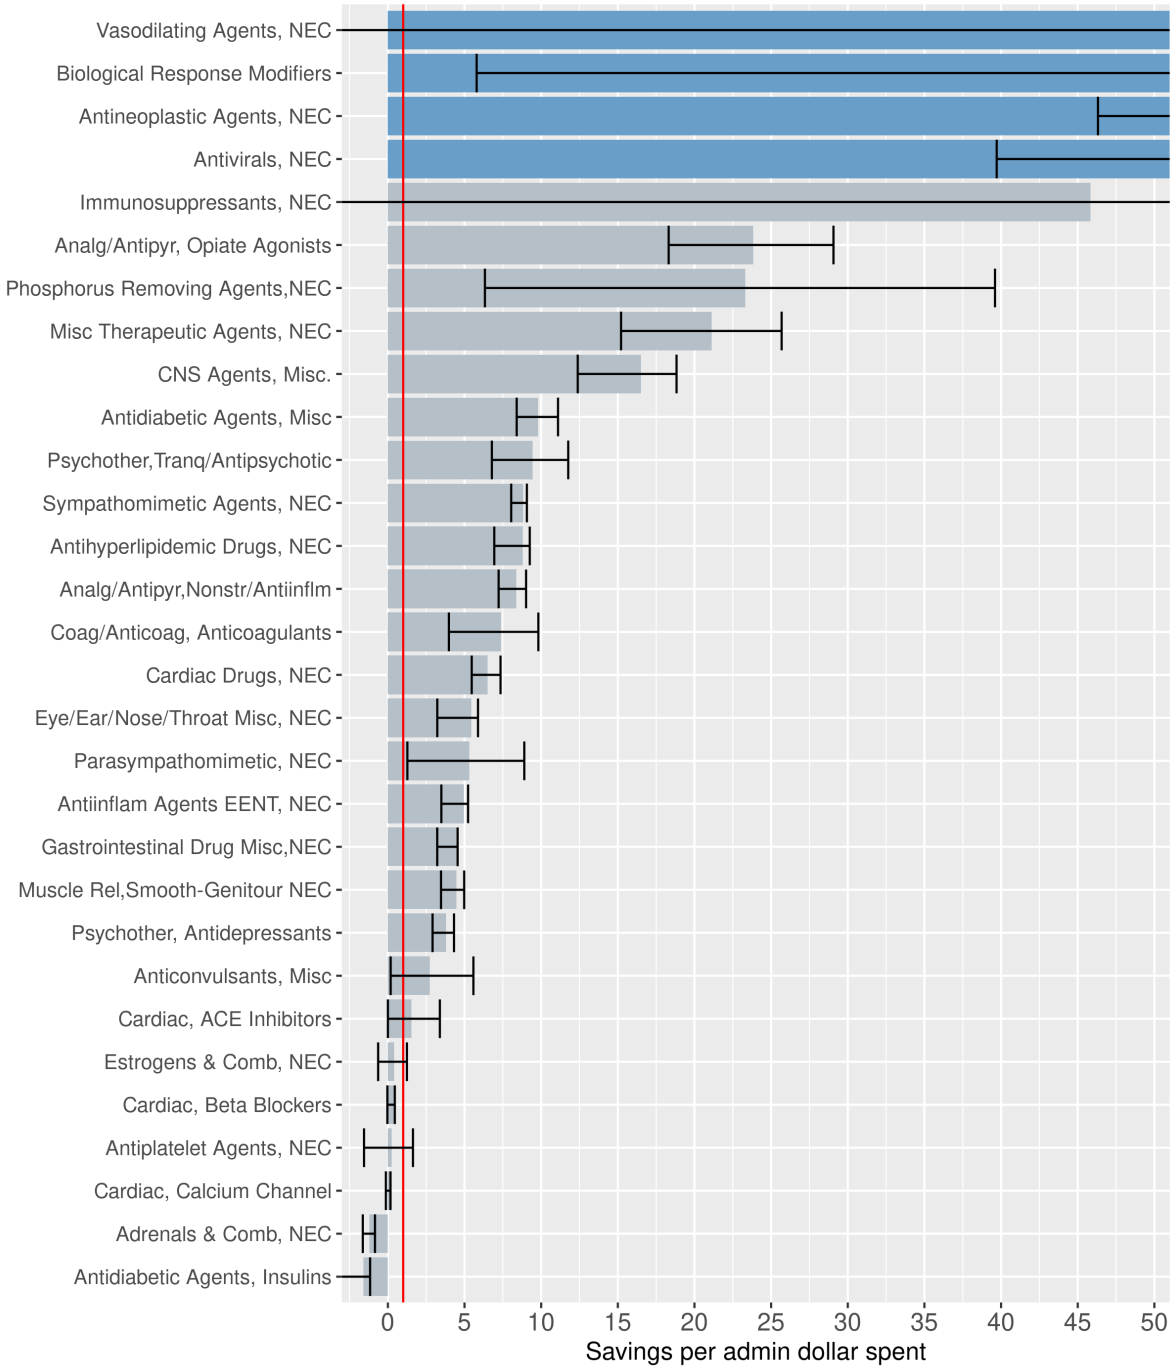

Notes: This figure reports, for each therapeutic class, estimates of the amount of spending reduced due to status quo prior authorization policies per dollar of administrative costs induced, under the calibration $a = \$22.48$ and $r = 0.04$. This is reported for the top 30 therapeutic classes by total 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.

8 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: Branded without generic bioequivalent	Drug type: Branded with generic bioequivalent	Drug type: Generic
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	
	Restrictions in Enrolled Plan	Exclusion in Enrolled Plan	Restrictions in Enrolled Plan	Exclusion in Enrolled Plan
Restrictions on Focal Drug in Assigned Plan	0.908	-0.001	0.873	0.000
	(0.002)	(0.000)	(0.005)	(0.001)
Exclusion of Focal Drug in Assigned Plan	0.000	0.905	0.002	0.849
	(0.000)	(0.003)	(0.001)	(0.006)
F-statistic	76,610	59,489	18,305	9,656
Number of drug-beneficiary-years	1,719,692,195		10,197,530	
Number of beneficiary-years	1,110,968		998,443	
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

	% Ever Filled	Days Supply	Spending
Restrictions on Focal Drug	-0.108	-0.170	-0.612
	(0.004)	(0.007)	(0.029)
Restrictions on Substitute Drugs	0.053	0.068	0.100
	(0.004)	(0.005)	(0.023)
PA % Effect	-26.9	-33.8	-21.9
Control Mean	1.305	1.527	2.811
Reweighted Control Mean	0.403	0.504	2.794
Number of drug $\times$ beneficiary-years		1,719,692,195	
Number of market years		210	
Average plans per market-year		6.6	
Average beneficiaries per plan		803	
Average drugs per year		1569.2	

Notes: This table presents coefficient estimates from instrumental variable regressions estimating the effect of whether a drug was put under prior authorization restrictions in the plan that each beneficiary was enrolled in a given year on that beneficiary's utilization of that drug 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 drug-market-year fixed effects. Prior authorization of substitute drugs is given by the average 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.

Table 6: Estimates of the Effect of Prior Authorization Status on Utilization of Substitute Drugs

All Drugs In Class Except Focal Drug			
	% Ever Filled	Days Supply	Spending
Restrictions on Focal Drug	0.021	0.109	0.083
	(0.007)	(0.023)	(0.455)
Restrictions on Substitute Drugs	-0.252	-0.814	-8.596
	(0.030)	(0.113)	(1.301)
Diversion to other drug	56.0	63.9	13.6

All Drugs In Class			
	% Ever Filled	Days Supply	Spending
Restrictions on Focal Drug	-0.048	-0.062	-0.529
	(0.007)	(0.022)	(0.456)
Restrictions on Substitute Drugs	-0.229	-0.746	-8.496
	(0.030)	(0.114)	(1.315)
PA % Effect	-0.7	-0.4	-1.0

Notes: This table presents coefficient estimates from instrumental variable regressions estimating the effect of whether a drug was put under prior authorization restrictions in the plan that each beneficiary was enrolled in a given year on that beneficiary's utilization of other drugs in the same therapeutic class as that drug 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 drug-market-year fixed effects. The first panel reflects regressions where the outcomes are the use of drugs in the therapeutic class *other than* the focal drug. The second panel reflects regressions where the outcomes are the use of drugs in the therapeutic class *including* the focal drug. Prior authorization of substitute drugs is given by the average 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.

Table 7: 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	(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 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. Parenthetical terms denote bootstrap standard errors for their associated estimate.

Table 8: 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
		(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 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. Parenthetical terms denote bootstrap standard errors for their associated estimate.

Table 9: Effect of Transitioning into the LIS Program on Out-of-Pocket Prices and Utilization

	Out Of Pocket Payment	% Ever Filled	Days Supply	Spending
Post-transition into	-192.28	0.047	0.089	0.570
Low Income Subsidy	(0.43)	(0.005)	(0.008)	(0.203)
Reweighted Mean Pre-transition	181.03	0.255	0.375	1.499
Reweighted Mean Post-transition	9.71	0.279	0.415	2.405
Number of drug $\times$ beneficiary years		34,078,924		
Number of market years		205		
Average drugs per year		1,054		

Notes: This table reports coefficient estimates from regressions of out-of-pocket payment, % beneficiaries who ever filled a prescription from the drug, days supply and net spending for a drug on a dummy equal to 1 if the beneficiary has transitioned into the Low Income Subsidy program. These regressions use our sample who transition into the LIS program and leverage the transition as a shock to out-of-pocket prices. All regression specifications include beneficiary fixed effects and drug-year fixed effects. Observations are weighted $w_{idt} = p_{dm(it)}(1 - p_{dm(it)})$, where $p_{dm(it)}$ is the probability. Weights are calculated using the main sample for the prior authorization analysis. Standard errors are clustered at the beneficiary level.

