Optimal Policy with Opposing Externalities: Evidence from Antibiotics

Edward L. Kong, PhD* Erica S. Shenoy, MD, PhD†
Alyssa R. Letourneau, MD, MPH†

March 22, 2026

Preliminary draft — Comments welcome; please do not cite or distribute.

Please click here for the most recent version.

Abstract

The threat of antibiotic resistance has generated substantial policy attention toward reducing antibiotic use. However, antibiotic use creates a second, opposing externality: innovation incentives rely on revenues from antibiotic sales. While direct payments for innovation have been proposed to decouple use and revenues, their optimal size and welfare consequences remain poorly understood. This paper introduces a framework for optimal antibiotic policy that accounts for use, resistance, and innovation. We estimate key parameters—the demand for antibiotics, the use-resistance externality, and the effect of use on mortality—using electronic health records data and quasi-experimental variation in antibiotic restrictions. We embed these estimates in a structural model to evaluate optimal demand incentives and innovation prizes. We find that antibiotic use lowers mortality but generates substantial resistance externalities: reducing use by 20% lowers resistance by 60%. However, the innovation externality is large and acts in the opposite direction. We show that, contrary to conventional wisdom, antibiotics may often be underused rather than overused. Optimal innovation prizes are complementary with greater utilization across a wide range of policies. We find that jointly optimal prizes and utilization policies decrease mortality in our sample by 5.9% and raise the present discounted value of welfare by $46 to $212 billion. More broadly, the results illustrate how policies that target opposing externalities can yield large welfare gains when multiple policy instruments are used in combination.

*Harvard University; edwardkong@hms.harvard.edu. We thank Mark Shepard, David Cutler, Tim Layton, Ariel Pakes, Jeffrey Pearson, David Kubiak, and Jason Abaluck for valuable feedback. We thank participants of the CEPR Health Economics Conference and seminar participants of the Harvard Labor/Public Economics, Industrial Organization, and Health Care Policy seminars for helpful comments and suggestions. Edward Kong received funding from award number T32AG51108 from the National Institute on Aging. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute on Aging or the National Institutes of Health. All errors are our own.

†Massachusetts General Hospital and Harvard Medical School.

1 Introduction

Antibiotics are a cornerstone of modern medicine, yet their effectiveness is increasingly threatened by antibiotic resistance. Antibiotic use generates resistance both for the individual patient and for others within the same healthcare environment, creating a classic externality. Although the precise magnitude of this externality is uncertain, the prevailing view is that antibiotics are overused relative to the social optimum (Årdal et al., 2019). In response, antibiotic stewardship programs (ASPs), which restrict and guide antibiotic prescribing within hospitals, have become central institutions for antibiotic management and were federally mandated for all U.S. acute-care hospitals by the Centers for Medicare and Medicaid Services in 2016 (CMS, 2019).

However, there is another externality to antibiotic use. Under the current reimbursement system, antibiotic profits sustain the revenues that fund research and development (R&D), implying that reduced use may weaken incentives to innovate (Outterson et al., 2015). Taken to the extreme, if new antibiotics were held entirely in reserve, the expected return to antibiotic R&D would vanish, reducing innovation to zero and ultimately worsening health outcomes. Balancing these multiple externalities poses a major challenge for both public health and the economics of innovation. Policymakers face a central question: how should society allocate resources between conserving existing antibiotics and stimulating the discovery of new ones?

Beyond antibiotics, examples of opposing externalities abound across multiple other important policy areas. For example, hunting can deplete wildlife populations but can also yield economic benefits and incentivize conservation efforts (Belant et al., 2025; Chapagain and Poudyal, 2020; Lindsey et al., 2007). Rent control may improve short-run housing affordability but decrease long-run investments in the quantity and quality of housing (Autor et al., 2014; Diamond et al., 2019). Climate change is a particularly salient example that has been explicitly likened to antibiotic resistance (Roope et al., 2019). The effect of economic activity on rising global temperatures parallels the effect of antibiotic use on resistance. Reducing economic output may avert large future damages (Bilal and Känzig, 2024; Fankhauser and S.J. Tol, 2005; Nordhaus, 2018), but also incur immediate costs (Dasgupta et al., 2002; Glavina et al., 2025), leading to debates around how to balance climate efforts with economic growth (Stern and Stiglitz, 2023).

While it is common for economic policies to incur trade-offs, these examples bear the common trait that expanding the space of available policies can improve welfare. Across the above settings, this could take the form of direct investments in new antibiotic development, environmental conservation, housing quality, or greener energy production. In the case of antibiotics, numerous proposed policies bear the common characteristic that they *de-couple revenues from sales volume*. In all of these cases, determining the optimal combination of policies requires quantifying the key opposing externalities and their joint contributions to welfare.

Our paper makes three contributions. First, we build an economic framework for understanding the trade-offs induced by the opposing externalities of antibiotic use. Second, we provide new reduced-form estimates of the use-resistance and use-mortality elasticities using novel data from electronic health records. Finally, we estimate a full structural model to simulate counterfactual

policies that vary the amount of antibiotic use and the size of innovation prizes.

Our first contribution is to introduce a simple conceptual framework of optimal antibiotic use that builds intuition for how observed antibiotic use might differ from the social optimum and how innovation prizes might increase welfare. In the framework, optimal use equates marginal benefits and marginal costs, but benefit and cost curves differ between private and social planner objective functions. These differences allow us to make specific predictions about how equilibrium antibiotic use, determined by private actors such as hospitals and physicians, might differ from socially optimal use.

Despite its simplicity, the model yields several important insights. First, it clarifies what we mean by *opposing externalities* and the role of direct innovation payments. Socially optimal antibiotic use trades off the marginal costs of higher resistance against marginal benefits, which include both contemporaneous mortality reductions from treating current infections and the dynamic benefits of increased innovation. Hospitals, following medical guidelines for antibiotic use, internalize the resistance externality but not the innovation externality. They also pay antibiotic prices that far exceed social marginal costs of production. These differences between hospital and social planner objectives lead to antibiotic *underuse* and suboptimally low levels of innovation. Direct innovation payments, such as innovation prizes, decouple use from innovation by compensating drug developers independently of sales volume.[1] More broadly, this reflects a general feature of environments in which a single action generates multiple opposing externalities: expanding the set of policy instruments can improve welfare by targeting each externality separately.

The model also delivers several seemingly counterintuitive implications. The conventional wisdom is that direct innovation payments lower antibiotic use by de-linking profits from sales (Årdal et al., 2019). However, we show that the effect of innovation payments on optimal antibiotic use is theoretically ambiguous: introducing an innovation prize could imply higher or lower use. Moreover, we show that optimal use reflects a classic marginal–inframarginal trade-off: the benefits of increased use accrue to marginal patients while resistance costs are borne by all inframarginal users. As a result, optimal antibiotic use is governed solely by the relationship between use and resistance; treatment effect magnitudes play little if any role in setting the optimal use level. Following this logic, the model yields a simple empirical test for assessing whether a given antibiotic is likely to be under- or overused. The test says that a given antibiotic is likely underused in a population if its resistance-use elasticity is less than the odds of non-resistance. Importantly, the model illustrates why optimal policy may sometimes involve expanding, rather than restricting, antibiotic use.

Finally, the conceptual framework highlights the key empirical elasticities that govern optimal policy: (1) the elasticity of resistance with respect to antibiotic use, (2) the effects of use and resistance on mortality, and (3) the responsiveness of innovation to market size.

Our second contribution is to provide new causal evidence on how resistance and health outcomes respond to antibiotic use in hospitalized patients, corresponding to elasticities (1) and (2)

[1] Challenges with implementing innovation prizes are well-described in the literature on optimal procurement (Weyl and Tirole, 2012, e.g.,). A further challenge in the antibiotics setting is distinguishing social value apart from market demand.

above. A major empirical challenge is the difficulty in observing antibiotic use and resistance among hospitalized patients, as this information is not available in standard health insurance claims datasets. To surmount this challenge, we assemble 12 years of electronic health records data on 386,498 hospital admissions and 416,203 bacterial isolates from two academic medical centers. For every isolate, we observe resistance to every relevant antibiotic in the choice set. This reduces selection into testing and allows us to directly observe individuals' resistance status under counterfactual antibiotic choices. Another challenge is identification: antibiotic use may be correlated with unobservable factors that affect both resistance and mortality. For example, sicker patients are more likely to both require antibiotics and harbor resistance. Reverse causality is also a concern, as higher resistance or lower mortality risk could result in lower use. We use two sources of identifying variation: for resistance, we use hospital-specific changes in antibiotic restriction policies. For mortality, we leverage variation in staffing for the providers in charge of approving antibiotic use within the hospital. We focus on gram-positive infections and the broad-spectrum antibiotic linezolid. While this limits external validity, our goal is obtain credible parameter estimates to embed in a modeling framework that can be adapted to other settings. We also test the sensitivity of our model results to a wide range of parameter estimates.

While it is conventional wisdom that higher antibiotic use leads to higher resistance, existing evidence in the medical literature is limited primarily to in-vitro experiments on bacterial cultures or single-center studies comparing resistance before versus after stewardship interventions, without control groups (Schuts et al., 2016; Adda, 2020). Thus, whether antibiotic use affects resistance in real-world hospital settings remains incompletely understood. We identify the causal effects of stewardship policies on antibiotic demand and resistance using a difference-in-difference framework leveraging two policy changes affecting antibiotic restrictions: policy 1 created formal paid positions for antibiotic stewardship of restricted antibiotics (including linezolid), while policy 2 established narrower approval criteria for linezolid in particular. Both policy changes occurred in one hospital but not the other. Additionally, we employ a triple-difference approach that leverages an additional comparison between linezolid and non-restricted antibiotics, which were not subject to either policy.

Our estimates show that the two antibiotic restriction policies, taken together, reduce use by about 20% and decrease resistance by about 60%. Moreover, while individuals' antibiotic use affects their future chance of resistance, at least 76% of the use-resistance relationship operates via hospital-wide resistance externalities affecting all patients. Our estimates are comparable in magnitude to meta analyses of before vs. after estimates from the medical literature, which find that stewardship programs are associated with a 51% to 63% decline in the incidence of resistant infections (Baur et al., 2017).

Next, we recover new causal estimates of how restrictions on antibiotic use affect short-run mortality among high-risk, hospitalized patients. Estimates of this effect are not available in the literature. While clinical trials of new antibiotics report mortality, they are ethically constrained to test non-inferiority to established, curative antibiotics, in pathogens known to be susceptible to both treatment and control antibiotics (e.g., Kohno et al., 2007; Falagas et al., 2008; Weigelt et al.,

2005; Wunderink et al., 2012; Wilcox, 2003).[2] Thus, the causal impact of increasing or decreasing the overall *level* of antibiotic use on the margin is not well understood. A notable exception is a recent cluster-randomized trial in rural Niger that compared blanket twice-yearly azithromycin treatment versus placebo in children under 5 years old—over two years, azithromycin use decreased mortality by 14% (O'Brien et al., 2024).

Identification of the mortality effects of antibiotics using observational data is challenging because treatment is selectively administered to sicker patients. We overcome this challenge using quasi-random variation in antibiotic approvals, in what is commonly known as a judge or examiner design (Kling, 2006; Goldsmith-Pinkham et al., 2025; Chyn et al., 2025). In many hospitals, including those in our sample, restricted antibiotics can only be prescribed after approval through an antibiotic pager. Treating physicians submit requests to this pager, which is staffed by a specialist infectious disease provider (hereafter, a "pager holder") who either approves or denies each request. At any point in time, only one pager holder is on duty, and the identity of this provider rotates weekly. We combine provider staffing schedules with persistent differences in approval propensities across pager holders to construct quasi-random variation in antibiotic use. We find that approval rates vary significantly across pager holders and are uncorrelated with patient characteristics. For patients with high predicted antibiotic use, we find that earlier antibiotic treatment significantly reduces 30-day mortality.

By combining multiple identification strategies, we provide new causal evidence for how antibiotic stewardship practices affect both hospital-wide resistance externalities and short-run health outcomes for individual patients. Together, these estimates clarify the trade-offs inherent in regulating antibiotic use.

Our third contribution is to develop and estimate a structural model to understand the impacts of counterfactual policies, including use restrictions and innovation prizes. The model parallels our conceptual framework, and we identify the key parameters of the model using the identification strategies from our reduced-form analyses. A key advantage of our approach is that estimated mortality effects allow us to measure welfare in terms of lives saved, rather than relying on consumer surplus inferred from demand. Combining our demand, resistance, and mortality estimates, we show that the entry of a single new antibiotic yields a 2.8% reduction in mortality among the patients in our sample. Scaled to the level of the U.S., this represents an increase in the present-discounted value (PDV) of welfare of $54 billion. Of note, this finding, that new antibiotics are highly socially valuable, does not depend on the innovation elasticity.

We model antibiotic entry as an endogenous decision based on expected profitability, following a long tradition in industrial organization (Berry and Reiss, 2007). We simulate expected profits by combining our demand model, scaled to the global level, with hospital administrative data on antibiotic procurement costs. We estimate the distribution of entry costs using data on FDA approvals of antibiotics, combined with net-of-rebate sales data for antibiotics between 1986 and

[2]These randomized trials are limited to pathogens that are susceptible to both the treatment antibiotic linezolid as well as the control antibiotic(s).

2024, following Pakes et al. (2007). Our preferred specification yields an innovation elasticity with respect to net revenues of 0.89; however, we conduct sensitivity analyses across a wide range of published elasticities, spanning 0.25 to 2.75 (Dubois et al., 2015).

Our counterfactual simulations uncover new evidence of complementarities between demand-based policies and innovation prizes—two widely used or proposed policy tools. To induce entry, the optimal innovation prize awards the innovator about $1.0 billion per year for 10 years and raises the PDV of welfare by $6.5 billion, net of prize costs. Even without direct innovation payments, demand-side policies, such as changes to stewardship restrictions, taxes, or subsidies, can increase welfare. We find that antibiotics are underused at baseline, such that optimal demand incentives *increase* antibiotic use. Resistance increases also occur, but these are outweighed by reductions in mortality and increased innovation. Raising antibiotic use to its optimal level increases present-discounted welfare by $89 billion. Benefits from optimizing antibiotic use exceed gains from introducing innovation prizes, although both effects are large.

Allowing for both policy instruments yields welfare improvements that exceed the sum of their individual gains, due to the complementarity between them. The optimal policy combination reduces mortality in our sample by 5.9% and raises present-discounted welfare by $100 billion. This ranges between $46 to 212 billion across sensitivity analyses varying the resistance, mortality, and innovation elasticities. Overall, the structural model simulations demonstrate two key insights: demand and innovation are often complements, and efforts to suppress resistance too aggressively can unintentionally raise mortality and dampen innovation. Our results reinforce the need for policies that balance conservation with access and discovery.

To our knowledge, this paper is one of the first to develop a structural framework in health care that jointly models provider treatment decisions, downstream mortality consequences, and the role of innovation in shaping the underlying tradeoffs. More broadly, this paper contributes to public economics by analyzing policy design in settings where addressing one externality can create or amplify another. In such environments, single-instrument solutions—such as Pigouvian taxes or subsidies—may fail to achieve the social optimum because mitigating one distortion (resistance) can induce offsetting effects elsewhere (innovation). By modeling these opposing forces, our paper demonstrates how multiple policy levers can act as complements, generating welfare gains that exceed the sum of each policy in isolation. This approach extends beyond antibiotics to other domains of public policy—ranging across climate change, environmental conservation, and housing—where interacting externalities make coordinated, multi-instrument interventions essential for maximizing social welfare.

Related Literature. This paper contributes to a growing economic literature recognizing the importance of antibiotic resistance, innovation, and the optimal use of antibiotics. Early work applied economic theory to the problem of antibiotic resistance, for example, by showing how property rights can affect the use-resistance relationship (Horowitz and Moehring, 2004) or how randomized treatment protocols could slow resistance evolution (Laxminarayan and Weitzman,

2002). This literature likened the problem of maintaining antibiotic effectiveness to natural resource management and recognized that optimal policy depends on the relationships between antibiotic use and resistance, the spread of infections, and the pace of innovation (Herrmann and Laxminarayan, 2010). We augment these theoretical contributions by developing a unified model that captures the opposing externalities linking antibiotic use, resistance, and innovation.

Empirical work on use and resistance is rare and largely limited to the outpatient setting. A large epidemiological literature has calibrated models of resistance dynamics under alternative utilization regimes, but typically abstracts from the economic incentives governing antibiotic development and use (Niewiadomska et al., 2019). Within economics, Adda (2020) identifies the causal effect of human antibiotic use on resistance using temperature-driven variation in prescribing. Recent work by Dubois and Gokkoca (2023) shows that resistance to an antibiotic reduces its use, controlling for endogeneity using antibiotic sales in veterinary medicine. They also show that restricting antibiotic use reduces consumer surplus, and that improving rapid antibiotic susceptibility testing is beneficial. Our estimates are consistent with these findings and extend this literature to the hospital setting, where most novel broad-spectrum antibiotics are first introduced. Additionally, by incorporating mortality and innovation, we provide new theory and evidence on situations where antibiotics may be underused and the complementarities between innovation prizes and utilization.

To model the innovation process, our paper leverages prior work on the elasticity of drug innovation with respect to profits and market size, both in the specific case of antibiotics (Kong and Zhao, 2025; Majewska, 2022) and for prescription drugs in general (e.g., Acemoglu and Linn, 2004; Finkelstein, 2004; Dubois et al., 2015; Agha et al., 2022; Blume-Kohout and Sood, 2013; Gaessler and Wagner, 2018). Our analysis of policy counterfactuals relates to theoretical work on the optimal design of rewards or procurement mechanisms for innovation, including intellectual property, prizes (Weyl and Tirole, 2012), and advanced market commitments (Kremer et al., 2020). In our model, we intentionally lean on this large body of prior work and refrain from specifying a detailed model of the innovation process itself. Instead, our detailed estimates of demand, resistance, and mortality allow us to focus on modeling (1) how counterfactual policies affect profits and (2) how innovation affects welfare.