Table 10: Revealed Preference Estimates of Consumer Surplus Loss

		(1) Linear	(2) Class-Specific Linear	(3) Constant Semi-Elasticity
Willingness to Pay	Best-Case	55.53	18.69	26.18
		(45.79, 69.68)	(-127.03, 161.43)	(21.36, 34.24)
	Random	206.60	69.54	176.45
		(170.39, 259.27)	(-472.66, 600.67)	(144.00, 230.83)
Net Financial Savings from Prior Authorization			86.12	

Notes: This table provides estimates of the loss in consumer surplus, in dollars per beneficiary-year, due to the present 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 9, as well as from estimates of how prior authorization changes drug use, given in Table 5. The two columns represent the consumer surplus measures derived from the two columns in Table 9, 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 7 and our preferred estimate of the average administrative cost of prior authorization given in bold in Table 8.

Table 11: Effects of Aggregate Prior Authorization Restriction Exposure on Utilization and Health Outcomes

All Beneficiaries				
	Spending	% Died in Year	Inpatient Spending	Non-Drug Medical Spending
Exposure to Prior	-591.54	-0.07	47.01	-236.96
Authorization	(93.55)	(0.19)	(262.10)	(361.30)
Control Mean	4,519.39	2.24	6,095.87	12,203.32
Beneficiary-Years			604,451	

Notes: This table presents coefficient estimates from a set of regressions of a beneficiary's utilization and health outcomes in a given year on their exposure to prior 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

(a) Weighted by Number of Beneficiaries

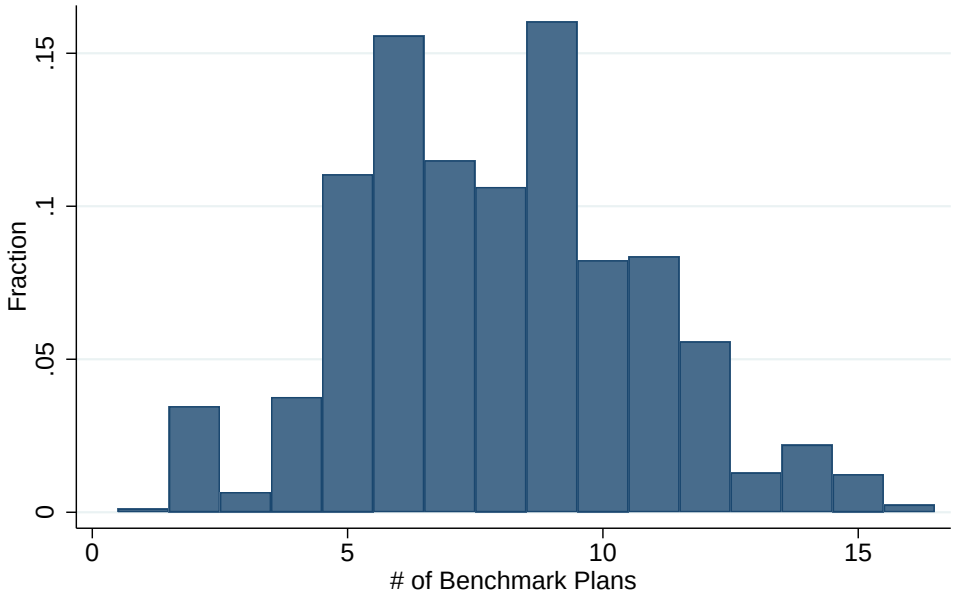

(b) Unweighted

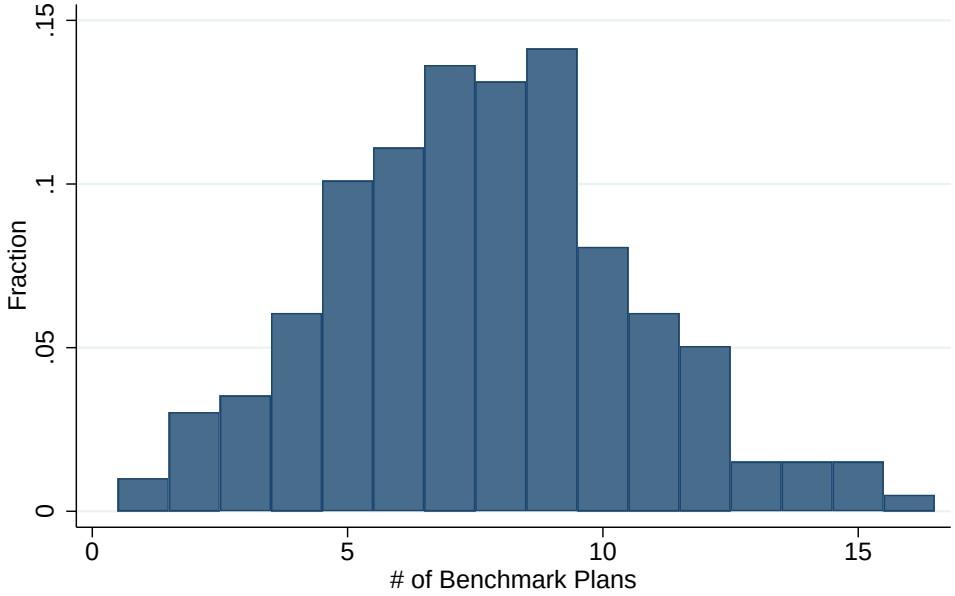

Notes: This set of figures plots the distribution in the number of benchmark plans across the pairs of Part D service region-years. The top figure presents this distribution weighting all Part D service region-year pairs equally, while the bottom weights Part D service region-year pairs by the number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