The remainder of the paper is organized as follows. Section 2 provides institutional background on antibiotic resistance, innovation, and pricing; ongoing policy proposals; and the specific antibiotics and hospital-acquired infections we focus on. Section 3 introduces our conceptual framework, which formalizes the relationships between antibiotic use, resistance, innovation, and mortality; offers graphical intuitions for optimal use; and derives additional results. Section 4 describes the electronic health records data and other datasets we use for estimation. Section 5 details the full structural model, including identification strategies, reduced-form evidence, estimating equations, and model parameter estimates. Each subsection addresses a single equation of the model, starting with demand, followed by resistance, mortality, and finally innovation. Section 6 uses the estimated structural model to simulate welfare implications for counterfactual use restrictions and innovation prizes. Section 7 discusses policy implications and concludes.

2 Background

Scope and mechanisms of antibiotic resistance. Antimicrobial resistance poses a major threat to global health. In 2019, bacterial resistance was estimated to cause 1.27 million deaths worldwide (Murray et al., 2022) and over $80 billion (in 2024 dollars) in additional healthcare costs and lost productivity in the U.S. alone (CDC, 2013). Higher antibiotic use can lead to increased resistance. One underlying mechanism is selective pressure: use of broad-spectrum antibiotics kills susceptible bacteria while allowing resistant strains to proliferate and spread within and across hospitals. In addition to the effect of use on resistance, another key empirical object is the fitness cost incurred by bacteria to become resistant (Lipsitch, 2001). With low fitness costs, resistance tends to persist even in the absence of continued selective pressure from antibiotic use. With high fitness costs, resistance tends to decline in the absence of continued use. In accordance with the literature on hospital-acquired infections, our results are consistent with high fitness costs.

In the U.S., ASPs have become the principal mechanism to curb resistance in hospitals and are now mandated for acute-care hospitals (CMS, 2019). The inpatient setting accounts for most broad-spectrum antibiotic use and represents the front line in the struggle between new drugs and emerging resistance. Hospital ASPs regulate the use of broad-spectrum antibiotics, drugs effective against a wide range of bacteria that are often administered intravenously to critically ill patients. However, the overall effectiveness and potential unintended consequences of specific stewardship policies, such as the antibiotic demand restrictions we study, remain incompletely understood.

Antibiotic innovation. In the antibiotic setting, there is increasing recognition among experts that the barriers to antibiotic innovation are economic rather than scientific in nature. Despite the over 400 diverse antibiotic projects preclinical in development globally, large pharmaceutical companies have largely abandoned this market, and smaller firms launching novel antibiotics have gone bankrupt (Årdal et al., 2019). Recent empirical estimates of antibiotic-specific innovation elasticities (Kong and Zhao, 2025; Majewska, 2022) likewise suggest a causal link between expected revenues and innovative activity, such as patenting and clinical trials. Motivated by this evidence, we model the antibiotic entry decision as a function of expected revenues.

The economic challenges affecting antibiotic innovation are inextricably linked to patterns of use and resistance. Firms' incentives to develop new antibiotics depend on expected use, which may be suppressed due to resistance concerns. Antibiotics are unique among drugs in this aspect, where the newest and most innovative therapies are intentionally held in reserve (Outterson et al., 2015, 2022). Short treatment courses (ranging from a few days to a few months) also reduce expected profits (Årdal et al., 2019). Because existing antibiotics are often close substitutes until resistance emerges, hospitals and health systems behave as elastic purchasers, further reducing expected returns to innovation. Policies that restrict antibiotic use thus create a trade-off: reducing consumption may slow resistance, but this may unintentionally weaken the commercial case for developing the next generation of antibiotics.

As pathogens continue to evolve new resistance mechanisms, a central question is whether

the supply of truly novel antibiotics is on the verge of exhaustion. Most clinically important antibiotics originated as natural products—compounds produced by microbes—and many were discovered through systematic screening of soil organisms. While the returns to this approach are diminishing, recent innovation has shifted toward synthetic compounds and combination therapies that pair antibiotics with adjuvants targeting resistance mechanisms (Brown and Wright (2016); Lewis (2013); Figure A1). The optimistic view holds that, for every mechanism of resistance, there exists a potential countermeasure: scientific ingenuity may therefore continue to outpace bacterial evolution (Tommasi et al., 2015; Walsh and Wencewicz, 2014). Resistance may also stabilize for some pathogens, as seen with *Treponema pallidum*, the pathogen responsible for syphilis, which has remained susceptible to penicillin despite decades of widespread use (Stamm, 2010). Rather than taking a position on whether we are "running out" of ideas for new antibiotics, we limit our analysis to understanding the implications of developing one additional broad-spectrum antibiotic.

Antibiotic costs and prices. A key institutional feature of hospital antibiotic use is that drug costs are borne by hospitals rather than patients, because inpatient drug expenses are bundled into fixed per-admission payments in Medicare. Any admission for which total input costs exceed this amount generates a loss. As a result, the bundled payment acts as an upper bound on all per-admission input costs, including drugs. For example, the average Medicare payment for pneumonia (DRG 194) was about $5,084 in FY2021.[3] The broad-spectrum antibiotics that we study are used almost exclusively in hospitalized patients and are therefore subject to these bundled inpatient reimbursement rules.

This stands in contrast to outpatient drug reimbursement, where drugs are paid for separately by insurers. In Medicare, most outpatient drugs are reimbursed at their average sales price plus an add-on percentage, such that higher drug prices mechanically translate into higher provider revenue. Thus, while outpatient reimbursement incentivizes higher-priced drugs, bundled inpatient payments encourage cost containment. These institutional details limit the ability of antibiotic manufacturers to raise prices too far beyond generic antibiotic options.

From the hospital's standpoint, limiting the use of expensive antibiotics may be optimal for patients if savings are reinvested toward other areas of patient care. Although hospitals likely optimize over input costs, individual physicians are unlikely to observe or consider prices directly. Accordingly, we do not model demand as a function of prices themselves; rather, we model demand as a function of ASP policies, which in turn may depend on prices.

This suggests that policies carving out antibiotics from bundled payments may have substantial effects on overall antibiotic use; we explore the implications of this in our counterfactual simulations. More generally, the prices faced by antibiotic purchasers will remain important for determining overall antibiotic use even under lump-sum innovation payments or antibiotic subscription schemes.

[3] From the Optum 2021 DRG National Average Payment Table: https://www.optumcoding.com/upload/docs/2021 $\mathrm{DRG}_{N}ationalAveragePaymentTable_{U}pdate.pdf$.

Ongoing policy proposals targeting resistance and innovation. A number of policies have been implemented to address antibiotic innovation, including "pull" incentives like the Generating Antibiotic Incentives Now (GAIN) Act, which increased FDA exclusivity for qualified antibiotics by 5 years, and "push" incentives, such as the Combating Antibiotic-Resistant Bacteria Biopharmaceutical Accelerator (CARB-X), which provides grants and other support for antibiotic R&D. Several proposed, but not enacted, U.S. policies include the Developing an Innovative Strategy for Antimicrobial-Resistant Microorganisms (DISARM) Act and the Pioneering Antimicrobial Subscriptions to End Upsurging Resistance (PASTEUR) Act (Gregory and Martin, 2022). Both aim to revitalize the antibiotic pipeline but operate through distinct economic mechanisms. The DISARM Act would introduce Medicare add-on payments for novel antibiotics; this acts as a demand-side incentive and may increase hospitals' willingness to use novel agents. The PASTEUR Act would create subscription-style pull contracts worth up to $6 billion in total funding to reward developers of critical antibiotics independent of sales volume. While these policies would incur significant fiscal costs, the magnitude of their benefits in terms of lives saved has not been rigorously quantified. Quantifying these benefits requires modeling not only their effect on innovation, but also their impacts on resistance, demand, and ultimately health outcomes. Thus, our estimates may serve as useful benchmarks for the value of implementing these policies.

Our setting: linezolid and vancomycin resistant *Enterococcus spp.* infections. Our primary analysis focuses on linezolid, a broad-spectrum antibiotic approved by the FDA in 2000, and on *Enterococcus spp.*, a major cause of hospital-acquired infections (Wilcox, 2003). Linezolid is often reserved for the subset of *Enterococcus spp.* infections that are vancomycin resistant, so called vancomycin-resistant *Enterococcus* (VRE). Bloodstream infections with VRE are serious; historical data describe VRE mortality rates of up to 30% before the approval of VRE-active agents such as linezolid (Murray, 2000; DiazGranados et al., 2005; Stosor et al., 1998). In addition to its relevance for mortality, we focus on this drug-organism pair for three reasons: (i) linezolid is subject to active ASP restrictions, and was the focus of multiple policy changes during our sample period, whereas the next broadest spectrum antibiotic and closest substitute, vancomycin, was not subject to restrictions; (ii) *Enterococcus spp.* infections are a key clinical indication for linezolid use; and (iii) linezolid resistance, while rare, is most commonly encountered in *Enterococcus spp.*, which is important to ensure statistical power. While our empirical analyses center on gram-positive infections, the framework is readily applicable to other pathogens.

3 Conceptual Framework

The following model relates antibiotic demand, resistance, innovation, and mortality to clarify the trade-offs involved in antibiotic regulation. The model delivers several insights. First, antibiotic use generates multiple opposing externalities: higher use raises resistance but also incentivizes innovation. Second, hospitals, which do not internalize innovation and face high antibiotic prices,

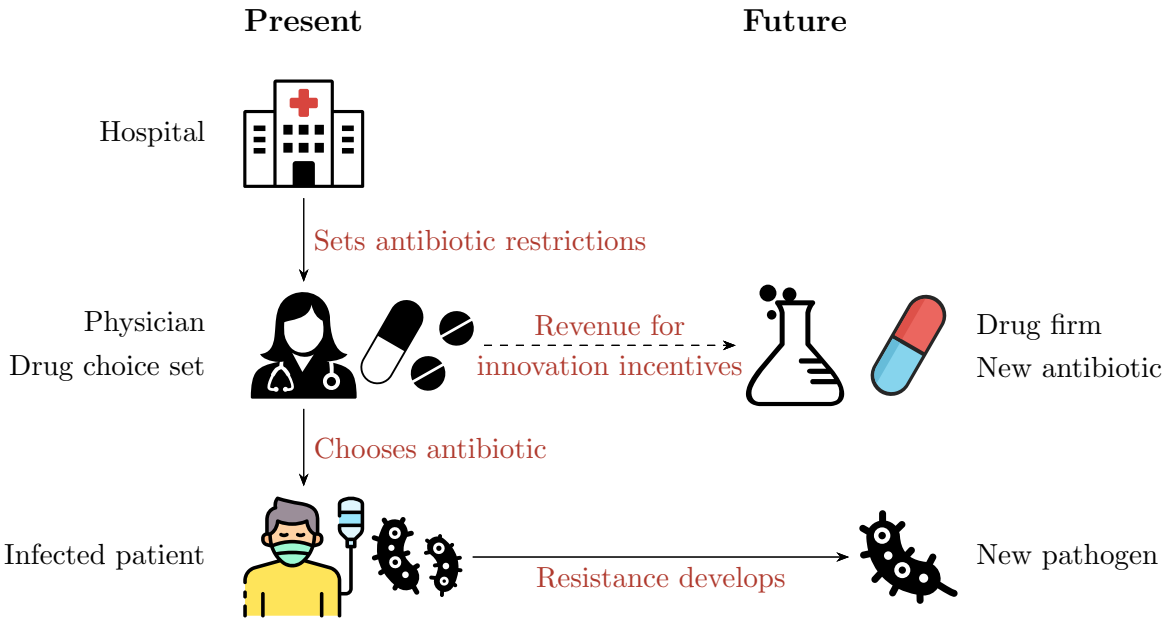
likely under-use antibiotics relative to the social optimum. Third, adding another policy lever—innovation payments—decouples use and innovation by allowing innovation incentives to be set independently of demand. However, we show that optimal use may be lower or higher with innovation payments, contrary to the intuition underlying most policy proposals that innovation payments lower antibiotic use. Fourth, the social planner can implement the optimal combination of demand and innovation policies by combining a fixed innovation payment with a subsidy that fully reimburses hospitals for antibiotic expenditures. Finally, by recognizing that resistance affects mortality by reducing antibiotic effectiveness, we derive several additional results, including a simple empirical test for suboptimal antibiotic use that depends only on the level of resistance and its elasticity with respect to demand.

In the following sections, we first present an overview of the agents and their actions. Then, we introduce a simple steady-state version of our model and graphical framework to build intuition about how optimal antibiotic use differs under hospital versus social planner objectives, with and without direct innovation payments. Finally, we derive a closed-form solution for optimal demand that highlights the underlying trade-off between marginal and inframarginal antibiotic users and yields an empirical test for suboptimal use.

Overview of Agents and Timing. We begin with an overview of the conceptual framework, summarized in Figure 1. The left column depicts how antibiotic demand and patient outcomes are determined each period. In the initial period, social planners (hospitals, medical associations, governments) set policies that affect demand, including antibiotic restrictions, taxes, or subsidies, as well as any innovation payments or prizes. Given these policies and current levels of resistance, physicians treat infected patients by choosing among available antibiotics.

The right column represents the innovation and resistance dynamics that result from these decisions. Pharmaceutical firms decide whether to enter with new antibiotics, which expand the choice set in future periods. Bacterial evolution, under selective evolutionary pressure from antibiotic use, gives rise to new, potentially resistant pathogens. Resistance then feeds back into future antibiotic choices, mortality, and innovation decisions. Together, the figure highlights how policy decisions shape both present-day antibiotic use and future innovation incentives and resistance patterns.

Figure 1: Framework linking use, resistance, and innovation



Notes: Figure shows a conceptual schematic of antibiotic use, innovation, and bacterial resistance. The left column ("Present") represents how antibiotic demand and patient outcomes are determined. From top to bottom, hospital policies affect physician antibiotic choice, which in turn affects patient outcomes. The right column ("Future") represents innovation and the emergence of potentially resistant pathogens. From top to bottom, pharmaceutical firms decide whether to enter with a novel antibiotic based on expected demand (dashed arrow). New pathogens emerge under selective pressure from prior antibiotic use in the hospital.

3.1 Setup.

To illustrate the important model results most clearly, we use a steady-state model with simplifying assumptions (we relax these assumptions in our full structural model described in Section 5). First, we assume that the social planner seeks to minimize mortality and innovation costs. We define the social welfare function as

$$W(D, P) \equiv -\lambda M(D, R(D, I(D, P))) - cI(D, P). \tag{1}$$

We assume that hospitals minimize mortality and antibiotic costs, such that the hospital objective function is

$$HW(D) \equiv -\lambda M(D, R(D)) - pD. \tag{2}$$

Each of the underlying choice variables and equations are defined as follows:

1. $D \geq 0$ represents the chosen level of steady-state antibiotic demand. Hospitals pay prices $p \geq 0$ for antibiotics, but drug revenues are considered transfers by the social planner.

2. $P \geq 0$ represents the magnitude of any direct innovation payments, which we consider transfers from the social planner perspective.[4]

3. $I(D, P) \geq 0$ represents the level of innovation R&D investment, which depends on steady-state demand D and innovation payments P. Each additional unit of R&D has marginal social cost $c \geq 0$. Importantly, hospitals consider I as exogenously determined and fixed.

4. $R(D, I) \in [0, 1]$ represents the steady-state share of antibiotic-resistant infections, and is assumed to be increasing in demand D and decreasing in innovation I.

5. $M(D, R) \geq 0$ represents steady-state mortality, which is decreasing in antibiotic use D and increasing in resistance R.

Both the social planner and hospitals are assumed to value mortality at λ dollars, and social marginal costs of antibiotic production are assumed to be zero. One simplification of our model is that we have assumed a constant stock of infected individuals, whereas in reality, aggressively treating infections may reduce their spread and reduce the number of infected individuals. This is an additional benefit of antibiotic treatment that is excluded from our model, but would tend to push in favor of expanded antibiotic use.[5]

Figure 2 illustrates the key trade-offs that determine optimal antibiotic use and the role of policy in aligning private (hospital) and social objectives. Appendix Section C reports the underlying derivations.

Without Innovation Payments, Use Trades Off Resistance and Innovation. As shown in Figure 2, Panel A, a social planner that lacks access to direct innovation payments equates the marginal costs of additional antibiotic use with marginal benefits. From the social perspective, marginal costs include only the resistance externality $\lambda \frac{\partial M}{\partial R} \frac{\partial R}{\partial D}$: higher use leads to higher resistance and thus higher mortality. Because innovation payments are not available, marginal benefits of higher use include both short-run mortality benefits and innovation incentives. Short-run mortality benefits are represented by $\lambda \left| \frac{\partial M}{\partial D} \right| \geq 0$. Innovation incentives increase welfare if the marginal benefits of R&D $\lambda \frac{\partial M}{\partial R} \left| \frac{\partial R}{\partial I} \right| \geq 0$ exceed marginal costs c. This highlights the tension between the negative externality of higher resistance and the positive externality of increased innovation.

Hospital Objectives Can Imply Under-use. As shown in Figure 2, Panel B, hospital marginal costs include antibiotic prices p, whereas hospital marginal benefits exclude the innovation externality. These two differences imply that hospitals' chosen level of antibiotic use will be lower than

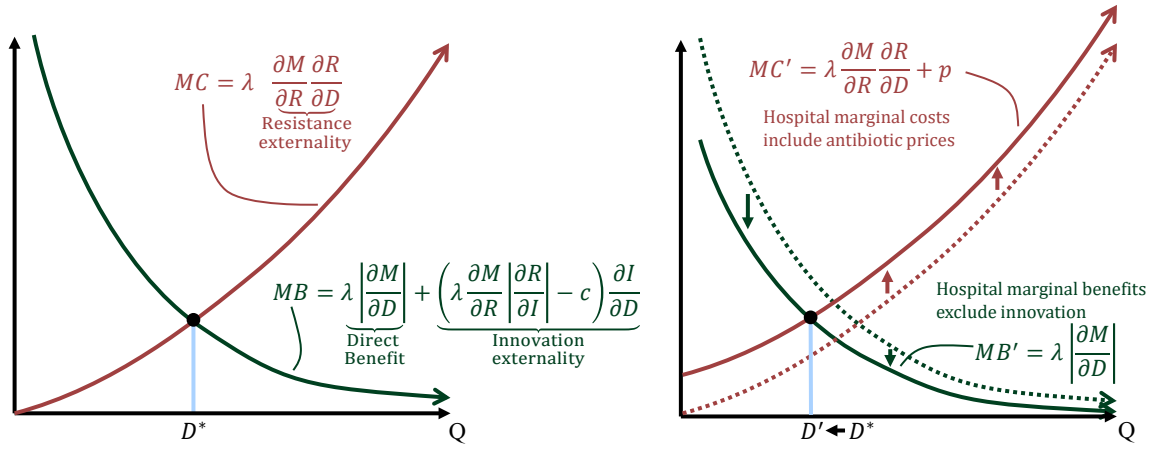
[4] Treating P as a social cost, rather than a transfer, does not change our model insights but somewhat complicates the mathematical expressions. In practice, innovation payments may take multiple forms, including grants, prizes, or R&D subsidies and may constitute varying mixtures of social costs and transfers. In our setting, the most important characteristic of innovation payments is that they offer a policy lever affecting innovation that is distinct from regulating the level of use.

[5] Moreover, given modern infection control practices in hospitals and the lower transmissibility of bacteria compared to most viruses, we expect the benefits of treatment on decreased spread to be modest.

the social optimum. This result depends on the assumption that hospitals fully internalize all resistance externalities. This is supported by our reduced-form finding that resistance remains primarily localized within institutions (Figure 5).

Figure 2: Antibiotic Use Without Innovation Payments

(a) Optimal social planner demand (b) Hospital demand



Notes: Figure shows marginal benefit and marginal cost curves for optimal antibiotic use without innovation payments. Panel A shows the social planner's objectives, where marginal benefits include both direct treatment effects and net innovation incentives. Marginal costs comprise resistance externalities. D^* denotes the socially optimal level of antibiotic use. Panel B shows hospitals' objectives. The dotted lines repeat the equilibrium in Panel A, whereas arrows and solid lines indicate hospitals' marginal benefits and marginal costs. Compared to the social planner, hospitals have lower marginal benefits because they do not consider innovation externalities. Hospitals have higher marginal costs because they internalize resistance externalities as well as antibiotic prices p. These changes lower the marginal benefit curve and raise the marginal cost curve, resulting in a lower level of demand, denoted by D'.

Innovation Payments Can Increase or Decrease Optimal Use. Figure 3 shows how socially optimal use may decrease or increase in the presence of innovation payments. We denote the optimal combination of demand and innovation payments by (D^{**}, P^{**}).

Case 1 – Innovation Payments Reduce Optimal Use (Figure 3, Panel A):

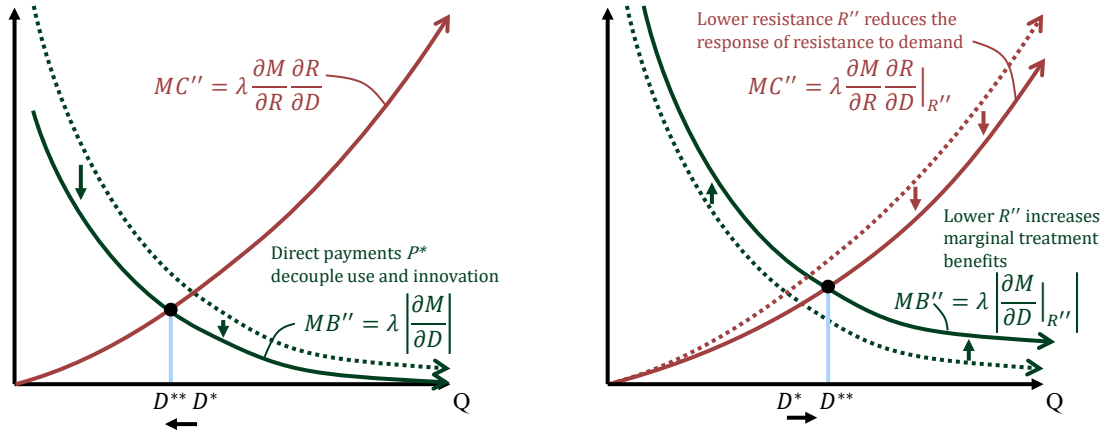
First, note that the optimal innovation payment P^{**} and optimal use D^{**} must satisfy the following pair of first-order conditions:

$$\begin{aligned} \text{FOC } P: \quad & \underbrace{\left(\lambda \frac{\partial M}{\partial R}\left|\frac{\partial R}{\partial I}\right| - c\right)\frac{\partial I}{\partial P}}_{\text{Innovation Externality}} = 0 \\ \text{FOC } D: \quad & \underbrace{\lambda\left|\frac{\partial M}{\partial D}\right|}_{\text{Direct Mortality Effect}} = \underbrace{\lambda \frac{\partial M}{\partial R}\frac{\partial R}{\partial D}}_{\text{Resistance Externality}} \end{aligned} \quad . \tag{3}$$

These conditions show that when innovation is separately optimized via the payment P^{**}, optimal use no longer considers innovation incentives on the margin. The social planner's marginal benefit expression in Figure 2, Panel A reduces to $\lambda \left|\frac{\partial M}{\partial D}\right|$. This formalizes the intuition that innovation payments can decouple use and innovation. This case is shown in Figure 3, Panel A. Here, the entire marginal benefit curve shifts downward, such that innovation payments lower optimal use. This reflects the motivation behind most current policy proposals for innovation payments: paying for innovation directly mitigates the reliance of innovation on drug sales. Note that this also implies that antibiotic use and innovation payments are strategic substitutes.

Figure 3: Optimal Antibiotic Use With Innovation Payments

(a) Case 1: Optimal Use Declines (b) Case 2: Optimal Use Increases



Notes: Figure shows marginal benefit and marginal cost curves for optimal antibiotic use, given optimal innovation payments satisfying the first equation in (3). Both panels show social planner objectives, but Dashed lines and D^* represent equilibrium social planner demand when innovation payments are zero. Solid lines and D^{**} represent social planner demand in the setting of optimal innovation payments. Panel A shows Case 1 where optimal use declines, due to the innovation externality dropping out of the marginal benefit expression. Panel B shows Case 2 where optimal use increases, due to a downward shift in resistance from R to $R^{''}$, which both increases marginal benefits and decreases marginal costs.

Case 2 – Innovation Payments Increase Optimal Use (Figure 3, Panel B):

Implementing an innovation payment $P^{**} > 0$ could also *increase* the optimal level of antibiotic use. This is because the partial derivatives shown in Figure 2 are *equilibrium* quantities that depend on the value of P^{**}. This case is shown in Figure 3, Panel B. Specifically, when innovation payments increase from \$0 to P^{**}, this causes an increase in the equilibrium level of innovation, which in turn decreases the equilibrium resistance level *at every level of demand.* Put another way, an increase in innovation from I to $I^{''}$ results in a decline in resistance from R to $R^{''}$, where $R^{''}(D) < R(D)$ for all D. This downward shift in the resistance curve has several effects. First, the marginal benefit of additional use rises via an increase in the direct treatment benefit $\left|\frac{\partial M}{\partial D}\right|$, since fewer individuals are resistant on the margin. Second, the marginal cost of use may fall if $\frac{\partial R}{\partial D}$ is smaller at lower

levels of resistance. This will be the case if resistance rates are convex in demand.[6] Figure 3, Panel B illustrates these two effects, which both push toward *higher* optimal demand. Here, antibiotic use and innovation payments are strategic complements.

Ultimately, whether innovation payments result in higher or lower optimal demand—in other words, whether innovation payments and use are substitutes or complements—is theoretically ambiguous and depends on whether the effects of Figure 3, Panel A or Figure 3, Panel B dominate. We address this empirically using our structural model estimates and counterfactual simulations. Previewing those results, we find that antibiotic use and innovation payments are complements in a large region around the status quo policies and over a wide range of model parameters.

Optimal Policies Can Be Implemented By Subsidizing Antibiotic Costs. How might the optimality conditions in (3) be implemented in practice? One solution would be for the social planner to directly set (D^{**}, P^{**}). However, this is unlikely to be feasible in practice as D^{**} could vary across hospitals and antibiotic-pathogen combinations. A more feasible solution exploits the fact that once P^{**} is set optimally, the first-order condition for D^{**} is nearly identical to the hospital's first-order condition in Figure 2, Panel B. Given our assumption that hospitals fully internalize resistance externalities, the only difference is that Panel B includes the price $p > 0$ paid by the hospital. Thus, combining optimal innovation payments P^{**} with a full subsidy for hospitals' antibiotic procurement costs would deliver the social optimum in (3), without requiring the social planner to control quantities directly.

This solution exploits the information possessed by individual hospitals about their patient populations and resistance characteristics. It also fits well with several existing policy proposals: for example, advanced market commitments and subscription models often exchange upfront payments for unit prices near production costs, and carving out antibiotics from inpatient bundled payments would shift responsibility for paying antibiotic prices from hospitals to insurers. One caveat is that if hospitals only partially internalize resistance externalities, this would call for a smaller subsidy or even a tax on use.

Optimal Antibiotic Use Represents a Marginal-Inframarginal Trade-off. Assuming a simple but reasonable functional form for mortality allows us to derive additional insights regarding the trade-off between direct treatment benefits and resistance externalities, which boils down to a classic comparison of marginal and inframarginal effects.

We assume that mortality takes the following form:

$$M(D, R) \equiv M_0 - \psi D(1 - R), \tag{4}$$

where M_0 is the (constant) baseline mortality rate without antibiotic treatment, ψ represents the (constant) mortality benefit from antibiotic treatment for share D of antibiotic users, of which

[6] This is plausible given that resistance is bounded between 0 and 1. Our structural model adopts a logit functional form for resistance. At low baseline levels of resistance, this implies a convex relationship of resistance with demand.

$(1-R)$ are non-resistant. Resistance affects mortality by rendering antibiotic treatment ineffective.

Setting aside innovation effects,[7] optimal demand D^{**} can be calculated as the first-order condition of (4) with respect to D. The marginal benefit of antibiotic use is the direct treatment effect $\left|\frac{\partial M}{\partial D}\right| = \psi(1-R) \geq 0$ and the marginal cost is the resistance externality $\frac{\partial M}{\partial R}\frac{\partial R}{\partial D} = \psi D\frac{\partial R}{\partial D} \geq 0$. Increasing use confers benefits $\psi(1-R)$ to the *marginal*, newly treated patients. Benefits are greater at low levels of baseline resistance R. Indeed, if $R \approx 0$, then the marginal users gain nearly the full benefit ψ.

Resistance costs $\psi D\frac{\partial R}{\partial D}$ are borne by the *inframarginal* patients D. Marginal costs clearly depend on the magnitude of the resistance externality $\frac{\partial R}{\partial D}$: if antibiotic use did not increase resistance, then there would be no reason to limit use. Less intuitively, marginal costs also increase with use: the cost of an additional unit of resistance is higher when more patients are relying on the antibiotic in the first place. One implication is that raising antibiotic use is most favorable when baseline levels of both resistance and use are low.

An Empirical Test for Suboptimal Antibiotic Use. Equating marginal benefits to marginal costs, canceling ψ, and re-arranging yields the following condition for optimal demand:

$$D^{**} = \frac{1-R}{\frac{\partial R}{\partial D}} \iff \varepsilon_{R,D} = \frac{1-R}{R}, \tag{5}$$

where $\varepsilon_{R,D}$ is the elasticity of resistance with respect to use, and $\frac{1-R}{R}$ is the inverse odds of resistance, or equivalently, the odds of non-resistance. This simple expression yields several insights. First, the optimal level of demand does not depend on the magnitude of the mortality effect ψ. The intuition is that antibiotic resistance only affects mortality to the extent that the antibiotic confers any survival advantage to non-resistant patients. This theoretical result is confirmed by the results from our structural model in Section 6.4: scaling the mortality effect matters for welfare but does not significantly affect the optimal level of demand.

Second, Equation 5 yields a simple empirical test for whether an antibiotic is underused or overused, with $\varepsilon_{R,D} < \frac{1-R}{R}$ implying underuse. For our antibiotic of interest linezolid, $\varepsilon_{R,D} \approx 4.0$ and $\frac{1-R}{R} \approx 82.3$, suggesting that raising linezolid use could improve welfare, a result that is consistent with our structural analysis of counterfactual policies in Section 6.[8]

We acknowledge that this simple framework omits important factors that affect optimal demand. For example, it may be the case that increasing resistance to one antibiotic increases resistance to others. That said, we do not find any evidence of this effect in our setting. Additionally, antibiotic benefits ψ may vary across individuals and may be greater for inframarginal patients. This would lower optimal demand relative to (5). While the toy model in this section is useful for building intuition, we turn to our structural model in Section 5 to account for treatment effect heterogeneity

[7] This could reflect a social planner that ignores innovation or implements a separate optimal innovation prize (Figure 3). This could also reflect hospital incentives, which ignore innovation, with fully subsidized antibiotic prices.

[8] Our underlying calculations are as follows. From Figures 4 and Figure 5, $\frac{\partial R}{\partial D} \approx 0.015/0.0087 = 1.7$. From Tables A1 and A4, $R \approx 1.2\%$ and $D \approx 2.8\%$. Hence, $\varepsilon_{R,D} \equiv \frac{\partial R}{\partial D}\frac{D}{R} \approx 4.0$ and $\frac{1-R}{R} \approx 82.3$.

and to evaluate counterfactual policies.

4 Data

This section describes the data sets we use to generate reduced-form evidence and to estimate the full structural model.

4.1 Electronic Health Records Data

Our EHR data form a panel dataset of admissions to two large academic medical centers in the same metropolitan area from 2008 through 2019.[9] The sample includes all admissions during which a bacterial culture was drawn.

Inpatient antibiotic use. Inpatient antibiotic use is not typically recorded in insurance or Medicare claims data, because billing is bundled for most hospitalizations and antibiotic administration is not itemized. In contrast, our EHR data include admission-level information on which inpatient antibiotics were administered, in addition to rich admission-level covariates, including age, sex, demographics, and admission-level characteristics (diagnoses,[10] prior admissions, prior antibiotics, prior positive cultures,[11] etc.) in the preceding 12 months.

Antibiotic resistance. Similarly, individual-level data on antibiotic resistance are not available in standard claims datasets.[12] Our EHR include microbiology reports for each admission, detailing the bacteria isolated and susceptibilities of those isolates to a panel of antibiotics. Importantly, cultures are drawn before antibiotics are administered, reducing concerns that resistance testing is conditional on antibiotic selection. Moreover, susceptibility panels test for resistance against a standard set of antibiotics, not just those ultimately prescribed. This allows us to observe counterfactual resistance outcomes even for drugs not given to the patient. This constitutes a major advantage compared to other types of treatment, where counterfactual treatment responses for a given individual are difficult to observe. For example, in the antibiotic context, a patient may have received antibiotic A, but their susceptibility report could reveal that antibiotics B and C would also have treated the infection.

[9] Although both hospitals are in the same city, they serve largely non-overlapping patient populations – only a small minority of patients are ever observed to have admissions in both hospitals across our sample window.

[10] Diagnoses include cancer, cystic fibrosis, catheterization, pneumonia, fever, abdominal-related diagnoses, post-operative infection, urinary tract infection, bacteremia/sepsis, cellulitis/abscess, heart disease, end-stage renal disease, and transplant.

[11] Microbiology results included whether a sample was positive and for what pathogen (e.g., prior culture positivity for *Enterococcus* spp. versus *Staphylococcus aureus*), as well as specimen type, which included abdominal, blood, tube/line, orthopedic, rectal, skin and soft tissue, urine, wound, and other.

[12] While some ICD codes for antibiotic resistance exist, these do not distinguish between different pathogen species or antibiotics. Also, like utilization, these codes may not be reliably recorded.

Health outcomes and pager request data. Our main health outcome of interest is within-admission 30-day mortality. To identify effects of antibiotic use on mortality, we leverage internal administrative data on antibiotic requests made via a pager system. The pager request data is available for Hospital 1 from 2011Q4-2016Q4. For each request, we observe the date, antibiotic requested, and identity of the pager holder, typically an infectious disease fellow or pharmacist, making the decision. We use these data to reconstruct the weekly shift schedule of pager holders.

4.2 Other Datasets

Antibiotic prices and restriction status. We obtained data on the cost-per-day of each antibiotic at Hospital 1. New entrant prices are assumed to equal the price of linezolid in 2008, midway between its approval in 2000 and the first generic entrant in 2015. All prices, revenues, and costs are reported in 2024 dollars.

Antibiotic FDA approvals and revenues. Our innovation model combines data on antibiotic entry and revenues. To measure new antibiotic entry, we obtain antibiotic approvals from the FDA *Orange Book* from 1996 through 2019 (Center for Drug Evaluation and Research, 2025). Figure A1 illustrates the sequence of major antibiotic approvals over our study period, separating drugs that primarily target gram-positive versus gram-negative bacteria. This timeline highlights the relatively limited number of new antibiotic entrants. We obtained data on annual U.S. revenues of approved antibiotics from *Evaluate Pharma*. These data are derived from 10-K filings by public pharmaceutical companies to the SEC and other sources. Importantly, the revenue data are net of manufacturer rebates. We assume that U.S. revenues account for 55% of global revenues (Ho and Pakes, 2024).

5 Structural Model

5.1 Motivation for Estimating a Structural Model

The conceptual framework presented in Section 3 provides intuition for understanding the relationships between antibiotic use, resistance, mortality, and innovation. To analyze the impacts of counterfactual policies, we implement a structural model of demand, resistance, and mortality estimated on individual-level data.