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

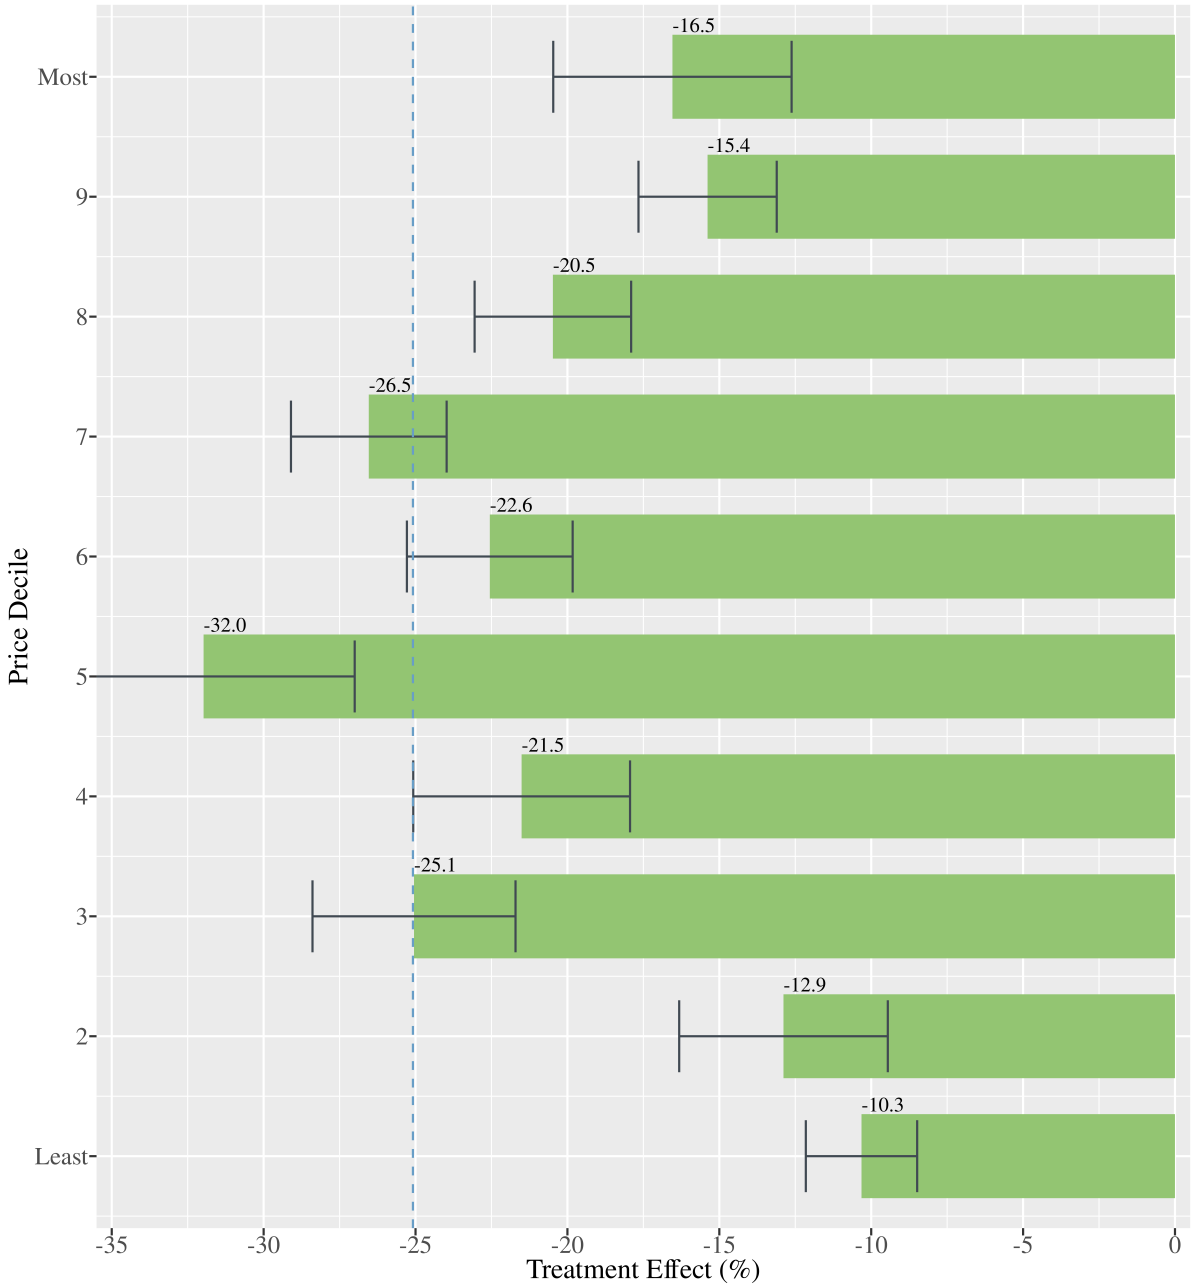

Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug. We run separate regressions on groups of drug-year pairs, where pairs are grouped into decile 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 3.2.

Appendix Figure A3: 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 3.2.

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

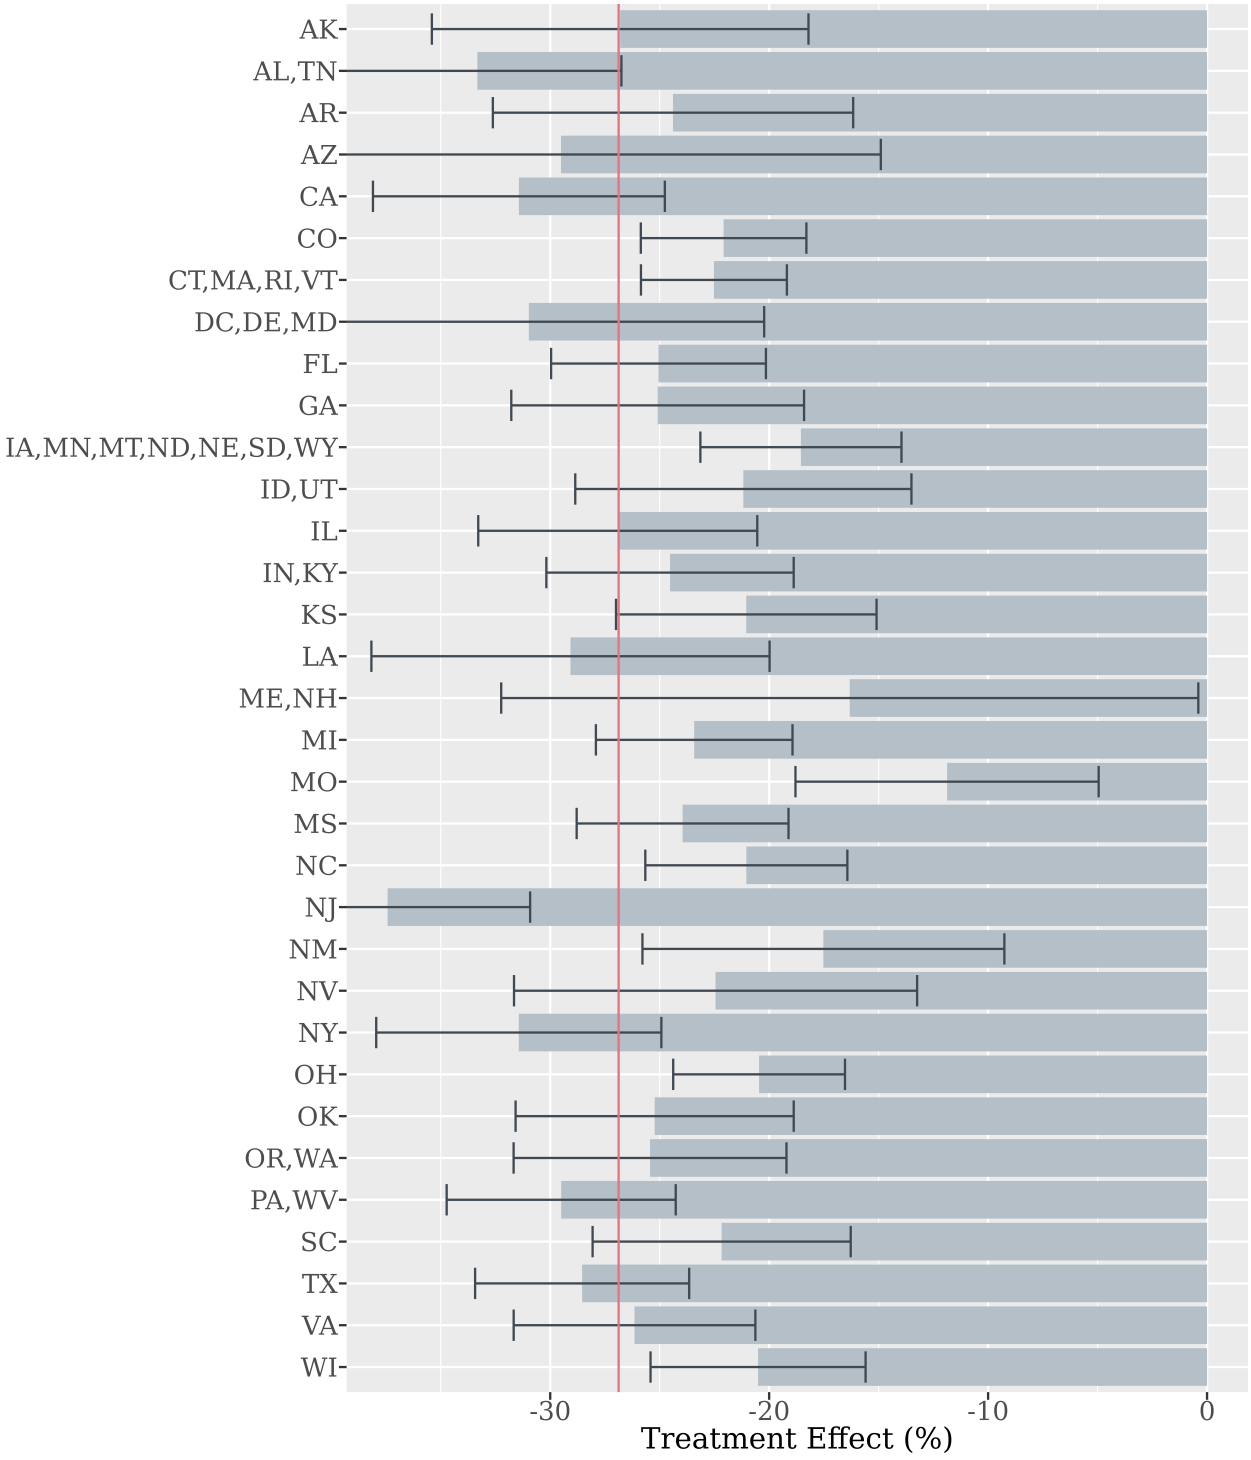

Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug. We run separate regressions for each geographically-defined service region in which plans compete. Effects are presented in the percent change due to prior authorization relative to a control mean, reweighted as described in 3.2.

Appendix Figure A5: Heterogeneous Effects of Prior Authorization on Utilization by Insurer