The structural model extends the basic setup introduced in Section 3 to model resistance and demand for the full choice set of antibiotics for gram-positive infections (comprising 5 antibiotic alternatives and 1 outside option). We focus our attention on utilization and resistance of linezolid, the highest-tier antibiotic in the choice set and the one most directly affected by our identification strategies.[13] We quantify welfare in terms of mortality, measured as the number of admissions

[13]The only other restricted gram-positive antibiotic is daptomycin, which is in the same tier as linezolid but is more rarely used and lacks resistance data for most of the sample period. Given its rarity, we exclude daptomycin from our demand and resistance models and drop a small number of admissions where daptomycin was the highest-tier

resulting in death within 30 days, scaled to the level of the U.S. based on the share of all U.S. inpatient beds accounted for by the treatment hospital. The effect of antibiotic use on mortality is identified via an instrumental variables design based on the identity of the provider approving restricted antibiotic requests each week.

The estimated model allows us to simulate forward demand, resistance, and mortality over time as a function of antibiotic stewardship restrictions on linezolid. Innovation takes the form of a single potential entrant that decides whether to enter based on expected future revenues.[14] Although our empirical focus on gram-positive infections and linezolid limits external validity, our goal is to develop a structural framework that connects treatment decisions (demand), health outcomes (resistance and mortality), and innovation (antibiotic entry).

5.2 Model Equations

Estimating the structural model amounts to estimating the following four separate equations describing antibiotic demand, resistance, mortality, and innovation.

1. Drug demand D_{ijt} models whether drug j is chosen for admission i at time t as a function of stewardship policies, resistance, and patient characteristics. Identification leverages the introduction of two antibiotic restriction policies (introduced in the treatment but not control hospital).

2. Resistance R_{ijt} depends on the prior year's hospital-level demand (capturing the externality of antibiotic use on hospital-level resistance), resistance (accounting for the potential of resistance to self-persist over time even in the absence of antibiotic use), and individual-level medical histories (including prior antibiotic use and resistance). Identification of the resistance externality leverages the same two policy changes used to identify the effect of antibiotic restrictions on demand.

3. Mortality M_{it} is determined by patient characteristics, the antibiotics administered, and the patient's resistance profile. The effect of antibiotic use on mortality is identified using a judges design that exploits variation in approval rates among stewardship pager holders, who rotate through this role in week-long shifts.

4. Innovation, I_{lt}, represents whether new antibiotic l enters the market at time t, which occurs when expected discounted profits exceed sunk costs of innovation.[15] The innovation model combines data on net-of-rebate antibiotic revenues with data on FDA antibiotic approvals.

Figure A3 summarizes the relationships between these equations. In the following subsections, we report our identification strategy, reduced-form evidence, and structural estimates for each

antibiotic. Including daptomycin in our models does not change our qualitative results.

[14] We assume that the entrant has the same drug characteristics as linezolid, including drug price, but no use or resistance history. Entrants are added to the choice set with initial resistance set to zero.

[15] The total entry cost includes the cost of failed development programs and clinical trials necessary to yield one FDA-approved antibiotic.

equation, starting with demand and followed by resistance, mortality, and innovation.

5.3 Demand

Identification Leveraging Changes in Antibiotic Restriction Policies. We use hospital-specific antibiotic stewardship program (ASP) policy changes to identify the causal effects of antibiotic restrictions on both antibiotic use and bacterial resistance, focusing on the antibiotic linezolid. During our sample period, two policy changes were implemented at the treatment hospital (Hospital 1), both aimed at reducing inappropriate antibiotic use. The control hospital (Hospital 2) did not adopt comparable ASP restrictions during this period, allowing us to isolate the effect of these policies using quasi-experimental methods.

1. The first policy, implemented in 2011Q4, formalized the prior authorization process for linezolid and other restricted antibiotics shown in Figure 6. Specifically, this policy created a dedicated, paid role for an infectious disease pharmacist or physician; previously, antibiotic approvals were handled part-time by first-year infectious disease fellows.

2. The second policy, implemented in 2014Q2, issued new linezolid-specific ordering guidelines that defined narrower clinical indications for its use. For example, the new guidelines recommended against empiric[16] linezolid use in almost all circumstances.

We evaluate the impact of these policies on (i) the utilization of linezolid and (ii) the prevalence of linezolid resistance. To do so, we estimate difference-in-differences (DiD) and triple-differences (DDD) specifications, paired with event-study specifications to assess dynamics. Appendix Section D.1.1 reports the underlying estimation specifications. We then embed this identifying variation in structural models of demand and resistance.

Reduced-form Evidence – Antibiotic Restriction Policies Decrease Use. Figure 4 illustrates our identification strategy by plotting linezolid use over time. We focus on linezolid because it is the only restricted antibiotic in our choice set and thus the only antibiotic directly affected by the policy changes. Panel A plots the mean share of admissions with any linezolid use in each hospital, aggregated to the half-year level.[17] The vertical dashed lines denote the introduction of each policy change. Prior to the first policy, trends in use are similar across hospitals. Immediately following policy 1, linezolid use at Hospital 1 declines sharply, while utilization at Hospital 2 remains stable. Our difference-in-differences (DiD) estimate implies a reduction of 0.009 percentage points in the share of admissions treated with linezolid after both policies (SE = 0.0016, $P < 0.001$), equivalent to about a 20 percent decline relative to baseline use.

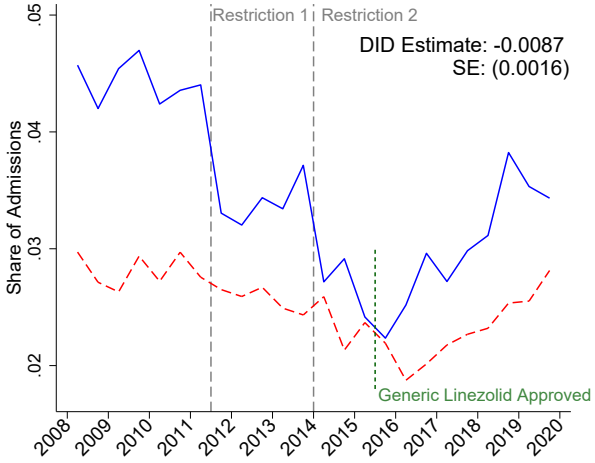
[16] Empiric use is defined as administering antibiotics for a suspected infection based on clinical signs, likely pathogens for a given infection site, local resistance patterns, and patient risk factors, aiming to prevent serious harm while waiting for lab results. Targeted therapy refers to antibiotic selection once the pathogen is known.

[17] Data are aligned such that the quarters prior to Policy 1 (2011Q2 and 2011Q3) are binned together, and likewise for the quarters prior to Policy 2 (2013Q4 and 2014Q1).

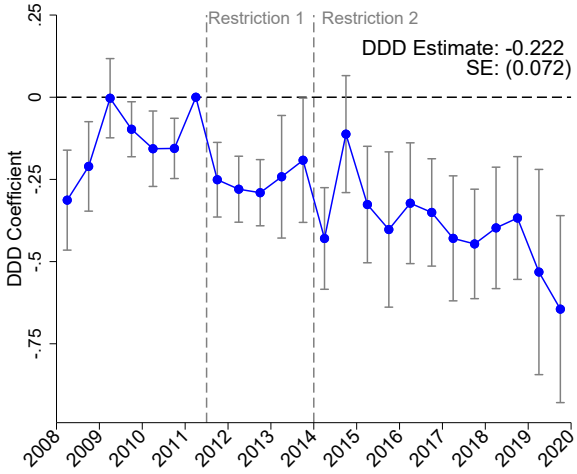
Figure 4, Panel B reports event-study estimates from a triple-difference (DDD) specification where the outcome is the log utilization rate at the hospital–drug–half-year level. This specification compares changes in linezolid use at Hospital 1 relative to Hospital 2 and relative to other antibiotics not subject to restriction. The estimated coefficients trace dynamic treatment effects surrounding the policy dates. Although noisy, the pre-period point estimates suggest a flat or somewhat increasing pre-trend. Following implementation, coefficients become significantly negative and decline over time, indicating a sustained reduction in antibiotic use. The average post-policy effect is –0.22 (SE = 0.07, $P = 0.004$), consistent with an approximately 25 percent decline. These results confirm that stewardship restrictions produced a large and persistent reduction in linezolid use, above and beyond contemporaneous trends in control antibiotics or hospitals.

Figure 4: Antibiotic Restrictions Lower Use

(a) Utilization in Hospital 1 vs Hospital 2



(b) Triple-difference Specification



Notes: Figure shows the reduced form relationship between antibiotic use and two antibiotic restriction policies (dashed vertical lines) implemented by the treatment hospital (Hospital 1). Panel A shows the share of admissions where the highest-tier antibiotic used was linezolid, separately for the treatment and control hospitals. Linezolid use sharply declines after each antibiotic restriction in the treatment but not control hospital. Both hospitals show increasing use after generic versions of linezolid were first approved in July 2015. Panel B shows coefficient estimates from the triple-difference event study specification in Appendix Section D.1.1, where the outcome variable is the log utilization share at the hospital-drug-half year level. For both panels, data are aggregated to the half-year level such that the two quarters prior to Policy 1 (2011Q2 and 2011Q3) are combined. The singleton quarters 2008Q1 and 2019Q4 are omitted from the plots but are included in the reported aggregate estimates.

Structural Model – Setup and Estimation. We embed the above policy-driven identifying variation in a conditional-logit model of demand. Patients indexed by i are hospitalized in year-month t. Their treating physicians choose antibiotic alternatives j from a choice set of five antibiotics, plus an outside option of no antibiotic, based on patients' observable characteristics and

any ASP policies in effect.[18] We divide antibiotics into tiers based on how they are used clinically, with higher-tier antibiotics reserved for more resistant pathogens (tier assignments are shown in Table A1).

We model only the highest-tier antibiotic administered during each admission.[19] Linezolid, the highest-tier antibiotic in the choice set, may be used early on as empiric treatment to cover for VRE before culture results are available, or patients may be escalated to linezolid if a culture returns resistant to lower-tier antibiotics. Patients may also be de-escalated to lower-tier antibiotics if culture results show susceptibility to lower-tier antibiotics. Modeling the highest-tier antibiotic administered accounts for all of these scenarios.

We assume that demand follows a conditional logit form (McFadden, 1974) for admission i, drug j, and year-month t:

$$U_{ijt} = \theta_j A_{hjt} + \tilde{\mathbf{u}}'_{\mathbf{ijt}} \beta^{\mathbf{D}}_{\mathbf{u,j}} + \tilde{\mathbf{r}}'_{\mathbf{ijt}} \beta^{\mathbf{D}}_{\mathbf{r,j}} + \mathbf{W}'_{\mathbf{it}} \delta_j + \varepsilon^D_{ijt}. \tag{6}$$

Choice utility U_{ijt} depends on the current ASP policy A_{hjt} for hospital h, antibiotic j, and year-month t. The effect of ASP policies on demand is captured by θ_j, which is allowed to vary by drug. We assume physicians do not directly consider antibiotic prices or *future* resistance when making treatment decisions; we assume that these considerations are captured by hospitals' ASP policies via $\theta_j A_{hjt}$. The $\tilde{\mathbf{u}}'_{\mathbf{ijt}} \beta^{\mathbf{D}}_{\mathbf{u,j}}$ term controls for the effects of past antibiotic use, which we allow to vary by drug j.[20] The $\tilde{\mathbf{r}}'_{\mathbf{ijt}} \beta^{\mathbf{D}}_{\mathbf{r,j}}$ term captures physicians' beliefs about resistance profiles before susceptibility results are available, captured by the predicted resistance probabilities from (7). Additionally, we include a rich set of admission-level controls $\mathbf{W}'_{\mathbf{it}}$, which incorporate drug fixed effects and their interactions with hospital, demographics, prior admissions, past ICD diagnosis codes, past positive cultures by specimen type and bacterial species, current-admission diagnosis codes, and current-admission culture draws. The error term ε^u_{ijt} is assumed to follow the standard Type I extreme value distribution.

Descriptive Statistics for Antibiotic Choices and Patient Characteristics. The demand estimation sample includes 386,498 unique hospital admissions with 5 choice alternatives each plus the outside option.

Table A1 shows characteristics for each antibiotic. At least one of the 5 antibiotics is chosen

[18] The choice set reflects antibiotics used to treat gram-positive infections and is identical to that of the resistance model (i.e., first-generation cephalosporins, clindamycin, trim/sulfa, vancomycin, and linezolid). Admissions where the outside option is chosen include cases where a different antibiotic outside of the choice set was administered. While our choice set is limited to antibiotics for gram-positive infections, patients may simultaneously receive antibiotics to treat other infections from gram-negative or anaerobic pathogens. Other categories of antimicrobial drugs include antibiotics for gram-negative, anaerobic, or "atypical" bacteria, or fungal pathogens. Choices for those categories can be considered independently from the choice of which anti-gram-positive antibiotic to use (if any).

[19] Patients may trial multiple antibiotics given initial uncertainty about the underlying pathogen. When multiple administered antibiotics occupy the same tier, we take the antibiotic with the latest administration date.

[20] The vector $\tilde{\mathbf{u}}'_{\mathbf{ijt}}$ includes past use of drug j, past use of higher-tier antibiotics, and past use of lower-tier antibiotics. When the set of higher- or lower-tier antibiotics contains multiple drugs, we take the average. For example, for patient i and alternative j = linezolid, the set of lower-tier antibiotics includes the four other antibiotics. If this patient only used one lower-tier antibiotic previously, then the lower-tier prior use variable would take value 1/4.

58% of the time, with linezolid chosen in 2.8% of admissions. Vancomycin, the most commonly chosen antibiotic (36.4% of admissions), is an IV infusion that is often used as the empiric (initial) treatment when the culprit pathogen is not yet known; it covers almost all gram-positive organisms except for VRE. Importantly, linezolid covers VRE in addition to the other pathogens covered by vancomycin. Prior use within the past 12 months is highest for vancomycin (14.9%) followed by cephalosporins (14.8%), which are the lowest tier of antibiotics in the choice set. Prior use of linezolid is rare (1.8%) reflecting its status as a restricted antibiotic. Predicted resistance ranges from 0.7% for linezolid to 48% for cephalosporins. Interestingly, predicted resistance is inversely related to antibiotic tier but is not clearly correlated with use.

Table A2 shows descriptive statistics for admission-level covariates including patient demographics, ICD diagnosis codes, and prior hospital admissions. Reassuringly for our identification strategy, these admission and patient characteristics are similar between both hospitals.

Demand Model Estimates. Table 1 presents the estimated coefficients from the demand model. The first set of coefficients shows that ASP policies significantly lowered demand for linezolid ($\hat{\theta}_{linezolid} = -0.37$, SE = 0.04) and increased demand for other antibiotics in the choice set, including vancomycin and lower-tier antibiotics, all relative to the outside option. Our estimate of $\hat{\theta}_{linezolid}$ implies that restriction policies reduced linezolid use by 31%, which is reassuringly similar in magnitude to our reduced-form results shown in Figure 4.

Table 1: Demand Model Estimates

	Linezolid	Vancomycin	Lower-tier
Policy Effect by Drug (θ_j)	-0.374**	0.167**	0.089**
	(0.037)	(0.018)	(0.017)
	Linezolid (self)	Vancomycin and Lower-tier	
Predicted Resistance ($\beta^D_{r,linezolid}$)	-1.291	3.209**	
	(0.876)	(0.312)	
Prior Use ($\beta^D_{u,linezolid}$)	1.745**	1.234**	
	(0.055)	(0.093)	

$^{**}p < 0.01$, $^{*}p < 0.05$.

Notes: Table shows conditional logit estimates of (6). Row 1 reports effects of the two antibiotic restriction policies on demand for linezolid, vancomycin, and lower-tier antibiotics, corresponding to $\theta_j A_{hjt}$. Rows 2 and 3 report effects of predicted resistance ($\beta^D_{r,j}$) and prior use ($\beta^D_{u,j}$) on demand for j = linezolid. The model also includes controls for admission-level covariates (ICD diagnosis codes, prior antibiotic use, prior cultures, and demographics), which are allowed to vary freely by antibiotic. Appendix Table A3 reports the full set of estimated coefficients. The estimation sample includes 2,318,988 admission-choice alternative pairs across 386,498 unique admissions. Standard errors (in parentheses) are clustered at the patient level.

The next set of coefficients indicates that predicted linezolid resistance lowers its demand, although

this is not statistically significant ($\hat{\beta} = -1.29$, SE = 0.88), consistent with linezolid's status as a last-line antibiotic for which few substitutes exist. In contrast, predicted resistance to lower-tier antibiotics is strongly and significantly associated with increased linezolid use ($\hat{\beta} = 3.21$, SE = 0.31). This pattern suggests that rising resistance to lower-tier antibiotics pushes clinicians toward prescribing higher-tier agents such as linezolid. Finally, prior use of linezolid strongly predicts subsequent linezolid choice ($\hat{\beta} = 1.75$, SE = 0.06), indicating substantial persistence in treatment patterns. Prior use of lower-tier antibiotics is also associated with higher linezolid demand ($\hat{\beta} = 1.23$, SE = 0.09). Together, these findings highlight strong state dependence in antibiotic prescribing and confirm that linezolid demand responds to ASP policy restrictions.[21] Coefficients corresponding to other antibiotics and drug-by-patient characteristic interactions are reported in Appendix Table A3.

5.4 Resistance

Setup: Individual- and Hospital-level Determinants of Resistance. Among positive bacterial cultures with antibiotic susceptibility results, we model the binary outcome of whether culture $c(i,t)$ for individual i obtained in year-month t is resistant to antibiotic j using a logit functional form. This framework has several advantages. First, it allows each culture to be resistant to a single antibiotic, multiple antibiotics, or none. Second, it implies convexity in resistance probabilities for low baseline levels of resistance. Finally, it facilitates modeling resistance at both the individual and hospital levels. We model the resistance status of each bacterial isolate as the following function of individual- and hospital-level factors:

$$R_{cj} = \tilde{\mathbf{r}}'_{\mathbf{cj}}\beta^{\mathbf{r}}_{\mathbf{j}} + \tilde{\mathbf{u}}'_{\mathbf{cj}}\beta^{\mathbf{u}}_{\mathbf{j}} + \mathbf{H}'_{\mathbf{c}}\omega_{\mathbf{j}} + \alpha_{hjy} + \varepsilon_{cj} \tag{7}$$

We define $\tilde{\mathbf{r}}_{\mathbf{cj}}$ as a vector that includes indicator variables for whether individual i had a prior resistant culture to drug j in the past 12 months, as well as i's prior resistance to higher- and lower-tier antibiotics. Analogously, $\tilde{\mathbf{u}}_{\mathbf{cj}}$ includes indicator variables for whether i used, in the past 12 months, antibiotic j, any higher-tier antibiotics, and any lower-tier antibiotics. The vector $\mathbf{H}'_{\mathbf{c}}$ includes a rich set of control variables for i's medical history in the last 12 months, including indicator variables for past hospitalization, prior ICD codes, prior positive cultures, and demographics. The full set of control variables is described above in Section 4. The coefficients on $\tilde{\mathbf{r}}_{\mathbf{cj}}$, $\tilde{\mathbf{u}}_{\mathbf{cj}}$, and $\mathbf{H}'_{\mathbf{c}}$ are allowed to vary flexibly by antibiotic. The term ε_{cj} denotes the *i.i.d.* logit error.

The α_{hjy} terms denote hospital-level effects, which capture the resistance externalities from antibiotic use as well as the persistence of resistance within a hospital over time. We model these effects using the following modified AR(1) process:

[21] The effects of predicted resistance and prior antibiotic use may reflect either *state dependence*—for example, inertia whereby exogenous changes to resistance or past use cause changes in demand—or *unobserved heterogeneity*, where estimated effects of past use capture persistent patient-level differences rather than causal effects. In our counterfactual simulations, we assume that past use effects reflect true state dependence, while effects of predicted resistance capture unobserved heterogeneity.

$$\alpha_{hjy} = \delta\alpha_{hjy-1} + \gamma\tilde{u}_{hjy-1} + a_{hj} + \eta_{hjy}, \tag{8}$$

where the $\delta\alpha_{hjy-1}$ term captures the dependency of this year's resistance on last year's resistance, with the decay factor δ. This allows shocks to resistance to persist, reflecting the biological reality that resistant strains of bacteria may persist in the hospital environment.[22] The $\tilde{u}_{hjy-1}$ term denotes the utilization share of drug j in hospital h in year $y-1$, with γ representing the resistance externality of antibiotic use. The a_{hj} term represents hospital-drug fixed effects and adjusts for fixed differences in resistance across different drug-hospital pairs.[23] Lastly, η_{hjy} is the error term.

Identification of the Use-Resistance Elasticity. At the individual-level (Equation 7), the coefficients of interest $\beta^{\mathbf{r}}_{\mathbf{j}}$, $\beta^{\mathbf{u}}_{\mathbf{j}}$, and $\omega_{\mathbf{j}}$ are identified using variation in patients' medical histories conditional on the drug-hospital-year fixed effects. At the hospital-level (Equation 8), the main identification challenge is that hospital antibiotic utilization $\tilde{u}_{hjy}$ may be correlated with omitted factors that affect resistance contained in the error term η_{hjy}. For example, changes in the patient population may increase both the share of patients needing antibiotics as well as the amount of antibiotic resistance. Moreover, hospital infection control practices may affect both resistance and use. To identify the hospital-level utilization coefficient γ, we leverage the same hospital-specific changes in antibiotic restriction policies introduced in Section 5.3.

Reduced-Form Evidence – Antibiotic Restriction Policies Lower Resistance. Figure 5 illustrates how the two antibiotic restriction policy changes affect linezolid resistance. Panel (a) plots the share of *Enterococcus spp.* isolates resistant to linezolid in each hospital over time. Prior to 2013, resistance rates trended similarly across hospitals. One year after the restrictions took effect, Hospital 1 experienced a pronounced decline in resistance, while resistance in Hospital 2 remained flat. The DiD estimate suggests a –0.015 percentage-point reduction in the share of resistant isolates (SE = 0.003, $P < 0.001$), a roughly 60 percent drop relative to the pre-policy mean. Panel (b) presents corresponding DDD event-study coefficients, estimated at the hospital–drug–pathogen–half-year level, where the outcome is the log share of resistant isolates. As with use, coefficients are flat or somewhat increasing prior to the restrictions, but fall sharply thereafter. The post-policy average effect equals –1.01 (SE = 0.09, $P < 0.001$), implying a 64 percent decline in resistance. The magnitude and persistence of these effects mirror those observed for utilization, consistent with a causal relationship whereby tighter antibiotic restrictions reduced use and subsequently lowered resistance rates.

The timing of the decline in resistance—lagging use by roughly one year—is consistent with a high biological fitness cost of resistance, whereby resistant bacterial strains are out-competed

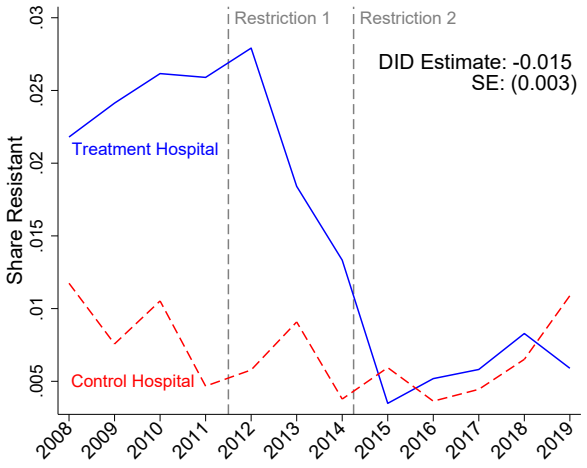
[22] As discussed in Section 2, persistence occurs when the fitness cost of developing resistance is low. In contrast, if resistance mutations confer significant fitness costs, resistance strains will not tend to persist in the absence of active antibiotic use, because they will be outcompeted by non-resistant strains.

[23] For example, for a given drug, utilization or resistance may be persistently higher in one hospital compared to the other. The hospital-drug fixed effects account for these differences and ensure that the identification of the utilization effect γ is driven only by changes over time occurring within each drug-by-hospital combination.

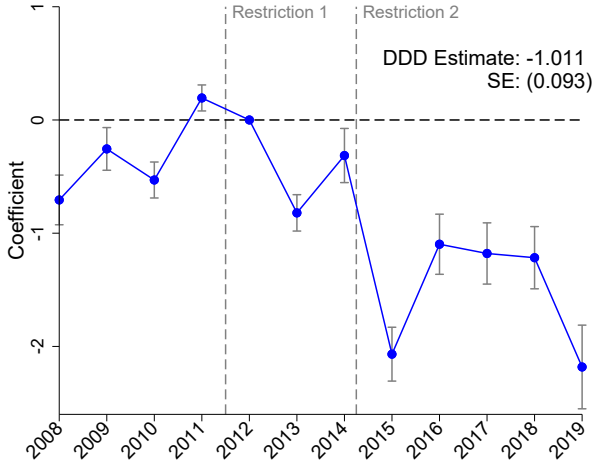
by non-resistant strains in the absence of evolutionary pressure from antibiotic use. Although we find some evidence that use and resistance was rising in Hospital 1 prior to the first policy, the sharp changes we observe after each policy lend support to a causal interpretation. Moreover, an increasing pre-trend would tend to bias our estimates toward zero. The findings imply substantial resistance externalities from antibiotic use, validating the efficacy of antibiotic restrictions in controlling resistance.

Figure 5: Antibiotic Restrictions Lower Resistance

(a) Resistance in Hospital 1 vs Hospital 2



(b) Triple-difference Specification



Notes: Figure shows the reduced form relationship between resistance and two antibiotic restriction policies implemented by Hospital 1 (shown in the vertical dashed lines). Panel (a) shows the share of *Enterococcus spp.* isolates resistant to linezolid in Hospital 1 and Hospital 2 over time. Linezolid resistance in Hospital 1 sharply declines one year after antibiotic restrictions are introduced, while resistance levels in Hospital 2 remain stable; text shows DiD estimates treating 2013 as the first year affected by the policies, with robust standard errors. Panel (b) shows estimates from a triple-difference event study specification where the outcome variable is log resistance at the hospital-drug-pathogen-year level. Here, we include an additional comparison of linezolid to other drugs (cephalosporins, clindamycin, trimethoprim/sulfamethoxazole, and vancomycin). Here, standard errors are clustered at the hospital-drug-pathogen level.

Structural Estimation. We use a two-step estimation approach for estimating (8). In step 1, we estimate the patient-level model in (7) using separate logistic regressions for each antibiotic j. In step 2, we estimate the hospital-level model in (8). A key challenge is that the AR(1) term is mechanically correlated with the drug-specific component of the error term (a_{hj}). This is a well-known issue with dynamic panel data models that leads to inconsistent estimates of δ (Roodman, 2009; Blundell and Bond, 1998; Arellano and Bover, 1995).To address this, we apply consistent estimators from the dynamic panel data model literature. Because our identification of γ depends only on between-hospital differences in utilization change over time, we use between-hospital differences as inputs to the estimators, akin to a two-by-two difference-in-differences design. Starting from Equation 8, we compute differences in all variables between Hospital 1 (treatment)

and Hospital 2 (control), yielding the following specification:

$$\alpha_{jy} = \delta\alpha_{jy-1} + \gamma\tilde{u}_{jy} + a_j + \eta_{jy}, \tag{9}$$

where $\alpha_{jy} \equiv \alpha_{1jy} - \alpha_{2jy}$ and likewise for the other variables. To ensure that γ is properly identified, we replace $\tilde{u}_{jy}$ in (9) by its predicted value from the following first-stage OLS regression:

$$\tilde{u}_{jy} = \theta\tilde{A}_{jy} + b_j + \varepsilon_{jy}, \tag{10}$$

where $\tilde{A}_{hjt}$ encodes ASP policy changes, b_j is a drug-specific fixed effect, and ε_{jy} is the error term.[24] In our preferred specification, we estimate (9) using the first-difference GMM estimator (Arellano and Bond, 1991). We confirm that our results are robust to instead using the system GMM estimator of Blundell and Bond (1998) and Arellano and Bover (1995). We also estimate a standard two-stage least squares (2SLS) version of (8) that replaces the AR(1) term with drug-year fixed effects to flexibly allow for drug-specific time trends.[25]

Descriptive Statistics for Bacterial Isolates. Table A4 shows descriptive statistics for the positive bacterial isolates in our resistance dataset, separately by antibiotic.[26] Each isolate is tested for resistance to a panel of antibiotics that potentially cover that pathogen.[27] Resistance rates to linezolid averaged 1.2% overall across both hospitals. Of isolates tested for linezolid resistance, 0.6% had a prior positive culture within the last 12 months that tested resistant to linezolid, whereas 12% had past resistance to a lower-tier antibiotic. Of isolates tested for linezolid resistance, linezolid was previously used in 11% of cases, and lower-tier antibiotics were previously used in 23% of cases.[28]

Individual-level Resistance Model Estimates. Table 2 shows estimates of the individual-level resistance model for linezolid and, for comparison, the average coefficients across all antibiotics. At the individual level, resistance to linezolid is highly persistent: patients with linezolid resistance in the past 12 months have a significantly higher likelihood of current resistance, with an estimated coefficient of 2.47 ($P < 0.01$). Prior use of linezolid is also significantly associated with current linezolid resistance (estimated coefficient = 1.61, $P < 0.01$). These effects are smaller in magnitude but remain statistically significant for non-linezolid antibiotics. This suggests that antibiotic use has significant internalities on patients' own future resistance, especially for linezolid, potentially because it is a more rarely used antibiotic.

Overall, we find limited evidence for cross-pathogen resistance. On average, resistance seems to only be significantly correlated with prior use of higher-tier antibiotics, with a modest mean coefficient of 0.18. We do not find significant associations with prior resistance to higher-tier antibiotics

[24] For the antibiotic of interest linezolid, $\tilde{A}_{hjt}$ takes value 1 after both antibiotic restriction policies, 0.5 between the first and second restriction policies, and 0 prior to the first restriction policy. For all other antibiotics, $\tilde{A}_{hjt} = 0$.

Table 2: Individual-level Resistance Model Estimates

	Linezolid coefficient	Average coefficient across drugs
Prior resistance to antibiotic j	2.47**	1.66**
	(0.43)	(0.04)
Prior resistance to higher-tier antibiotic	—	0.17
		(0.12)
Prior resistance to lower-tier antibiotic	-0.02	0.05
	(0.54)	(0.04)
Prior use of antibiotic j	1.61**	0.54**
	(0.27)	(0.02)
Prior use of higher-tier antibiotic	—	0.18**
		(0.06)
Prior use of lower-tier antibiotic	-0.52	-0.23
	(0.44)	(0.03)
Number of isolates	416,203	416,203

$^{**}p < 0.01$, $^{*}p < 0.05$.

Notes: Table shows logistic regression estimates for the individual-level resistance model. Each observation represents a different positive bacterial isolate, with each observation representing a unique patient-hospital-antibiotic-specimen source-pathogen-collection date combination. Reported coefficients represent either: (i) drug-specific estimates for linezolid (Column 1), or (ii) the average of drug-specific coefficients across all other antibiotics (cephalosporins, clindamycin, trim/sulfa, and vancomycin) (Column 2). The model also includes controls for admission-level variables (ICD diagnosis codes, prior cultures, and demographics); these effects are allowed to vary freely by antibiotic (coefficients not shown). Fixed effects for each drug-hospital-year combination are also included; these fixed effect estimates are used in step 2 of the resistance model estimation. Standard errors (in parentheses) are clustered at the patient level. Coefficients for higher-tier antibiotics are omitted for linezolid because it is already in the highest tier. Coefficients for lower-tier antibiotics exclude cephalosporins, which are already the lowest tier.

or prior use or resistance of lower-tier antibiotics.[29] Reassuringly, predicted resistance probabilities closely match observed resistance, both overall and separately by antibiotic, suggesting that the included covariates capture substantial individual-level heterogeneity in resistance (Figure A4).

Hospital-level Resistance Model Estimates. Table 3 shows estimates of the hospital-level resistance model for linezolid. Column 1 shows a strong first stage: with lagged ASP policies

Figure A5 plots $\tilde{A}_{hjt}$ for linezolid over time, separately by hospital.

[25]This specification is shown in Appendix Section D.2.1. Although this specification lacks the dynamic dependency of current resistance on past resistance, our results in Table 3 show that δ does not differ significantly from zero.

[26]Often, cultures may be repeated for the same patient across multiple days. To arrive at a set of independent isolates, we group cultures within the same patient, specimen type, and pathogen that occur less than 7 days apart.

[27]For example, linezolid resistance is only tested in *Enterococcus spp.*, because resistance in other pathogens is exceedingly rare. In contrast, trimethoprim-sulfamethoxazole resistance is tested in a wider range of pathogens.

[28]By comparison, of isolates tested for vancomycin resistance, 4.5% had prior use of a higher-tier antibiotic (i.e., linezolid), 37% had prior use of vancomycin, and 19% had prior use of a lower-tier antibiotic.

[29]Cross-tier associations may reflect causal biological effects—where use of one antibiotic induces co-resistance to others—or selection, where patients exposed to higher-tier antibiotics are generally more likely to have prior antibiotic exposure. For the purposes of our counterfactual simulations, we treat only the own-antibiotic effects as causal and hold the higher- and lower-tier antibiotic effects fixed.

leading to a 1.0 percentage point decline in lagged linezolid use ($SE = 0.001$, $P < 0.001$). Column 2 shows 2SLS results that imply a large resistance externality, with a 1.0 percentage point increase in the share of admissions using linezolid resulting in a 1.37 increase in the hospital-level fixed effect. This magnitude is consistent with our reduced-form plots in Figure 5 and is the same order of magnitude as the individual-level effects of prior linezolid use (1.61) or prior linezolid resistance (2.47) shown in Table 2. Put another way, exposure to resistance externalities from a 1pp increase in hospital-level linezolid use is nearly equivalent to the resistance effect of using linezolid.

Table 3: Hospital-level Resistance Model Estimates

	First Stage	2SLS	First-Difference GMM	System GMM
Lagged ASP Policy	-0.010^{**} (0.001)			
Lagged Utilization (share)		136.9^{**} (26.49)	159.7^{**} (13.94)	142.9^{**} (21.67)
AR1 Coefficient δ			-0.07 (0.16)	0.03 (0.15)
Fixed Effects	Drug × Hosp Drug × Year	Drug × Hosp Drug × Year	Drug	Drug
N (drug-hospital-year)	108	108	53	53

** $p < 0.01$, * $p < 0.05$, + $p < 0.10$

Notes: Table shows estimates of the hospital-level part of the resistance model in (8). Column 1 shows results from the first stage regression of lagged utilization on the lagged ASP policy. Column 2 shows two-stage least squares (2SLS) estimates of the resistance fixed effects α_{hjy-1} on lagged utilization, instrumented by the lagged ASP policy as in Column 1, assuming $\delta = 0$. Columns 3-4 show estimates of (8) using the between-hospital differences defined in (9), allowing $\delta \neq 0$. Column 3 shows estimates from the first-differences estimator of Arellano and Bond (1991). Column 4 shows the results from the System GMM of Blundell and Bond (1998) and Arellano and Bover (1995). Robust standard errors are shown in parentheses.

Columns 3 and 4 show that both the first-difference and system GMM estimators yield estimates of the AR(1) coefficient δ that are near zero and not statistically significant (-0.07 and 0.03 respectively). This suggests that, once the causal effects of antibiotic utilization are accounted for, resistance does not significantly persist from year to year on its own, at least for linezolid.[30]

Decomposing Resistance Into Externalities and Individual-level Effects. To better illustrate the relative magnitude of the hospital-level externality, Appendix Section D.3 and Table A5 show how we decompose effects into hospital- and individual-level contributions. We find that individual-level effects account for at most 24% of the total observed decline in linezolid resistance caused by the ASP policy, implying that at least 76% of the effect of ASP policies on resistance

[30]This is consistent with the rapid decline in resistance we observe in Figure 5. Thta said, our 95% confidence intervals include values of δ as large as 0.33.

operates through hospital-level externalities.