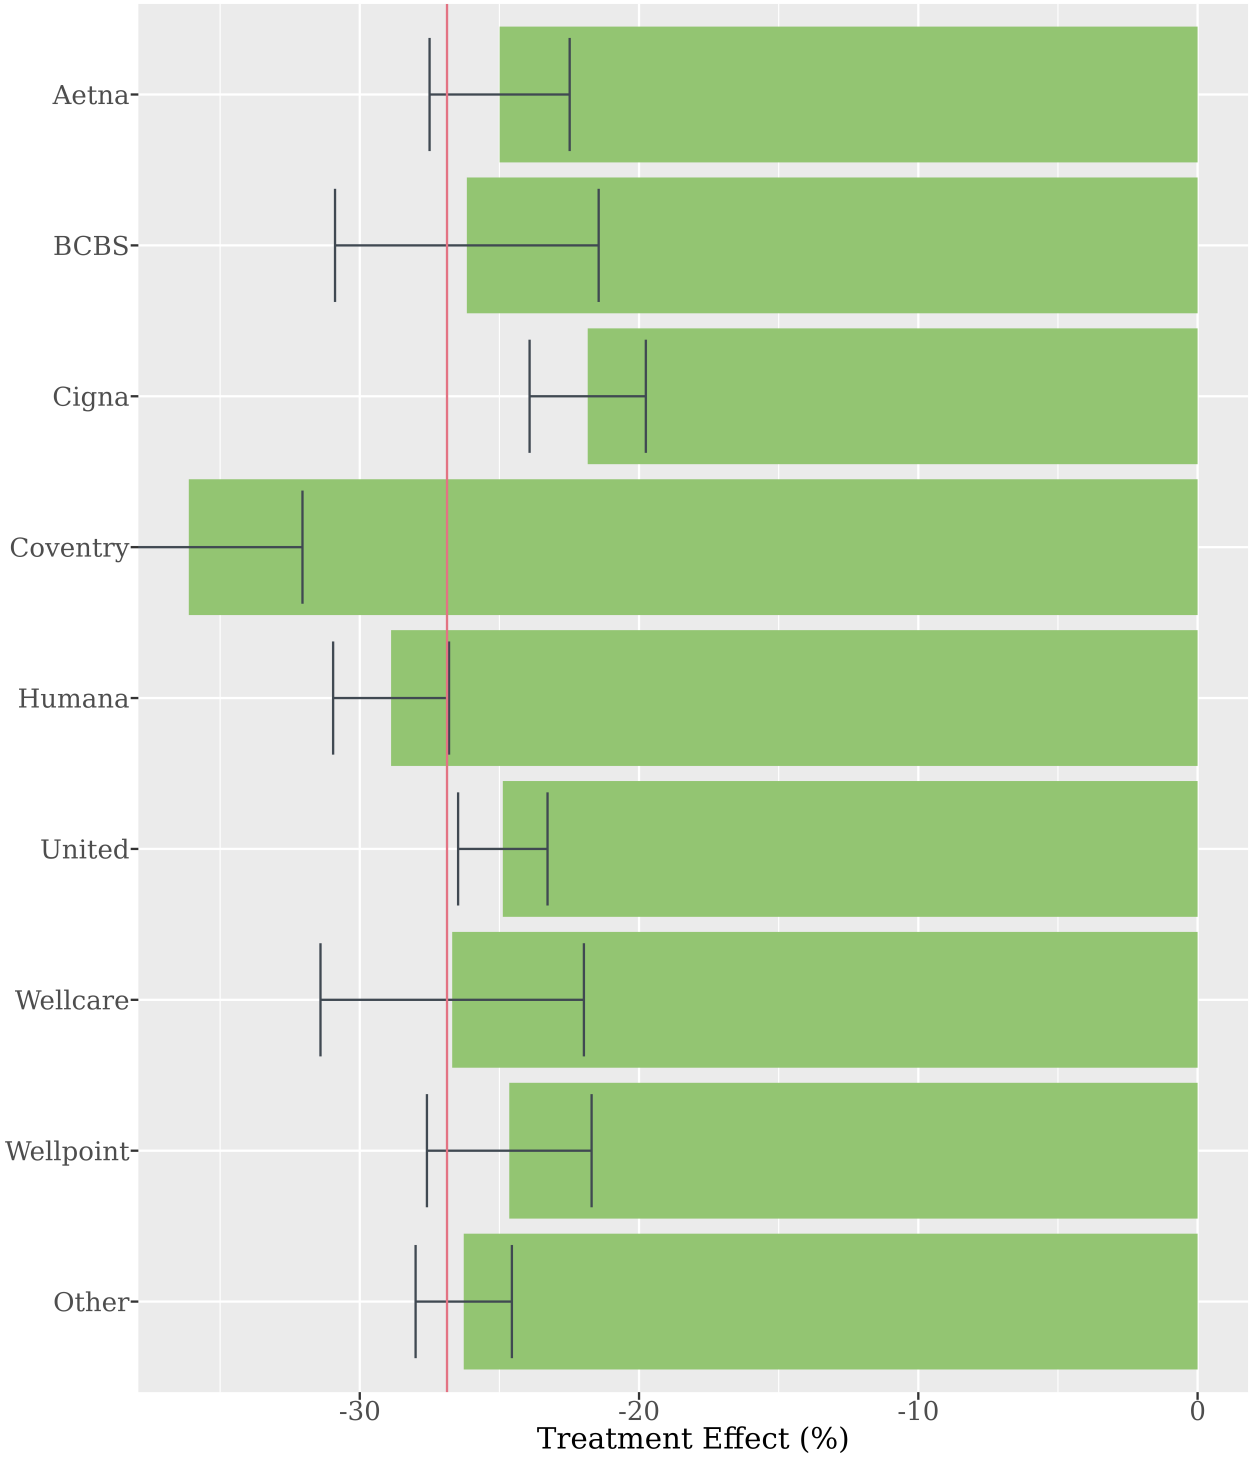

Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug. We interact the effect of prior authorization with the identity of the parent insurance carrier that sponsors the plans. Effects are presented in the percent change due to prior authorization relative to a control mean, reweighted as described in 3.2.

Appendix Figure A6: Heterogeneous Effects of Prior Authorization on Utilization by Drug-Year Restriction Rate Deciles

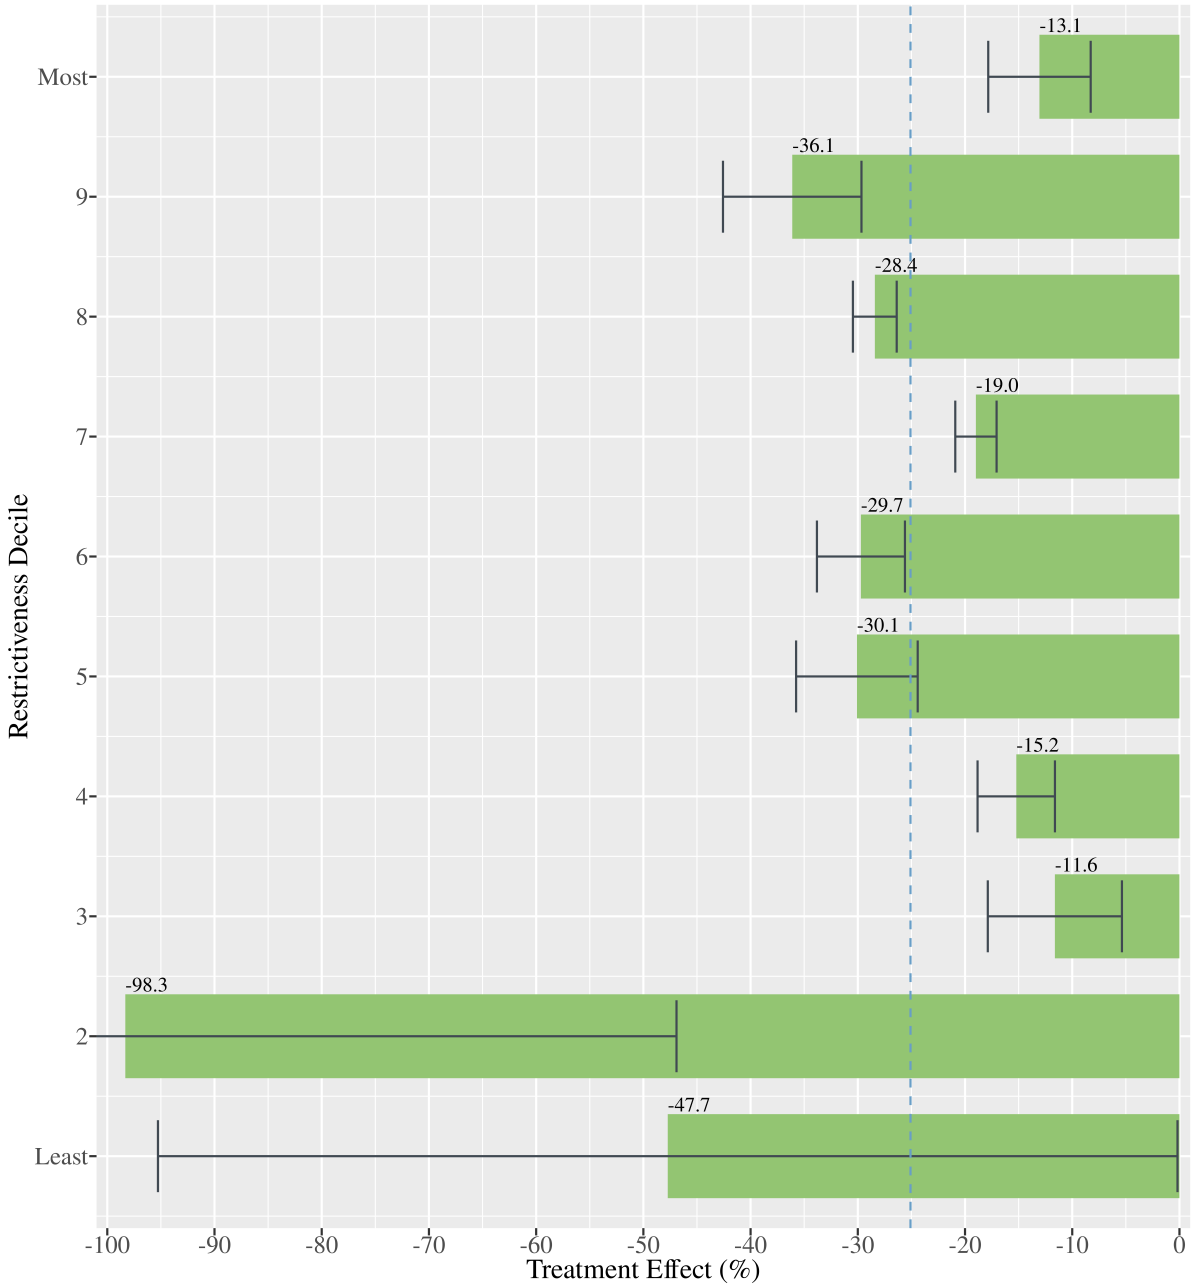

Notes: This figure presents the results from regressions of beneficiary utilization of a given drug on an indicator for whether the beneficiary's assigned plan put a prior authorization restriction on that drug. We run separate regressions on groups of drug-year pairs, where pairs are grouped into decile 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 3.2.

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

	Restrictions on Focal Drug in Enrolled Plan				
Restrictions on Focal Drug in Assigned Plan	0.950	0.913	0.908	0.908	0.908
	(0.002)	(0.003)	(0.003)	(0.002)	(0.002)
Exclusion of Focal Drug in Assigned Plan	0.002	-0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
	Exclusion of Focal Drug in Enrolled Plan				
Restrictions on Focal Drug in Assigned Plan	0.003	0.002	-0.001	-0.001	-0.001
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Exclusion of Focal Drug in Assigned Plan	0.950	0.918	0.905	0.905	0.905
	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)
Drug FEs		X			
Drug-year FEs			X		
Drug-market-year FEs				X	X
Substitution Controls					X
Number of drug-beneficiary-years	1,719,692,195				
Number of beneficiary-years	1,110,968				
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

	Restrictions on Focal Drug in Enrolled Plan in $t-1$	Exclusion of Focal Drug in Enrolled Plan in $t-1$
Restrictions on Focal Drug in Assigned Plan in t	-0.005	0.001
	(0.002)	(0.001)
Exclusion of Focal Drug in Assigned Plan in t	-0.000	-0.002
	(0.000)	(0.002)
F-statistic	4	2
Number of drug-beneficiary-years	1,212,656,940	
Number of beneficiary-years	833,643	
Number of market-years	207	
Average plans per market-year	6.0	
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
Restrictions on Focal Drug in Assigned Plan	0.015	0.000	0.003	0.002
	(0.018)	(0.000)	(0.002)	(0.001)
Reweighted Control Mean	2.280	0.016	0.462	0.368

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
Restrictions on Focal Drug in Assigned Plan	-0.000	-0.000	-0.000	0.000
	(0.003)	(0.003)	(0.005)	(0.001)
Exclusion of Focal Drug in Assigned Plan	-0.001	0.000	-0.000	-0.001
	(0.005)	(0.005)	(0.006)	(0.002)
Control Mean	0.582	0.605	64.8	3.95

Notes: This table presents estimates from a set of 'placebo' utilization regressions, where we regress indicators for a 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

	(1)	(2)	(3)	(4)	(5)
$\text{Auth}^{\text{Enrolled}}$	-1.169	-0.136	-0.098	-0.099	-0.108
	(0.012)	(0.005)	(0.004)	(0.003)	(0.004)
Auth^{Sub}					0.053
					(0.004)
PA % Effect	-290.0	-33.7	-24.3	-24.5	-26.9
Control Mean			1.299		
Reweighted Control Mean			0.403		
Drug FEs		X			
Drug-year FEs			X		
Drug-market-year FEs				X	X
Substitution Controls					X
Number of drug $\times$ beneficiary-years			1,719,692,195		
Number of market years			210		
Average plans per market-year			6.6		
Average beneficiaries per plan			803		
Average drugs per year			1569.2		

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..

Appendix Table A7: Estimates of the Effect of Prior Authorization Restrictions on Utilization: Additional Specifications

	(6)	(7)	(8)
$\text{Auth}^{\text{Enrolled}}$	-0.108	-0.091	-0.105
	(0.004)	(0.003)	(0.004)
Auth^{Sub}	0.048	0.391	0.048
	(0.004)	(0.019)	(0.007)
PA % Effect	-26.7	-22.5	-26.6
Control Mean		1.305	
Reweighted Control Mean	0.403		0.395
Drug-market-year FEs	X	X	X
Plan-by-cost FEs	X		
Plan-by-class FEs		X	
Number of drug $\times$ beneficiary-years	1,719,692,195		1,234,508,812
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 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 what we call the hourly wages method, the researchers convert employees' salaries to hourly wage equivalents, then price their time using those hourly equivalents.

[b] In what we call the opportunity costs 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. We thank Sarig (2024) for pointing this mistake out.

[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	All services	4.1%
	Advantage MCOs	and drugs	
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	Lower-extremity	6.1%
	New York University Hospital, 2017	venous procedures	
Wallace et al. (2020)	Single rheumatology clinic	Infusable drugs	
Initial application			21%
Final application			4%
Schwartz et al. (2021)	Large private insurer	Hosp. services	4.2%
		and drugs	
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: Savings per Administrative Dollar From Authorization Restrictions

Paperwork Cost	Request Rejection Rate 0%	1.5%	4%	7.5%	15%
$11.62	19.80	19.50	19.01	18.31	16.83
	(4.83)	(4.76)	(4.64)	(4.47)	(4.11)
$18.19	12.65	12.46	12.14	11.70	10.75
	(3.09)	(3.04)	(2.96)	(2.86)	(2.62)
$21.72	10.59	10.43	10.17	9.80	9.00
	(2.58)	(2.55)	(2.48)	(2.39)	(2.20)
$22.48	10.23	10.08	**9.82**	9.47	8.70
	(2.50)	(2.46)	**(2.40)**	(2.31)	(2.12)
$31.30	7.35	7.24	7.06	6.80	6.25
	(1.79)	(1.77)	(1.72)	(1.66)	(1.52)
$50	4.60	4.53	4.42	4.26	3.91
	(1.12)	(1.11)	(1.08)	(1.04)	(0.95)
$100	2.30	2.27	2.21	2.13	1.96
	(0.56)	(0.55)	(0.54)	(0.52)	(0.48)
$200	1.15	1.13	1.10	1.06	0.98
	(0.28)	(0.28)	(0.27)	(0.26)	(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 A11: 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 A12: 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 A13: 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 5.1.

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 CHA_2DS_2-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)
☐ Eliquis ☐ Xarelto
☐ Pradaxa

Nonpreferred
☐ Savaysa

IAPEC-X1664-19 December 2019

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.

ANTIPSYCHOTICS PRIOR AUTHORIZATION FORM

(form effective 1/5/21)

Keystone First | PERFORMRX℠ Next Generation Pharmacy Benefits

Fax to PerformRx℠ 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		☐ perphenazine tablet	☐ risperidone solution	

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
☐ amitriptyline/perphenazine	☐ Fanapt tablet	☐ Latuda tablet	☐ pimozide tablet	☐ Seroquel XR tablet	☐ Zyprexa Zydis
☐ aripiprazole ODT	☐ Fazaclo 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:

KF_201073500-7

D Demand Estimation: Additional Details

To overcome computational hurdles, we estimate our nested demand system in Section 3.4 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

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

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

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

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 $E[\log(\hat{s}_{gd})] \not\approx \log\left(E[s_{gd}]\right)$; 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 $P_{gd} \to 0$ 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.[47]

In Section 3.4, 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.[48]

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

[47] 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.

[48] 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:

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

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, $\hat{\tilde{\beta}}, \hat{\tilde{\delta}}$, and $\hat{\tilde{\Delta}}\kappa$, 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 $\hat{\mathcal{V}}_{gt} = \log\left(\sum_{k\neq 0} \exp(\hat{\tilde{V}}_{gdt})\right)$, the inclusive value of the inside goods, and a group fixed effect $\omega_{gt} = -\log\left(1 + \left(\sum_{k\neq 0} \exp(\tilde{V}_{gdt} + \tilde{\kappa}_{1m(gt)})\right)^{\lambda_C}\right)$. 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.[49]

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:

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

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,

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

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,

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

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

[49] While 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 beneficaries 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:

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

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

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:

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

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 $\hat{u}_{gdt}^2$.

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:

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

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).[50] 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

[50] We 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 $w_{gt} \sim \text{Gamma}(n, 1)$.[51]

For each therapeutic-class-specific drug demand estimation routine, we replace the Poisson pseudo-maximum-likelihood method with a **weighted** psuedo-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 A14, 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 A14 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 A7 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 A8 and A9 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 A14, 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 A10 and A11, 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.