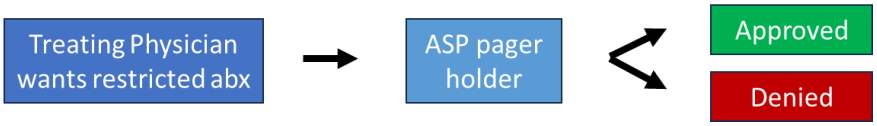
5.5 Mortality

The causal effect of antibiotics and resistance on health outcomes, including mortality, is challenging to identify due to endogenous selection of patients into treatment. This is especially true for restricted broad-spectrum antibiotics, which are only given to patients satisfying strict criteria. To estimate the mortality equation of our structural model, we use a judges design that leverages heterogeneity in antibiotic approval propensities across different providers.

In both of the hospitals in our sample, physicians who wish to administer restricted antibiotics must obtain prior approval by submitting a request to an antibiotic stewardship pager (Figure 6).[31] Because holding the stewardship approval pager is shift-based, the individual holding the pager (a physician, fellow, or pharmacist specializing in infectious disease) changes from week to week. Because different approvers vary in their linezolid approval rates, this generates quasi-random variation in the probability a given patient will receive linezolid.

There are several key advantages to our identification strategy. First, a single individual manages the approval pager each week, and we observe 36 different approvers in our data from 2011Q4-2016Q4. This yields nontrivial week-to-week variation in approval rates. Second, many different approvers are active in a given year, allowing us to control for year fixed effects.[32] Lastly, this pager variation induces individual-level variation in antibiotic use, which is better suited for identifying direct, short-run mortality effects on treated individuals. This complements the hospital-wide policy variation we used to estimate resistance externalities.

Figure 6: ASP Pager Approval Process



Notes: Figure illustrates the prior-authorization process for restricted antibiotics. Treating physicians must request approval from the on-call ASP pager holder before administering a restricted drug. The pager holder, an infectious diseases physician, fellow, or pharmacist, reviews the clinical case and either approves or denies the request.

Instrument Construction. We exploit variation in different providers' propensity to approve linezolid as an instrument for linezolid use. To do this, we use the pager request data to reconstruct weekly shift schedules and compute leave-one-out (LOO) linezolid approval rates for each pager holder. We denote the LOO linezolid approval rate by Z_i and restrict to admissions i beginning

[31] Not all restricted antibiotics are obtained via the pager. Alternative methods include calling an infectious disease consult, or dosing the antibiotic overnight, as the ASP pager is only active during the day. Patients may also receive one dose of restricted antibiotics overnight, but continued use must be justified via the ASP pager the following morning.

[32] Our estimates are robust to whether or not year fixed effects are included.

on the first day of each ASP pager shift.[33] We compute Z_i for 343 weeks and for 36 unique pager holders. Z_i varies about 30 percentage points across approvers in the data. [34]

To improve statistical power, we identify a subset of admissions, denoted by $T_i = 1$, where patients are more likely to require linezolid. This exploits the reasoning that Z_i should only affect outcomes for patients for whom linezolid might be beneficial (e.g., those with more serious infections with gram-positive organisms). Specifically, we use a machine learning algorithm to compute the predicted probability of receiving linezolid during each admission, and assign $T_i = 1$ to admissions in the top decile of this predicted probability (Table A7).[35]

Table A6 reports descriptive statistics for the mortality sample, separately by T_i. Rates of linezolid use are higher for the $T_i = 1$ group (17%), compared to the $T_i = 0$ group (1%). Average 30-day mortality is higher in the $T_i = 1$ group (5.8%) relative to the $T_i = 0$ group (3.3%). Values of Z_i did not differ between the two groups, supporting our assumption that pager-holder shift schedules are random with respect to patient characteristics. Demographics (age, sex, language, ethnicity, and race) are similar between both groups. However, the $T_i = 1$ group has several additional features that suggest that linezolid use would be beneficial for this group, including a much higher VRE positivity rate (32% vs. 3%), longer average lengths of stay (32 days vs. 13 days), and a higher rate of bacteremia or sepsis ICD codes (40% versus 8%). Intuitively, these covariates associated with high predicted linezolid use reflect physicians' assessments of which patients stand to benefit most from the antibiotic.

Identification. We assume that the LOO approval rate Z_i is as good as randomly assigned from the perspective of patients admitted in a given week. Appendix Figure A2 shows that Z_i is uncorrelated with a range of placebo outcomes, including the admission length, predicted mortality, predicted linezolid use, and bacterial culture results. This lends strong support to the validity of our identification strategy.

Reduced-form Evidence of Linezolid's Effect on Mortality. Figure 7 shows binned scatterplots of linezolid use and mortality, separately for patients with high- $T_i = 1$ and low predicted linezolid use $T_i = 0$. Panel A shows the relationship between linezolid use and the leave-one-out pager approval rate Z_i. We code linezolid use as the post-linezolid share of the admission to capture both the extensive margin (whether linezolid was used at all), as well as the intensive margin (how quickly linezolid was used after the start of each admission).[36] Consistent with our identification

[33]That is, for each weekly shift, we compute the average linezolid approval rate for the same pager holder, only including pager requests from that pager holder's other shifts.

[34]Specifically, the LOO approval rate has a 10th percentile of 0.59 and a 90th percentile of 0.90.

[35]The machine learning predictor is a random forest trained on a 10% hold-out sample of patients. More than 200 covariates are available, including demographics, prior admissions, ICD codes, positive culture results, and resistance testing from the past 12 months, as well as current admission characteristics such as ICD codes and culture results. The most important predictors include indicators of whether the patient has a positive screen or culture for vancomycin-resistant *Enterococcus*, the primary pathogen for which non-linezolid antibiotics are ineffective. We constructed predicted mortality using the same random forest predictor.

[36]For example, this variable takes value 1 for admissions where linezolid was given on the first day, 0.5 if linezolid was given halfway through the admission, and 0 if linezolid was never given. Using a binary measure of whether

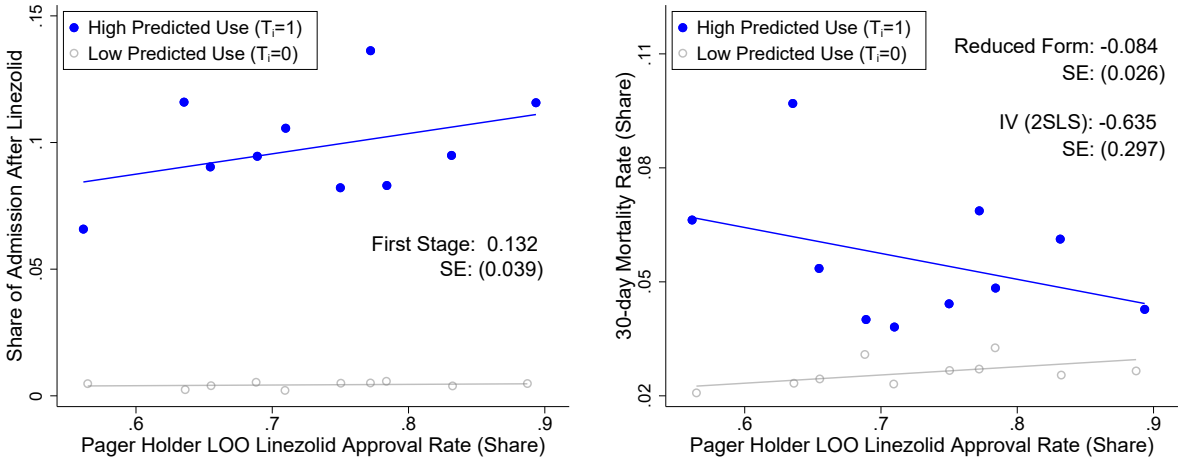
strategy, Z_i strongly predicts linezolid use in the $T_i = 1$ group but not the $T_i = 0$ group. In the $T_i = 1$ group, each 10 percentage point increase in Z_i raises linezolid use by 1.32 percentage points (7.8%).

Panel B shows the relationship between within-hospital 30-day mortality and Z_i for, separately for $T_i = 1$ and $T_i = 0$. In the $T_i = 1$ group, each 10 percentage point increase in Z_i reduces 30-day mortality by 0.8 percentage points (14%). Taken together, these results imply that for the marginal individual, being approved for linezolid earlier in their admission reduces their 30-day mortality risk.

Figure 7: Linezolid Use and Mortality by Providers' Leave-One-Out Approval Rate

(a) First Stage: Linezolid Use

(b) Reduced Form: 30-day mortality



Notes: Figure shows binned scatterplots representing the first stage and reduced form of our instrumental variables framework relating antibiotic use and mortality. Data are shown separately for patients with high ($T_i = 1$) and low predicted linezolid use ($T_i = 0$). Panel A shows that linezolid utilization increases with the leave-one-out (LOO) linezolid approval rate, but only for patients with high predicted use. Panel B shows that 30-day in-hospital mortality declines with the LOO linezolid approval rate, only for patients with high predicted use. Data are limited to admissions beginning on the start date of each approving provider's shift. First stage, reduced form, and two-stage least squares (2SLS) IV estimates are shown based on the specification in Appendix Section D.5. Standard errors are clustered at the approver level.

Our 2SLS estimate of -0.635 (SE = 0.297, $P = 0.032$) implies that receiving linezolid 1 percentage point earlier during an admission decreases mortality by 0.64 percentage points. The large size of this effect may reflect the unique properties of the marginal individuals, whose treating physicians were concerned enough about infection to request linezolid. Our IV estimate is also quite noisy; our standard errors are consistent with a local average treatment effect as small as -0.053.[37]

The goal of these reduced-form results is to illustrate the identification strategy underlying our structural model of mortality, rather than to exactly quantify the magnitude of the mortality effect.

linezolid was used in a given admission does not change our qualitative results.

[37] Another reason for the large size of this effect stems from linearly extrapolating beyond a modest range of linezolid use (from about 0.08 to 0.12 in Figure 7, Panel A), such that raising linezolid use to 100% would result in a negative mortality rate. In contrast, our structural model specification constrains both mortality and linezolid use within [0,1].

In our analyses of counterfactual policies, we test sensitivity to a range of mortality effects.

Structural Estimation. We estimate the following structural model of mortality as an IV probit following Newey (1987).[38]

$$m_i^* = \psi_0 + \psi_1 u_{i,\text{linezolid}} + \mathbf{X_i'}\psi_\mathbf{2} + \mu_i, \quad M_i = \begin{cases} 1 & \text{if } m_i^* \geq 0 \\ 0 & \text{if } m_i^* < 0. \end{cases} \tag{11}$$

The dependent variable M_i is an indicator for whether admission i resulted in in-hospital mortality within 30 days, determined by the latent index m_i^*. The constant term ψ_0 captures baseline mortality risk. The coefficient ψ_1 measures the causal effect of linezolid use on mortality, where $u_{i,\text{linezolid}}$ is an indicator for whether admission i ever received linezolid.[39] The vector $\mathbf{X_i}$ contains exogenous admission-level covariates, including age-by-sex bins, demographics (non-English speaking, non-White, Hispanic), year fixed effects, predicted mortality,[40] predicted mortality squared, and the main effects of Z_i and T_i. Finally, the error term μ_i represents unobserved factors affecting mortality.

We instrument $u_{i,\text{linezolid}}$ using the interaction of the leave-one-out linezolid approval rate $Z_i \times T_i$, an indicator for high versus low predicted linezolid use. We include the main effects of T_i and Z_i in the control vector $\mathbf{X_i}$. The first stage regression therefore takes the form:

$$u_{i,\text{linezolid}} = \delta_0 + \delta_1 Z_i \times T_i + \mathbf{X_i'}\delta_\mathbf{2} + \varepsilon_i. \tag{12}$$

To account for selection into treatment, we assume that the error terms μ_i and ε_i are distributed joint normal; their correlation reflects the endogeneity of antibiotic coverage with respect to mortality risk. The estimated mortality model allows us to evaluate the mortality consequences of both linezolid use and resistance, under the assumption that only non-resistant individuals will benefit from linezolid.

Mortality Model Estimates. Table 4 reports IV probit estimates for the mortality model. Column 1, Row 1 shows that linezolid coverage significantly decreases 30-day mortality, consistent with our reduced-form results in Figure 7. Column 2 establishes the relevance of the instrument $Z_i \times T_i$ for predicting linezolid use. Columns 3 and 4 demonstrate robustness to whether we include year fixed effects (Column 3) and predicted mortality controls (Column 4).

Our baseline estimates in Column 1 translate to an average partial effect (APE) of use on mortality of 0.077 within the $T_i = 1$ group. This implies that, for marginal patients in the $T_i = 1$ group, receiving linezolid reduces 30-day mortality by 7.7 percentage points relative to not receiving

[38] We use the ivprobit function in Stata 17 (StataCorp, 2021).

[39] Here, we use a binary measure of linezolid use to maintain internal consistency with the demand and resistance models and for computational tractability.

[40] We predict mortality using a random forest predictor trained on exogenous admission-level covariates. Our results are also robust to the exclusion of predicted mortality controls.

Table 4: Mortality Model Estimates

	(1) Baseline	(2) First Stage	(3) No Year FE	(4) No Pred. Mortality
Outcome variable:	M_i	$u_{i,\text{linezolid}}$	M_i	M_i
Linezolid Coverage ($u_{i,\text{linezolid}}$)	-0.902**		-0.815**	-0.737**
	(0.233)		(0.262)	(0.237)
Instrument ($Z_i \times T_i$)		1.168*		
		(0.506)		
Z_i	0.072	-0.514	-0.027	-0.005
	(0.147)	(0.379)	(0.151)	(0.161)
T_i	0.050	0.291	0.020	0.447**
	(0.084)	(0.375)	(0.089)	(0.083)
Constant	-2.723**	-1.742**	-2.660**	-2.541**
	(0.249)	(0.267)	(0.230)	(0.217)
Correlation (μ_i, ε_i)	0.51		0.45	0.49
Year Fixed Effects	Y	Y	N	N
Predicted Mortality Controls	Y	Y	Y	N
Outcome Mean ($T_i = 1$)	0.06	0.17	0.06	0.06
N Admissions	18,917	18,917	18,917	18,917

$^{**}p < 0.01$, $^{*}p < 0.05$, $^{+}p < 0.1$.

Notes: Table shows mortality model estimates for the effect of ever using linezolid during an admission on 30-day mortality. Columns 1 and 2 report baseline IV probit and first stage results respectively. Column 3 reports IV probit results excluding year fixed effects. Column 4 reports IV probit results without year fixed effects and without controls for predicted mortality. All specifications include controls for age-by-sex bins, demographics (non-English speaking, non-White, Hispanic). Outcome means are reported for the $T_i = 1$ group. Standard errors are clustered at the pager holder level (N = 36 clusters).

linezolid.[41] The error terms (μ_i, ε_i) are highly positively correlated at 0.51, suggesting that the ASP pager instrument addresses a significant amount of endogeneity. Our model estimates imply that absent linezolid use, mortality rates for patients who received linezolid would have increased from 6% to 27%. Even within the $T_i = 1$ group, higher linezolid use probabilities are associated with greater mortality benefits (Figure A6). All of these findings are consistent with physicians' targeting linezolid to patients with a high mortality risk not fully captured by observables.

Our mortality estimates are best interpreted as *short-run* effects on 30-day mortality per admission. We are not able to assess whether these estimated effects translate into longer-run changes in life expectancy for individual patients. In Section 6.4, we assess sensitivity to varying the effects of antibiotic treatment on mortality.

[41] This estimated effect is smaller in magnitude than the corresponding linear 2SLS estimates in Figure 7, partly because the probit model constrains predicted probabilities to lie between 0 and 1, which compresses large treatment effects near the boundaries of the outcome distribution.

5.6 Innovation Model

A key component of the structural model is the response of innovation to expected profits, for which there are a wide range of existing empirical estimates in the literature (Dubois et al., 2015; Kong and Zhao, 2025; Acemoglu and Linn, 2004; Finkelstein, 2004, e.g.,). Our goal is not to estimate a detailed model of all steps involved in bringing a drug to market, but rather to parsimoniously capture the well-established relationship between market entry and expected profits in our counterfactual simulations. With regard to innovation, our focus is on modeling (1) how counterfactual policies affect profits and (2) how innovation affects welfare. These goals lean more heavily on our other estimates of demand, resistance, and mortality.

Accordingly, we specify a simple entry model where potential entrants draw a random sunk cost of entry and enter the market if their expected profits exceed the entry cost, following Berry and Reiss (2007). To measure profits, we use net-of-rebate U.S. revenue data from marketed antibiotics between 1986 and 2024. To measure entry, we use drug approval data from the FDA Orange Book. Despite its simplicity, the estimated model implies entry costs that are similar to prior estimates and yields innovation elasticities that fall within published ranges. That said, we test sensitivity to alternative innovation elasticities spanning a wide range reported in the literature.

Setup: computing expected profits. We partition our data into three groups, indexed by g: (i) antibiotics treating gram-positive infections only, (ii) antibiotics treating both gram-negative and gram-positive infections, and (iii) antifungals. For each potential entrant, the value of entering in year t is the discounted present value of flow profits:

$$VE(g,t) = \sum_{\tau=1}^{T} \beta^{\tau} \pi(g,t), \tag{13}$$

where $\pi(g,t)$ is the expected annual profit for group g entering in year t, and T is fixed to 10 years, reflecting the FDA exclusivities granted to antibiotics.[42] We assume that marginal costs are zero, so profits equal net revenues. We also assume that when the exclusivity period expires, generic versions enter and profits fall to zero. Our setup assumes that conditional on the group g and entry year t, firms do not have private information about their profits before deciding whether to enter the market. Thus, we first compute average net sales across all antibiotics entering in group g and year t, denoted by $\bar{\pi}_{gt}$. Because these averages are undefined for years with zero entrants, we interpolate by regressing $\bar{\pi}_{gt}$ on group fixed effects and a linear year trend that varies by group. The predictions from this regression are scaled to the global level under the assumption that U.S. revenues are 55% of global revenues (Ho and Pakes, 2024). These scaled revenues, denoted by $\pi(g,t)$, are then used to calculate $VE(g,t)$ using (13).

[42] This includes the 5-year new chemical entity exclusivity plus another 5 years of exclusivity from the Generating Antibiotic Incentives Now (GAIN) Act.