[51] Note 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 A14: Class Level Nested Logit Summary Statistics for Top 30 Classes by Part D Spending

Class	Unweighted Market-Year Survival	Weighted Market-Year Survival	Share of Focal Days Supply, multi-drug users	Share of multi-drug users
Antihyperlipidemic Drugs, NEC	93.8%	98.6%	59.5%	23.9%
Psychother,Tranq/Antipsychotic	95.2%	99.0%	60.4%	31.9%
Antidiabetic Agents, Insulins	95.2%	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%	25.6%
Antineoplastic Agents, NEC	92.9%	99.1%	67.3%	6.6%
Biological Response Modifiers	79.5%	96.7%	72.6%	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	95.2%	99.3%	66.3%	23.9%
Cardiac Drugs, NEC	96.2%	99.3%	73.0%	20.2%
Antiplatelet Agents, NEC	81.9%	95.4%	63.2%	12.6%
Immunosuppressants, NEC	91.4%	98.7%	62.0%	15.4%
Misc Therapeutic Agents, NEC	95.7%	99.1%	59.8%	24.8%
Anticonvulsants, Misc	92.4%	98.9%	57.0%	19.0%
Cardiac, Calcium Channel	93.3%	98.6%	71.6%	10.5%
Coag/Anticoag, Anticoagulants	88.6%	95.6%	85.5%	15.6%
Cardiac, Beta Blockers	90.0%	98.6%	71.9%	7.2%
Parasympathomimetic, NEC	83.3%	95.3%	72.2%	7.0%
Eye/Ear/Nose/Throat Misc, NEC	91.0%	98.7%	53.7%	35.0%
Muscle Rel,Smooth-Genitour NEC	95.7%	99.2%	74.7%	15.3%
Analg/Antipyr,Nonstr/Antiinflm	96.2%	99.2%	71.2%	22.1%
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%	74.5%	7.9%
Vasodilating Agents, NEC	71.0%	92.9%	79.1%	18.5%
Phosphorus Removing Agents,NEC	74.3%	94.4%	68.7%	7.9%
Cardiac, ACE Inhibitors	70.5%	90.6%	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 A7: Share of Drug-Users Whom Take Multiple Drugs

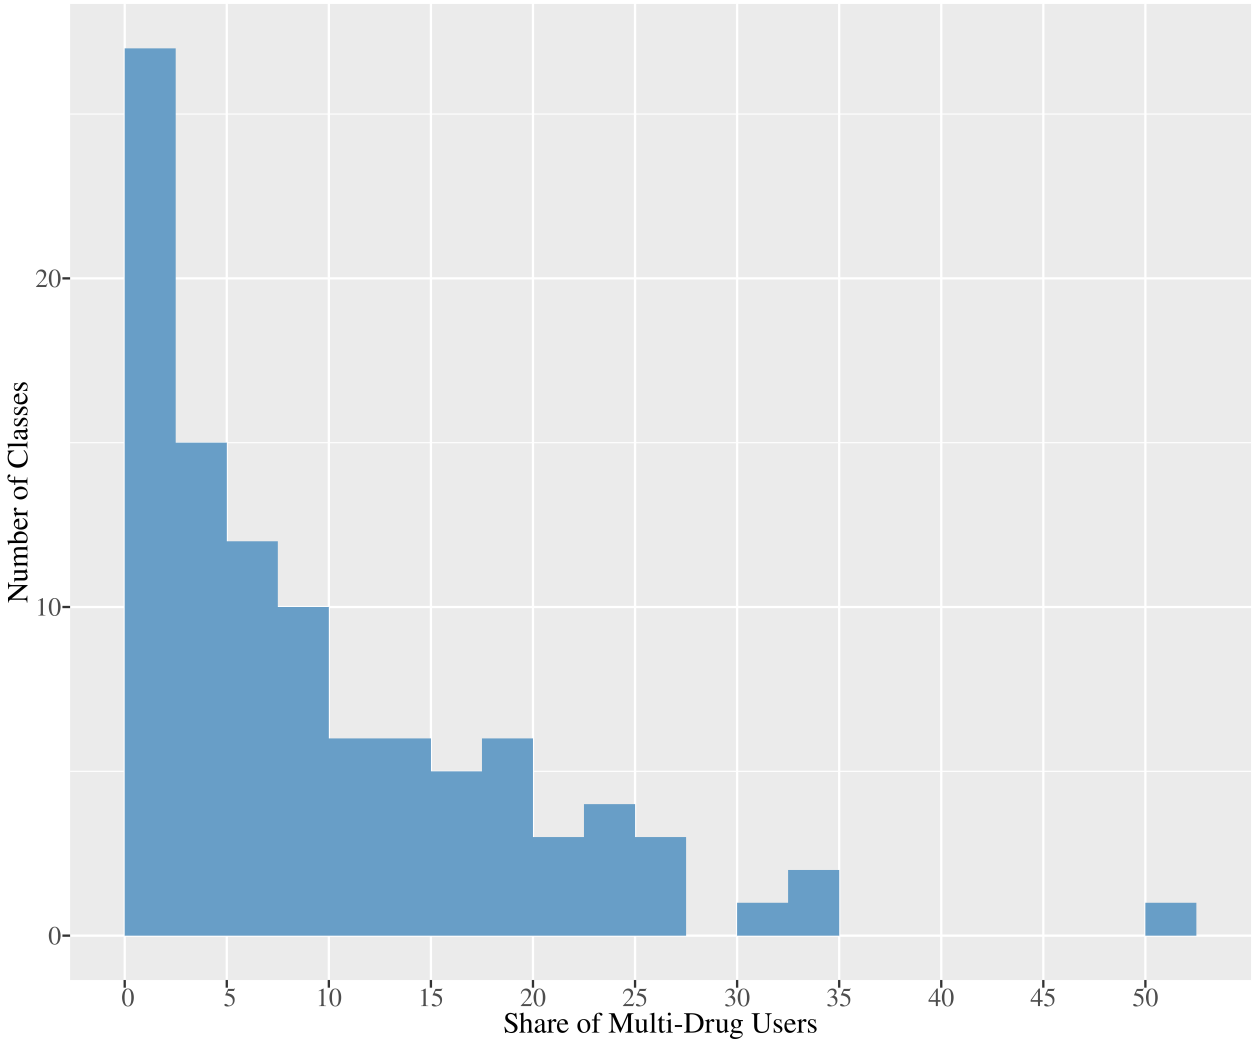

Notes: For each therapeutic class, we measure, out of the set of beneficiary-year pairs where the beneficiary took at least one drug within the class in that year, how many beneficiary-year pairs were ones in which the beneficiary took at least two drugs in the class. This figure plots the distribution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

Appendix Figure A8: Focal Drug Days Supply Share for Multi-Drug Users

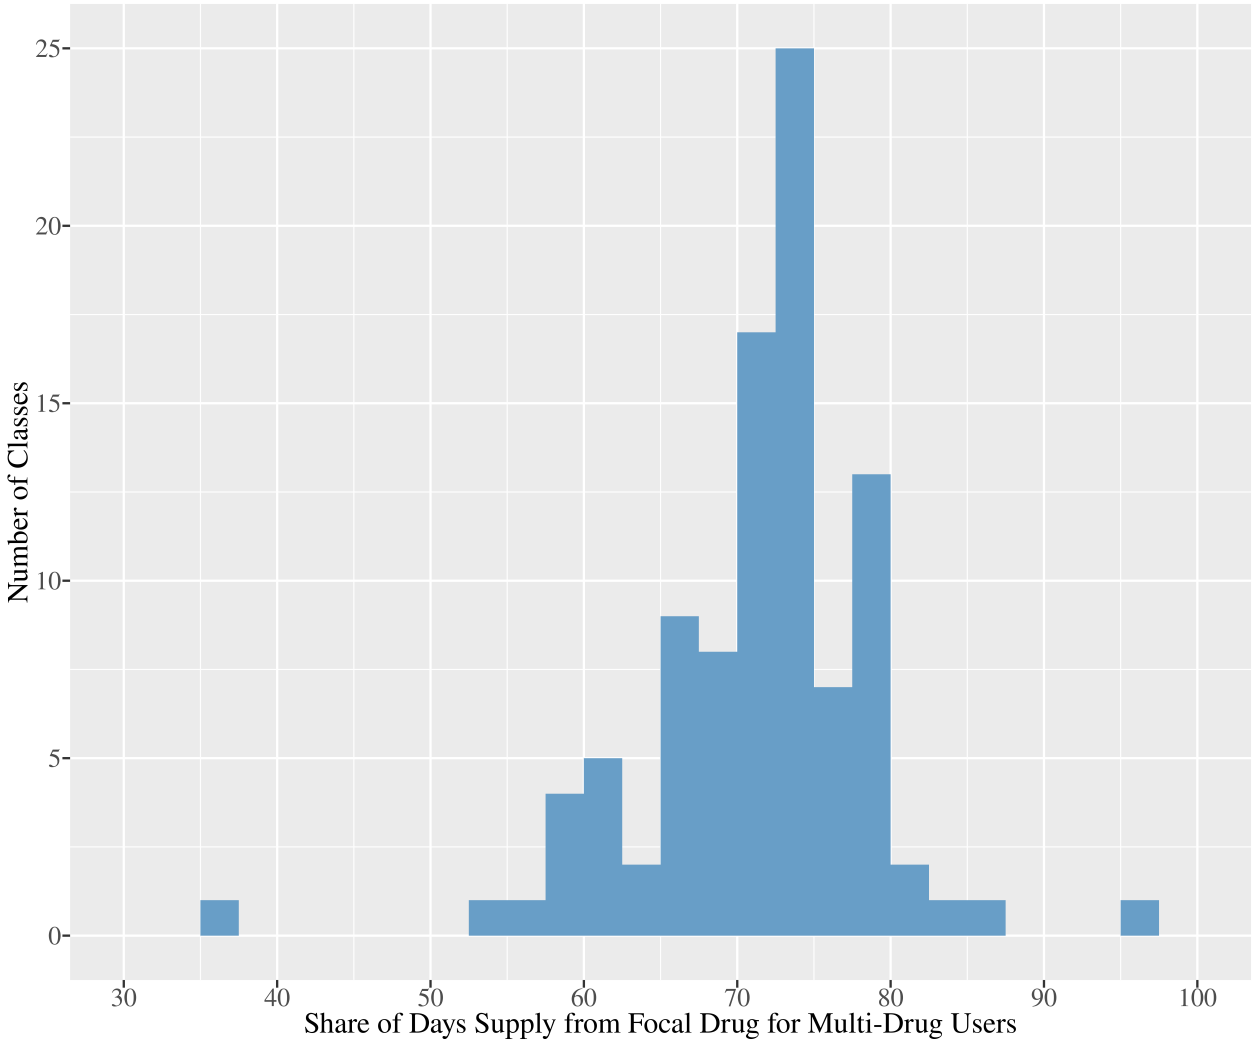

Notes: For each therapeutic class, we measure, out of the set of beneficiary-year pairs where the beneficiary took at least two drugs within the class in that year, what share of days supply in that class were accounted for by the focal (most-used) drug. This figure plots the distribution of that statistic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

Appendix Figure A9: Focal Drug Days Supply Share for All Drug Users

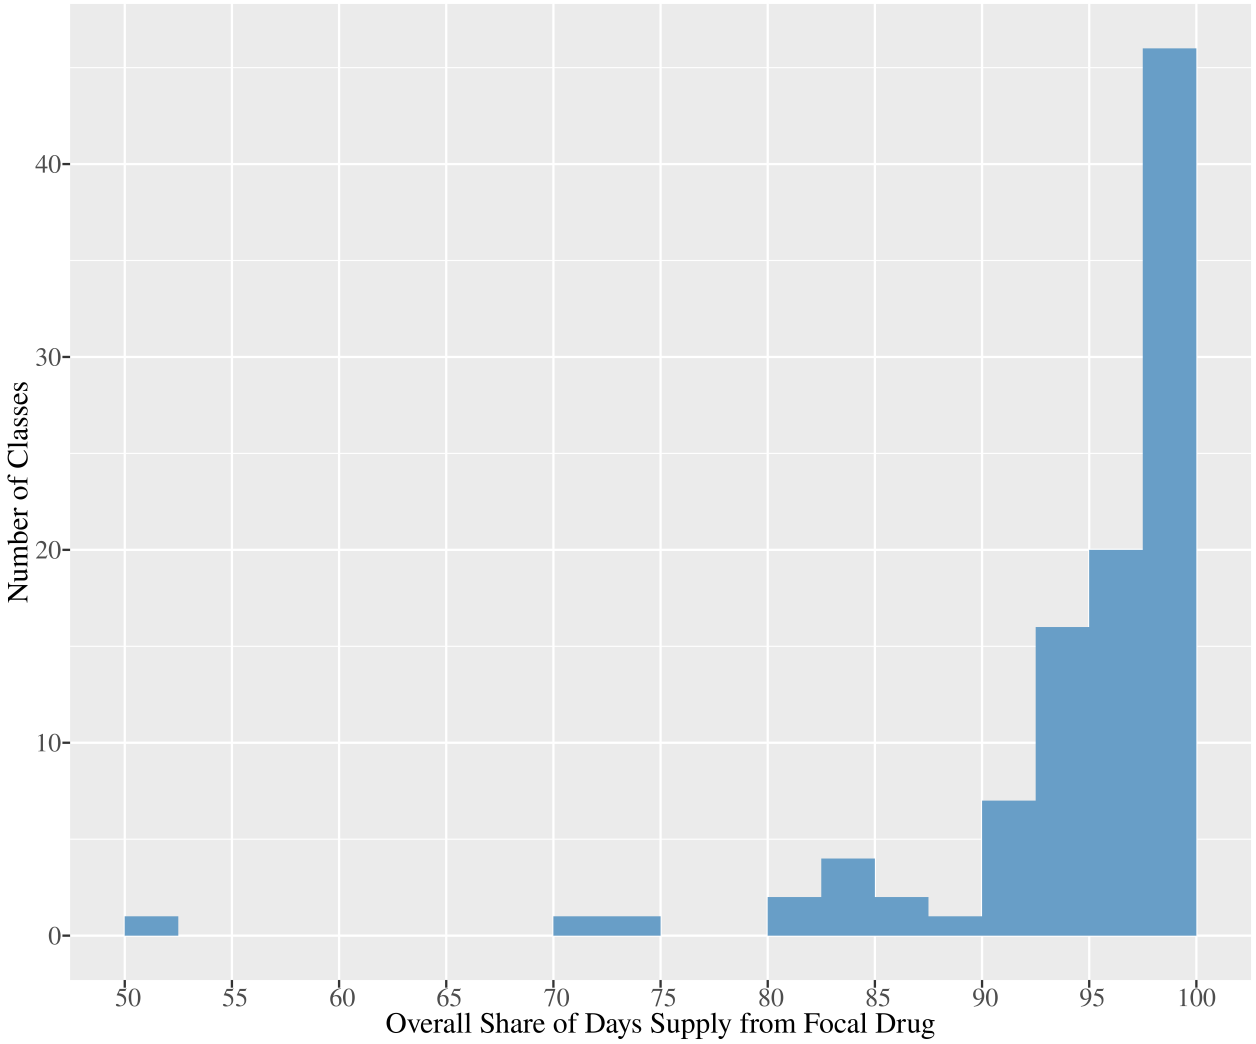

Notes: For each therapeutic class, we measure, out of the set of beneficiary-year pairs where the beneficiary took at least one drug within the class in that year, what share of days supply in that class were accounted for by the focal (most-used) drug. This figure plots the distribution of that statistic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

Appendix Figure A10: Unweighted Market-Year Survival After Logit Restrictions

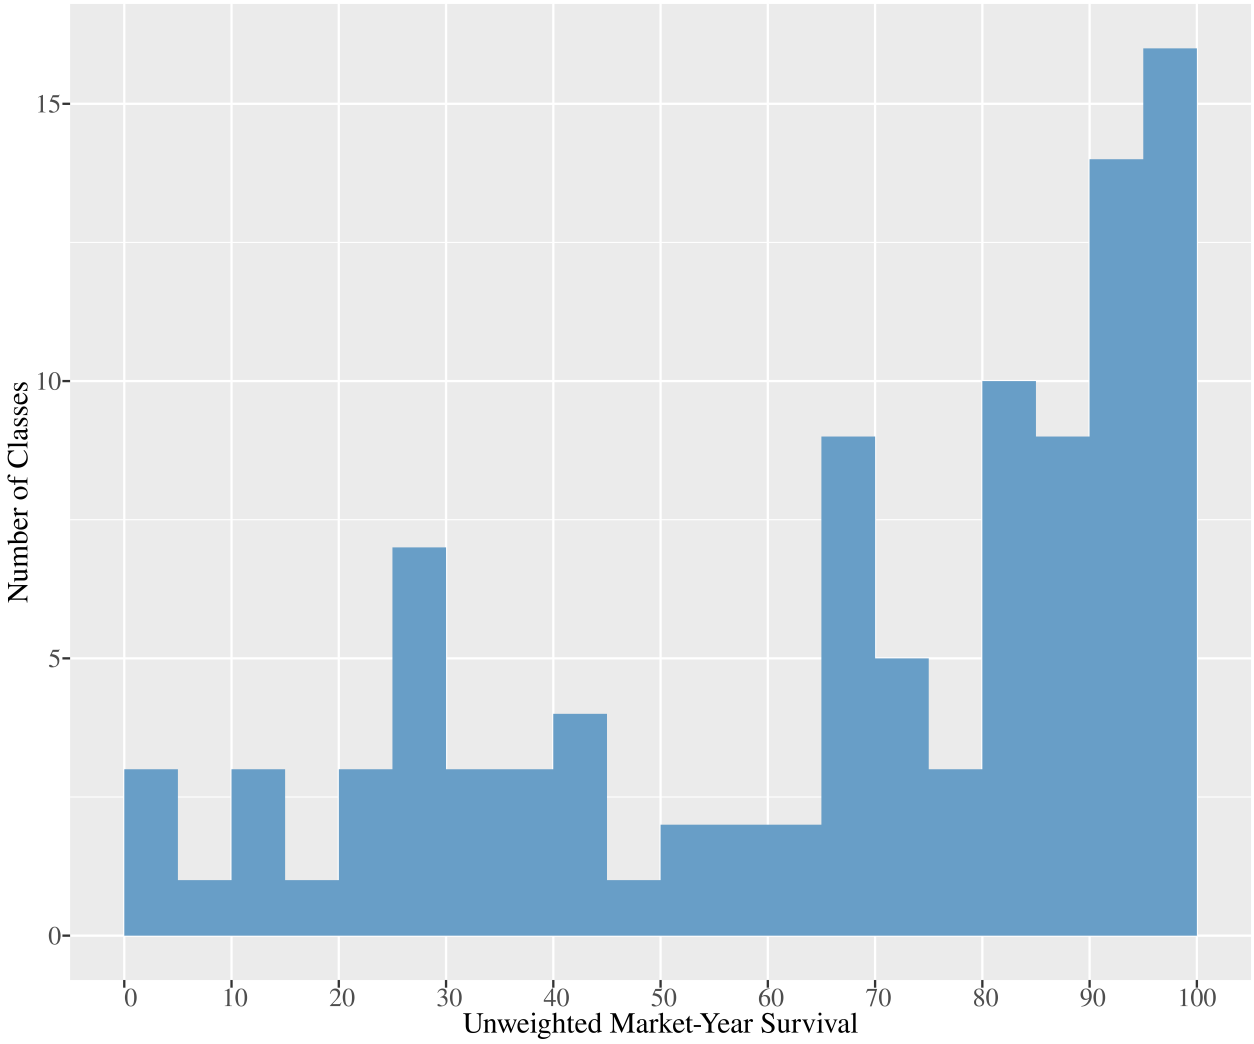

Notes: This figure plots the distribution, across classes, of the share of markets (region-year pairs) that have (1) at least one drug in the class that is under a prior authorization restriction in between 0 and 100% of plans; and (2) where at least two drugs in the class are filled.