Descriptive Statistics on Antibiotic Revenues and Entry. The distribution of net U.S. revenues for approved antibiotics is shown in Figure A7. U.S. antibiotic revenues average $198M per year, but the distribution is highly skewed, with 75% of approved antibiotics earning below-average revenues. The highest-revenue antibiotics, ciprofloxacin and azithromycin, earned $1.4B and $1.8B per year, respectively.

Figure A8 shows the distribution of $VE(g,t)$ in the data. The prediction and averaging steps result in less variation and skew at the group-by-year level, and the average global PDV of profits is $2.3B. The observed share of potential entrants that enter the market is 15.4%.

Estimation. We assume that each potential entrant draws an entry cost C from an exponential distribution with mean μ.[43] Our method of moments estimator uses the following aggregate moment condition, obtained by averaging over group-year observations:

$$\frac{1}{N}\sum_{g,t} F(VE(g,t);\mu) = \frac{1}{N}\sum_{g,t}\frac{e(g,t)}{K_g}. \tag{14}$$

where $F(VE(g,t);\mu)$ represents the cumulative distribution function of the entry cost distribution evaluated at $VE(g,t)$, $e(g,t)$ is the observed number of entrants in group g and year t, and K_g denotes the fixed number of potential entrants per year, which is equal to the maximum number of entrants observed for each group (3 for gram-positive only, 6 for gram-negative/positive, and 2 for antifungals). The estimation procedure returns the the estimate of μ that matches the average model-implied entry rate to the average observed entry rate.

Innovation Model Estimates. Figure A9 shows the exponential entry cost distribution, which has an estimated mean of $\bar{c}$ =$13.63B. The average *realized* entry cost, defined as the average entry cost conditional on costs being below the expected value of entry, is $1.13B in 2024 dollars. This estimate is consistent with, although somewhat smaller than, previous estimates of new drug development costs during similar time periods. For instance, DiMasi et al. (2003) estimate average development costs (across all therapeutic areas) of $1.46B (in 2024 dollars) per approved compound.

Figure A10 plots how the entry probability varies with profit based on the estimated model. The arc elasticity near the average level of profits is 0.89. This elasticity falls within the range of values reported in the innovation literature, albeit on the higher side. Prior estimates of the elasticity of drug innovation with respect to market size cluster near 0.5 but include values ranging from 0.25–2.75 (Dubois et al., 2015; Finkelstein, 2004). Given the prevailing sentiment of under-investment in antibiotic development (Outterson et al., 2015) and prior empiric work documenting larger elasticities in the antibiotic context (Kong and Zhao, 2025; Majewska, 2022), we view our elasticity of 0.89 as a reasonable benchmark. That said, we also test the sensitivity of our counterfactuals to values at the extreme ranges of 0.25 to 2.75.

[43] We assume that C also includes all costs associated with failed candidates necessary to generate one new FDA-approved antibiotic.

6 Counterfactuals

This section first describes how we run simulations using the estimated structural model. Next, we specify our social welfare function and use it to calculate the value of a new antibiotic entrant. Finally, we solve for optimal counterfactual policies that maximize long-run welfare.

6.1 Simulation Methods

To simulate from the full structural model, we adopt a finite-horizon, forward simulation approach, akin to Igami (2018). Briefly, we simulate demand and resistance jointly, one month at a time, using (6) for demand and (7)-(8) for resistance, and repeat this process $S = 10$ times across different draws of the underlying *i.i.d.* shocks. This captures the lagged structure of the resistance and demand equations and accounts for underlying sampling variation. We then simulate mortality using (11), under the assumption that linezolid use only reduces mortality for non-resistant individuals.[44] For innovation, we model a single potential entrant deciding whether to enter in January 2009, one year after the start of our sample, using our estimates of (14). We assume that the entering antibiotic begins with zero resistance but otherwise bears the same observed characteristics as linezolid and is subject to the same restriction policies. Higher simulated demand and innovation prizes both increase expected profits and thus the probability of entry. For further details on our simulation methods, see Appendix Section D.7.

Appendix Section D.8 describes our baseline simulation results and evaluates comparative statics. We confirm that relaxing antibiotic restrictions increases demand and resistance but lowers mortality. Introducing a novel antibiotic significantly decreases mortality; this effect is larger if restrictions are removed and smaller if restrictions are doubled.

6.2 Social Welfare and Hospital Objectives

Solving for optimal policies first requires specifying the social planner's objective function, which we define as:

$$W \equiv -\sum_{t=2008}^{\infty} \beta^t E\left[\lambda \underbrace{M_t^N(D_t(\theta), R_t)}_{\text{Mortality}} + \underbrace{P \cdot \mathbb{1}\{\text{patent}_t\}}_{\text{Innovation Prize}}\right]. \tag{15}$$

As in Section 3, the social welfare function equals the expected present-discounted value of mortality plus innovation costs, represented here by the innovation prize amount P. M_t^N denotes the yearly number of infection-related deaths scaled nationally to the entire U.S.[45] D_t and R_t represent demand and resistance across all available antibiotics and admissions in year t. Following Igami (2018), we assume that D_t, R_t, and M_t stop changing after the last period of the model. Because averted deaths

[44] We do not observe resistance for patients without positive bacterial cultures; for these patients, we assume that their probability of resistance equals the overall share of resistant cultures in the hospital and year.

[45] We scale Hospital 1 mortality totals by 1,000, which is approximately equal to the number of total inpatient beds in the U.S. divided by number of staffed inpatient beds in Hospital 1.

may occur among individuals with limited life expectancy at baseline, We use a conservative value of statistical life (VSL) of λ =\$1 million, compared to other estimates estimates in the literature of \$7 million or more (Viscusi and Aldy, 2003). For reference, the U.S. Department of Health and Human Services uses a VSL of about \$13 million in their regulatory impact analyses, with a range between \$6 and \$20 million (Kearsley, 2024). We use an annual discount factor of β =0.95. We use θ to denote policies affecting use. The innovation prize $P \geq 0$ is paid annually conditional on $\mathbb{1}\{\text{patent}_t\}$, which takes value 1 for the 10 years after a new antibiotic enters the market.[46] Here, we consider P as a social cost rather than a transfer to drug manufacturers, in order to capture the real costs incurred when bringing a drug to market.[47] Of note, θ affects not only affects mortality via demand but also indirectly via resistance and the entry probability. In contrast, P only affects the probability of entry, allowing for a separation of resistance and innovation externalities.

The Social Value of a New Antibiotic. Equation 15 allows us to compute the social value of a new antibiotic by calculating (15) with and without a new entrant. We find that a new antibiotic reduces mortality by about 2.7 deaths per year in Hospital 1, a 2.8% decrease in average annual deaths. This translates into 2,700 averted deaths per year when scaled up to the level of the U.S. At a VSL of \$1 million, this implies that a new antibiotic generates \$54.7 billion in additional (present-discounted) social welfare, or \$2.74B per year, not including any additional deaths averted globally. Given our estimated development cost of \$1.13 billion per new antibiotic, our findings suggest that antibiotic R&D investments yield high social returns.

Of note, our estimate of social value depend on utilization policies θ, because the effect of the new antibiotic on mortality depends on the degree to which its use is restricted. While holding new antibiotics in reserve may provide insurance against resistance increases, our findings suggest that the costs of doing so can be significant. In addition, our estimate somewhat overstates the marginal value of a single entrant. This is not because of growing resistance for the entrant, which is accounted for by the model. Rather, it is because the value of each entrant declines with subsequent entry.

6.3 Optimal Antibiotic Stewardship Policies

The large social value of a novel antibiotic suggests that innovation is highly welfare relevant. We consider two types of policies available to the social planner that could incentivize innovation: demand-focused policies, innovation prizes, and combinations of the two.

Optimal innovation prizes. A social planner may address innovation directly by awarding a prize P for a successful entrant. While we assume that this prize takes the form of an annual lump-sum payment over 10 years, our findings do not hinge on the exact formulation of this prize.

[46] Recall that we consider only one potential entrant in 2009.

[47] These include, for example, the labor and capital costs of running clinical trials and drug marketing. P may overstate real costs if it includes transfers to firms, or understate real costs if the marginal cost of public funds is greater than 1 (Browning, 1976).

The important aspect of the prize is that it affects innovation directly, without directly distorting treatment decisions.

Table 5 reports optimal policies as a function of the policy levers available to the social planner, along with the probability of entry, long-run social welfare relative to baseline, and average mortality, linezolid demand, and linezolid resistance. In the baseline case, with status quo antibiotic restrictions and no innovation payments, the potential entrant has a 58.2% chance of entering the market.[48]

We find that the optimal innovation prize (holding demand restrictions at status quo levels) awards the innovator $1.04 billion per year for the first 10 years, which increases the entry probability to 85.4% and raises long-run welfare (net of innovation prize payments) by $6.5 billion. Dividing the PDV of prize costs (about $8 billion) by the PDV of lives saved yields a cost-per-life-saved of $541,000. We can calculate the Marginal Value of Public Funds (MVPF) of this policy by dividing the social willingness to pay (i.e., the PDV of lives saved times a $1 million VSL) by the net costs of the policy (Hendren and Sprung-Keyser, 2020). This yields a welfare-improving MVPF of 1.9. Innovation prizes have the added benefit of achieving welfare increases while *decreasing* linezolid resistance, as individuals substitute away from linezolid to the new entrant.

Optimal demand restrictions or subsidies. Next, we consider the case where the social planner controls demand but does not have access to an innovation prize. Currently, demand restrictions θ are determined by private institutions such as hospitals, which have budget constraints and thus may rationally consider antibiotic prices. However, because the marginal cost of production of most drugs is near zero, the private value of θ may differ from the social optimum. A social planner could influence demand in a number of ways, including taxes or subsidies, direct allotments of novel antibiotics, or carve-outs of antibiotic costs from inpatient bundled payments. We denote the net effects of these policies by a "policy scaler" that scales the choice utility of linezolid (and entrants) by the absolute value of the ASP policy effect $\hat{\theta}_{linezolid}$ shown in Table 1,. A policy scaler of 0 maintains status quo restrictions, a policy scaler of 1 simulates removal of the ASP policy on linezolid use, and values larger than 1 induce even higher demand for restricted antibiotics.

We find that the optimal policy scaler is 2.6, resulting in a 43% increase in use from 3% of admissions to 4.3% of admissions. Resistance also increases substantially, from 1.4% to 7.5%, but this is outweighed by declines in mortality from treatment expansion, raising the PDV of welfare by $88.6 billion. While relaxing utilization restrictions incur social costs in terms of higher resistance, it does not entail fiscal costs.[49] Thus, the cost per life saved is $0, and the MVPF is infinite. If we consider all additional antibiotic spending (relative to baseline) as a social cost rather than a transfer, this generates an additional present-discounted cost of $18.7 billion, which increases the cost-per-life-saved to $211,000 and lowers the MVPF to 4.7.

[48] This entry probability is significantly higher than the average entry rate per year reported in Section 5.6. This is because the potential entrant we consider has drug characteristics equivalent to linezolid, which had significantly above-average revenues.

[49] This is because the social marginal cost of production is negligible for antibiotics once they are invented.

Why Increasing Antibiotic Use Raises Welfare. The welfare increase from greater utilization exceeds that of the optimal innovation prize. One reason is because the innovation prize involves direct fiscal costs: entering firms are paid a discounted sum of $8.4B, which is directly subtracted from the PDV of welfare. That said, optimal use only results in a modest increase in the entry probability (52.9% versus 42.8%). This implies that most of the welfare benefit of expanded treatment is due to the mortality benefits of additional antibiotic coverage.

Although optimal use only modestly increases the entry probability, accounting for the *possibility* of innovation plays a pivotal role. Setting the probability of entry to zero (i.e., holding the antibiotic choice set fixed) results in a lower optimal policy scaler of 2.0 (compared to 2.6). This makes sense because innovation partially replenishes the overall stock of antibiotic efficacy, so accounting for future innovation allows for treatment expansion in the present. This also illustrates how pessimistic beliefs about the future stock of discoverable antibiotics can rationalize more conservative use. That said, even with an entry probability of zero, the optimal policy scaler of 2.0 suggests that there may still be room to expand antibiotic treatment.

Another key insight is that resistance affects mortality only for patients who require linezolid coverage. However, beyond a linezolid demand level of 4.6%, resistance increases steeply above 11%, and the associated mortality costs exceed the benefits of additional treatment (Figure 8). This reflects the nonlinear response of resistance to utilization, as shown in Figure A15. This sharp reversal highlights a challenge with demand-based policies: they require careful modeling and monitoring of demand responses and resistance dynamics. As we show in our sensitivity analyses in Section 6.4, the use-resistance relationship is the most important elasticity for determining optimal use. Together, these results show that while expanding use can increase welfare significantly, any utilization increase requires ongoing surveillance to ensure that antibiotic resistance remains within an acceptable range.

Optimal policy with two policy levers. The optimal combination of policies includes a policy scaler of 2.6 combined with a large innovation payment of $1.29B per year (Table 5, Row 4). Combined, these yield a large increase in expected firm profits, reflected in an entry probability of 93.0%. The combination of incentivizing entry and treatment expansion leads to a 5.9% decline in average mortality and a $99.6B increase in the PDV of welfare. Linezolid resistance is also lower than when demand policies are used in isolation (6.0% versus 7.5%). This is because the entering antibiotic absorbs a good share of linezolid demand and hence lowers expected resistance. Given a PDV of innovation prize costs of about $10 billion, the cost-per-life-saved is $90,500, and the MVPF is 11.1. Considering the additional antibiotic spending as a social cost generates an additional present-discounted cost of $22.8 billion, increases the cost-per-life-saved to $298,000, and lowers the MVPF to 3.4.

The heatmap shown in Figure 8 breaks down how different combinations of demand and innovation prizes translate into long-run welfare.[50] Low levels of demand (policy scaler < 0) lead to under-

[50]Figure A13 shows corresponding heatmaps for other outcomes, including mortality, linezolid resistance, entrant

Table 5: Optimal Policies by Available Levers

	(1) Policy scaler	(2) Innov. payment (B$/yr)	(3) Entry prob. (%)	(4) Avg. mortality (%)	(5) Welfare gain (B$)	(6) Linezolid demand (%)	(7) Linezolid resistance (%)
Baseline	0	0	58.2	3.90	0.0	3.0	1.4
Innovation Prize	0	1.04	85.4	3.86	6.5	2.9	1.2
Optimal Strictness	2.6	0	73.9	3.72	88.6	4.3	7.5
Both Policies	2.6	1.29	93.0	3.67	99.6	4.2	6.0

Notes: Table reports simulation results for optimal policies under different available policy levers, where the available policies are chosen to minimize (15). Row 1 shows the baseline with no policy intervention; Row 2 shows the optimal innovation payment given status-quo antibiotic restriction policies; Row 3 shows the optimal demand-side policy scaler, holding the innovation payment fixed at zero; Row 4 shows the joint optimum using both policy instruments. Columns 1 and 2 report the optimal demand policy scaler and innovation prize. Subsequent columns report equilibrium outcomes implied by the model, including the probability of antibiotic entry (Column 3), the population-weighted annual mortality rate (Column 4), the increase in long-run welfare, defined in (15) (Column 5), average linezolid demand (Column 6) and resistance (Column 7). Data are limited to the treatment hospital (Hospital 1).

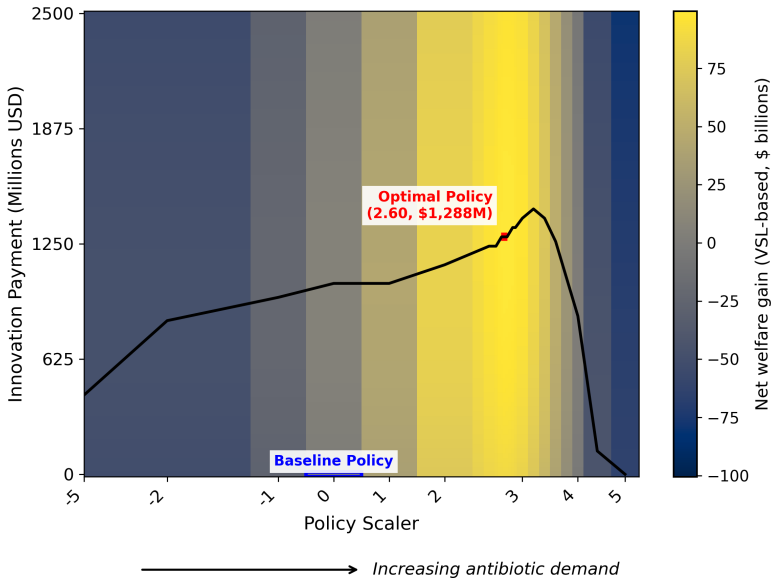
treatment and low entry probabilities, resulting in negative welfare changes compared to the status quo. In contrast, very high levels of demand yield higher entry probabilities but also lead to higher resistance and overall mortality. The black curve shows the optimal innovation payment for each level of use. For most levels of use, the optimal innovation payment is non-zero.

Are policies complements or substitutes? We find that over a large region of policy combinations, demand shifters and innovation prizes are complementary. As the policy scaler increases, so does the optimal innovation payment, reflected in the upward-sloping optimal prize curve in Figure 8. Another result of this complementarity is that the long-run welfare gain from the optimal joint policy ($99.6B) exceeds the sum of gains from the individual policies ($6.5B + $88.6B = $95.1, Table A8).

The intuition is that the value of a new antibiotic is higher the more it is used. Another source of complementarity between use and innovation stems from the convexity of the use-resistance relationship. Spreading the same amount of antibiotic utilization over several antibiotics, including entrants, improves welfare.

demand, entrant profit, and entry probability.

Figure 8: Welfare as a function of demand and innovation policies



Notes: Figure shows changes in simulated welfare (Equation 15), relative to the status quo (*policy scaler* = 0, *innovation payment* = 0), across two policy dimensions: the magnitude of the policy scaler (X-axis; higher values correspond to higher antibiotic demand) and the annual innovation payment (Y-axis). Lighter shading indicates higher net welfare. The black curve denotes the optimal innovation payment as a function of the policy scaler. The boxed cells denote the baseline and welfare-maximizing policy combinations.