Appendix Figure A11: Beneficiary-Weighted Market-Year Survival After Logit Restrictions

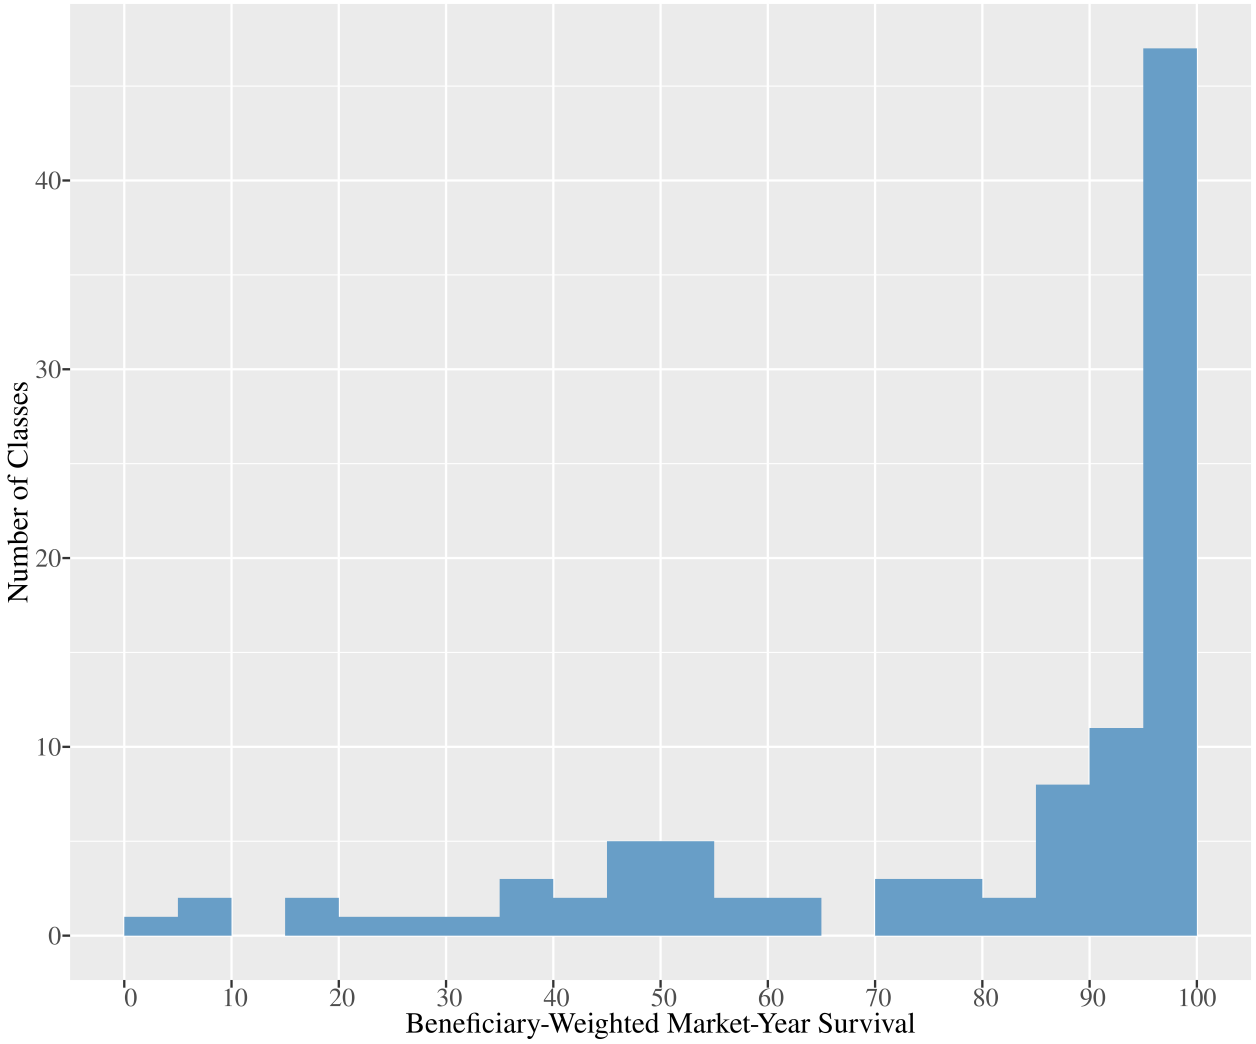

Notes: This figure plots the distribution, across classes, of the share of markets (region-year pairs) that have (1) at least one drug in the class that is under a prior authorization restriction in between 0 and 100% of plans; and (2) where at least two drugs in the class are filled. In this figure,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

E Revealed Preference Analysis: Additional Details

E.1 Aggregation of Drug-Level Consumer Surplus

In Section 5.1 we describe the construction of drug-specific measures of willingness-to-pay. Here we describe the exact procedure by which we construct the aggregate measures underlying Table 10.

We begin by constructing a demand curve for each drug. Each drug has demand equal to

$$D_d(p_d) = D_d(0) - \gamma p_d$$

where $D(0)$ is the amount demanded when the price is zero, and γ is the demand slope. We obtain γ by dividing the effect of the LIS transition on fraction of beneficiaries with a prescription for the drug by its effect on out-of-pocket prices faced. In our main specification, there is a single estimate of γ that is applied to each drug individually. In our specification where we allow for class-specific demand slopes, we have class-specific estimates of both of these objects that are assigned to each drug.

Given this structure, total willingness to pay for the foregone drug Δq_d in the "best-case" scenario is given by

$$\Delta CS_d^{\text{Best Case}} = \frac{1}{2}\Delta q_d^2 \frac{1}{\gamma}$$

We compute Δq_d as the fraction of beneficiaries with at least one prescription for drug d in a year among those who do not face prior authorization (this average is based on our main sample and appears as the"Control Mean" in Table 5), multiplied by the share of beneficiaries who face prior authorization on drug d, multiplied by -0.269, the proportional reduction in fraction of beneficiaries with a prescription for the drug due to prior authorization. We use the proportional change rather than the absolute level effect because different drugs have different base rates, and simply applying the estimated -0.108 effect will result in negative average quantities for many drugs. We compute $\Delta CS = \sum_d \Delta CS_d$. To construct confidence intervals, we generate random draws of the coefficients from the LIS transition regressions and calculate consumer surplus loss as described above.

Similarly, when we consider the random case, we compute $D(0)$ as the mean days supply for drug d among those who do not face prior authorization multiplied by the share of beneficiaries who face prior authorization on drug d. We then compute

$$\Delta CS_d^{\text{Random}} = 0.269\frac{1}{2}D(0)^2\frac{1}{\gamma}$$

We also estimate constant-semi-elasticity demand function rather than a linear one. For this, we estimate effects on quantity using Poisson regression, and compute the semi-elasticity as the Poisson quantity effects divided by linear price effects. This assumes that the relevant demand curve is of the form $D(p) = D(0)e^{\epsilon p}$ where ϵ is the price semi-elasticity of demand. Under this demand curve, the total willingness to pay for a quantity change δq_d is

$$\int_{q=D(0)-\delta q_d}^{D(0)} \frac{1}{\epsilon}\log\left(\frac{q}{D(0)}\right)dq = \frac{1}{\epsilon}\left[q\log\left(\frac{q}{D(0)}\right)-q\right]_{q=D(0)-\delta q_d}^{D(0)}$$

noting that $D^{-1}(q) = \frac{1}{\epsilon}\log\left(\frac{q}{D(0)}\right)$ and that, while its antiderivative is undefined at $q=0$, $\lim_{q\to 0} q\log(q) - q = 0$.

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 5.1, $W_d(\theta_i)$ is drawn from a zero-inflated uniform distribution with a mass at zero of $1 - D(0)$ and a cumulative density function slope of ζ, then, as in the prior section, this structure gives rise to a demand curve that depends on administrative costs, $D_d(a) = D(0) - a\zeta$. 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 demand slope, we simply use the demand response to prior authorization restrictions that we observe in Sections 3 and 3.4. In response to prior authorization, providers prescribe restricted drugs 0.108pp less. Our baseline calibration of provider-facing cost is \$22.48. These two numbers imply that the administrative cost semi-elasticity of prescription is $\epsilon = \frac{0.108}{22.48} \approx 5 \times 10^{-3}$. In contrast, the slope of the demand curve revealed in Section 5.1 is approximately 2×10^{-4}. By this calibration, providers are an order of magnitude more responsive to administrative costs than patients are 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 \$1.21and \$4.52 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. For the best-case and random case scenarios, the maximum ρ to make prior authorization inefficient is 0.01 and 0.05 respectively. ρ has a stronger economic interpretation in this case. This calibration reveals that, to rationalize prior authorization being inefficient *and* to think that the quantity changes reflects physician decision-making, the fact that physicians respond strongly to small administrative costs implies that they put very low weight on their patients' welfare relative to their own. Indeed, screening mechanisms that rely on physician agency may be inefficient if physicians are poor agents for their patients.