Complementarity holds for a large region of the policy space. The optimal innovation prize curve continues to slope upward past the optimal level of demand until a policy scaler value of 3.2. Past that point, the optimal innovation payment falls toward zero. For policy scaler values above 3.2, use and innovation are substitutes. Thus, if baseline use were higher than optimal, then reducing use would both increase welfare and imply higher optimal innovation prizes. A general insight is that whether antibiotics are under- or over-used at baseline is important for determining whether use and innovation are complements or substitutes.

6.4 Sensitivity Analyses.

We test the sensitivity of our policy counterfactuals to varying the parameters encoding each of the key relationships in the model: the response of resistance to use, the response of mortality to antibiotic coverage, and the response of innovation to expected demand. Each modification delivers intuitive effects on our results, shown in Table 6. Our major conclusions remain unchanged in a large region surrounding our preferred specification. First, new antibiotics remain highly valuable in all scenarios. Second, combining demand and innovation policies delivers large welfare gains, with complementarity across most parameter combinations (Figure A14). The sensitivity analyses provide additional intuition for how each component of the model affects our final results. For example, the magnitude of the optimal innovation prize varies and is not always positive.

Table 6: Sensitivity Analyses

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Prize Only		Demand Only		Both Policies		
	Optimal Innovation Prize Only (B$/yr)	Welfare Gain (B$)	Optimal Policy Scaler Only	Welfare Gain (B$)	Optimal Innovation Prize (B$/yr)	Optimal Policy Scaler	Welfare gain (B$)
Baseline	1.04	6.5	2.60	88.6	1.29	2.60	99.6
Varying Resistance Elasticity							
Low ($\lambda_R = 0.5$)	1.04	6.4	**4.00**	165.8	1.16	**4.50**	174.0
High ($\lambda_R = 2.0$)	1.04	6.6	**1.50**	43.6	1.29	**1.50**	54.6
Varying Mortality Elasticity							
Low ($\lambda_M = 0.5$)	**0.35**	0.6	2.60	44.3	**0.61**	2.60	46.2
High ($\lambda_M = 2.0$)	**1.72**	23.7	2.60	177.2	**1.97**	2.65	211.9
Varying Innovation Elasticity							
Low ($\lambda_I = 0.28$)	**0.00**	0.0	2.50	77.3	**1.31**	2.60	79.1
High ($\lambda_I = 3.09$)	**0.71**	14.6	2.70	107.1	**0.45**	2.70	111.6

Notes: Table reports optimal policies under different resistance, mortality, and innovation elasticities. Columns 1-2 show the optimal prize and welfare gain when only innovation prizes are available. Columns 3-4 show the optimal policy scaler and welfare gain when only demand-based policies are available. Columns 5-7 show the optimal policy scaler, innovation prize, and welfare gain when both policy levers are available. The rows show alternative elasticity magnitudes described in Section 6.4.

Varying the Use-Resistance Relationship. We scale the effect of use on resistance by multiplying the estimated resistance elasticity (the estimated $\hat{\gamma}$ coefficient from Equation 8) by $\lambda_R \in \{0.5, 2.0\}$. These values represent cases where resistance externalities are half or twice as large as what we have estimated. Lower resistance elasticities imply higher optimal use (optimal policy scaler = 4.00) and much higher welfare gains from expanded demand and innovation. Higher resistance elasticities reduce the optimal level of use (from a policy scaler of 2.6 to 1.5) and decrease the welfare gains from treatment expansion. A sufficiently high resistance elasticity would reverse our finding that linezolid may be underused at baseline. These patterns make sense as antibiotic use directly trades off current mortality and future resistance.

One useful takeaway is that optimal innovation prizes and their welfare gains are stable across different values of λ_R. This suggests that innovation prizes of about $1B per year are likely to improve welfare regardless of the magnitude of the use-resistance relationship.

Varying Mortality Effects. We vary the mortality effects of antibiotic coverage by directly scaling the estimated linezolid treatment effect for each individual by $\lambda_M \in \{0.5, 2.0\}$. Letting $TE_i \equiv \Pr(M_i|u_{i,\text{linezolid}} = 0) - \Pr(M_i|u_{i,\text{linezolid}} = 1) > 0$ represent the mortality decline from linezolid treatment for individual i and holding the untreated mortality probability constant, we

obtain the new treated mortality probability by adding $(1 - \lambda_M)TE_i$. Thus, $\lambda_M = 0.5$ halves the effect of linezolid on mortality, whereas $\lambda_M = 2.0$ doubles the mortality effects.

Modifying the effect of antibiotic coverage on mortality greatly affects the optimal size of the innovation prizes, but not the optimal use level. This makes sense based on insights from the conceptual framework in (5), which showed that optimal use does not depend on the mortality effect.

Doubling the mortality effect of antibiotics increases the optimal prize level from \$1.04B to \$1.72B and raises the associated gain in the PDV of welfare from \$6.5B to \$23.7B. With both policies, welfare gains increase roughly linearly from \$99.6B to \$211.9B when mortality effects are twice as large. With smaller mortality effects, optimal prize levels and welfare gains are smaller.

Varying the Innovation Elasticity. We vary the innovation elasticity by scaling any changes in expected profits relative to the baseline scenario by λ_I. For example, to reduce the effect of an innovation prize by half, we would represent the effect of prize P by $P/2$. This also affects any profit changes from higher or lower demand. Starting from our estimated innovation elasticity of 0.89, we test a lower bound elasticity of 0.25 using $\lambda_I = 0.28$.[51] We test an upper bound innovation elasticity of 2.75 using $\lambda_I = 3.09$.

Moving to a small innovation elasticity of 0.25 reduces the optimal innovation prize (without changes in use) to \$0. With optimal expansions of use, innovation prizes become helpful once again, but long-run welfare gains are somewhat lower (\$79.1B versus \$99.6B).

Moving to a high innovation elasticity of 2.75 modestly increases welfare gains. Interestingly, optimal innovation prizes are lower. Without expanding use, an optimal prize of \$710M per year leads to an entry probability of 95% and net long-run welfare increase of \$14.6B. With both policies, the optimal prize falls to \$450M, implying that use and innovation are substitutes rather than complements. This is because innovation becomes saturated as entry probabilities approach 1. The point where use and innovation policies transition from complements to substitutes occurs at a threshold innovation elasticity of 1.4.

In addition to the embedded elasticities, other important assumptions include the information structures we have imposed on the entrant and social planner. We assumed a single potential entrant that knows its drug characteristics prior to entry. With multiple potential entrants and heterogeneity in antibiotic quality, promising antibiotics will tend to enter at higher rates even without innovation prizes, and prizes may have larger effects on antibiotics with lower expected demand. Even without heterogeneity in drug characteristics, considering multiple potential entrants would require specifying how the planner would allocate prizes among them. Richer information structures may also involve asymmetric information between the social planner and firms. The design of optimal prize contracts is a complex, active area of research (e.g., Weyl and Tirole, 2012; Che et al., 2021). While outside of the scope of this paper, it remains an important area for future work.

[51] We obtain this value of λ_I because multiplying profit changes by 0.28 achieves an elasticity of 0.28 * 0.89 = 0.25.

7 Discussion

In this paper, we have sought to identify and quantify the rich set of trade-offs involved in regulating antibiotic use and innovation, combining theory, reduced form empirics, and structural approaches.

Our main contribution is to develop a parsimonious framework that characterizes optimal antibiotic policy as a trade-off between direct, short-run mortality effects and future externalities on resistance and innovation. These externalities act in opposite directions: greater use increases resistance but also expands market size and raises the probability of new antibiotic entry. The model delivers several novel insights. Because hospitals do not internalize innovation benefits—and often face high prices for novel, non-generic drugs—antibiotic use may fall below socially optimal levels. Introducing a second policy lever in the form of innovation payments decouples use from innovation incentives. However, while innovation payments are commonly motivated by their potential to reduce antibiotic use by preserving innovation independently of sales revenue, we show that socially optimal use could go up or down in the presence of prizes. Finally, the model demonstrates that antibiotic resistance induces a marginal-inframarginal trade-off: increasing use confers benefits to marginal patients but imposes resistance costs on all inframarginal users. We used this insight to derive a simple empirical test for suboptimal use.

Empirically, we used electronic health record data to estimate reduced-form effects of antibiotic use on resistance and mortality. The EHR data allowed us to observe details on use and resistance for specific antibiotic and pathogen combinations, surmounting limitations of standard insurance claims data. We leveraged hospital-specific policy changes and provider staffing schedules to identify the effects of antibiotic use on resistance and mortality. Taken together, these findings provided an empirical foundation for evaluating policies that jointly target utilization and innovation.

We then estimated a full structural model linking antibiotic use, resistance, mortality, and innovation to conduct policy counterfactuals. This produced several additional insights. First, we found that most of the use-resistance relationship operates through hospital-level externalities rather than individual effects. Second, we found large effects of antibiotic coverage on mortality, establishing the magnitude of the "lives at stake" in the policy trade-off between antibiotic use and resistance. Indeed, our estimates imply that a novel antibiotic would generate $54.7 billion in long-run social value from averted deaths in the U.S. alone. These large mortality benefits are consistent with persistent underinvestment in antibiotic R&D.

Consistent with prior work, our results confirm that direct financial incentives, such as innovation prizes, could materially influence antibiotic innovation. Our innovation model, though simple, produced credible estimates of entry costs and innovation elasticities that align closely with prior literature (DiMasi et al., 2016; Dubois et al., 2015). In our simulations, we found that optimal innovation prizes are large and welfare increasing.

Our results also underscore the effect of antibiotic use regulations on welfare and innovation. The conventional wisdom holds that, due to resistance externalities, broad-spectrum antibiotics are overused and demand restrictions can improve welfare. This may be a valid approach if the set of available antibiotics is fixed and resistance responds rapidly to utilization. However, we show

that if there is even one potential entrant, the intuition that antibiotics are overused may not hold. In our counterfactual simulations, expanding antibiotic use by 43% lowers long-run mortality and increases the PDV of welfare by $88.6B, despite a sizable increase in steady-state resistance. The key mechanisms are treatment expansion and innovation: relaxing antibiotic restrictions expands treatment for both incumbent and novel antibiotics, raises the probability of entry, and mitigates resistance by expanding the antibiotic choice set.

Another key contribution is our finding that use and innovation are complementary over a large region of the policy space, implying that combinations of innovation prizes and higher demand yield the largest welfare gains. This stands in contrast to most arguments for policies that disentangle revenues from use (Årdal et al., 2019), which view policies such as prizes, subscription models, and advanced market commitments as substitutes for higher use. We show that these intuitions only hold when antibiotics are significantly overused at baseline.

For gram-positive infections, we find that the optimal combined policy corresponds to an increase in linezolid demand of about 40% and an innovation prize of $1.29 billion per year. We estimate that this policy would avert over 4,500 deaths in the U.S. per year, for a long-run welfare gain of $99.6 billion. This welfare gain is $4.5 billion larger than the sum of the welfare gains from the two policies considered in isolation.

Antibiotic innovation has become a central focus of recent policy proposals, including the PASTEUR Act, GAIN Act, and DISARM Act (Gregory and Martin, 2022). These initiatives introduce direct incentives—such as lump-sum innovation prizes or exemptions from bundled inpatient payments—to stimulate antibiotic development. Yet, the magnitude of their potential benefits has not been rigorously quantified. Our results provide the first theoretical and empirical evidence that such policies could generate substantial welfare gains through averted infection deaths, with long-run benefits of nearly $100 billion in the United States alone. Our finding that optimal demand produces larger welfare gains than innovation prizes relates to prior work showing that prizes primarily serve to supplement intellectual property arrangements, which scale rewards with innovation quality (Weyl and Tirole, 2012; Che et al., 2021).

Our findings also yield guiding principles for how to structure optimal antibiotic use policies. Frequently proposed policies to address antibiotic innovation—including prizes, subscription models, and advanced market commitments—still require some way of determining the optimal quantities of antibiotics used. To induce optimal quantities, potential approaches range from carving out antibiotics from bundled hospital reimbursement; to taxes or subsidies for novel antibiotics; to direct, fixed allotments of novel antibiotics. Our results show that the effect of antibiotics on mortality has little impact on the level of optimal use, which depends more crucially on the shape of the resistance externality. This insight highlights the value of research efforts that estimate pathogen-drug specific resistance externalities, especially for the most deadly pathogens where timely and appropriate antibiotic choices are likely to greatly influence mortality.

More broadly, we show that antibiotic resistance and innovation are not distinct problems but opposing and interdependent externalities that must be addressed jointly. This insight extends

beyond antibiotics to other areas of public policy where addressing one externality through a Pigouvian tax or restriction could create or worsen an opposing externality elsewhere. Examples include climate policy, where restrictions on economic output could crowd out investment in clean energy and carbon capture technologies; environmental conservation, where restricting land use can weaken incentives for habitat preservation; and housing policy, where rent controls can reduce the quality and supply of housing. Across these domains, the key lesson is that when externalities interact, combinations of policy instruments may yield complementary welfare gains significantly larger than any single policy in isolation.

References

Acemoglu, D. and J. Linn (2004): "Market Size In Innovation: Theory And Evidence From The Pharmaceutical Industry," *Quarterly Journal of Economics*, 119, 1049–1090, iSBN: 00335533.

Adda, J. (2020): "Preventing the Spread of Antibiotic Resistance," *AEA Papers and Proceedings*, 110, 255–259.

Agha, L., S. Kim, and D. Li (2022): "Insurance Design and Pharmaceutical Innovation," *American Economic Review: Insights*, 4, 191–208, publisher: American Economic Association.

Arellano, M. and S. Bond (1991): "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *The Review of Economic Studies*, 58, 277–297, publisher: [Oxford University Press, Review of Economic Studies, Ltd.].

Arellano, M. and O. Bover (1995): "Another look at the instrumental variable estimation of error-components models," *Journal of Econometrics*, 68, 29–51.

Autor, D. H., C. J. Palmer, and P. A. Pathak (2014): "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts," *Journal of Political Economy*, 122, 661–717, publisher: The University of Chicago Press.

Baur, D., B. P. Gladstone, F. Burkert, E. Carrara, F. Foschi, S. Döbele, and E. Tacconelli (2017): "Effect of antibiotic stewardship on the incidence of infection and colonisation with antibiotic-resistant bacteria and Clostridium difficile infection: a systematic review and meta-analysis," *The Lancet Infectious Diseases*, 17, 990–1001.

Belant, J. L., K. C. Elliott, J. E. Hill, and K. F. Kellner (2025): "Does hunting benefit wildlife conservation?" *Frontiers in Ecology and the Environment*, 23, e2864, _eprint: https://esajournals.onlinelibrary.wiley.com/doi/pdf/10.1002/fee.2864.

Berry, S. and P. Reiss (2007): "Empirical Models of Entry and Market Structure," Elsevier, Handbook of Industrial Organization, 1845–1886.

Bilal, A. and D. R. Känzig (2024): "The Macroeconomic Impact of Climate Change: Global vs. Local Temperature," .

Blume-Kohout, M. E. and N. Sood (2013): "Market size and innovation: Effects of Medicare Part D on pharmaceutical research and development," *Journal of Public Economics*, 97, 327–336, publisher: Elsevier B.V.

Blundell, R. and S. Bond (1998): "Initial conditions and moment restrictions in dynamic panel data models," *Journal of Econometrics*.

Brown, E. D. and G. D. Wright (2016): "Antibacterial drug discovery in the resistance era," *Nature*, 529, 336–343, publisher: Nature Publishing Group.

Browning, E. K. (1976): "The Marginal Cost of Public Funds," *Journal of Political Economy*, 84, 283–298, publisher: University of Chicago Press.

CDC (2013): "Antibiotic Resistance Threats in the United States," Tech. rep.

Center for Drug Evaluation and Research (2025): "Approved Drug Products with Therapeutic Equivalence Evaluations | Orange Book," *FDA*, publisher: FDA.

Chapagain, B. P. and N. C. Poudyal (2020): "Economic benefit of wildlife reintroduction: A case of elk hunting in Tennessee, USA," *Journal of Environmental Management*, 269, 110808.

Che, Y.-K., E. Iossa, and P. Rey (2021): "Prizes versus Contracts as Incentives for Innovation," *The Review of Economic Studies*, 88, 2149–2178.

Chyn, E., B. Frandsen, and E. Leslie (2025): "Examiner and Judge Designs in Economics: A Practitioner's Guide," *Journal of Economic Literature*, 63, 401–439.

CMS (2019): "Medicare and Medicaid Programs; Hospital and Critical Access Hospital (CAH) Changes To Promote Innovation, Flexibility, and Improvement in Patient Care; Extension of Timeline for Publication of the Final Rule," in *Federal Register*, vol. 84, 27069, issue: 112.

Dasgupta, S., B. Laplante, H. Wang, and D. Wheeler (2002): "Confronting the Environmental Kuznets Curve," *Journal of Economic Perspectives*, 16, 147–168.

Diamond, R., T. McQuade, and F. Qian (2019): "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco," *American Economic Review*, 109, 3365–3394.

DiazGranados, C. A., S. M. Zimmer, M. Klein, and J. A. Jernigan (2005): "Comparison of mortality associated with vancomycin-resistant and vancomycin-susceptible enterococcal bloodstream infections: a meta-analysis," *Clinical Infectious Diseases: An Official Publication of the Infectious Diseases Society of America*, 41, 327–333.

DiMasi, J. A., H. G. Grabowski, and R. W. Hansen (2016): "Innovation in the pharmaceutical industry: New estimates of R&D costs," *Journal of Health Economics*, 47, 20–33.

DiMasi, J. A., R. W. Hansen, and H. G. Grabowski (2003): "The price of innovation: new estimates of drug development costs," *Journal of Health Economics*, 22, 151–185.

Dubois, P., O. de Mouzon, F. Scott-Morton, and P. Seabright (2015): "Market size and pharmaceutical innovation," *RAND Journal of Economics*, 46, 844–871, iSBN: 0741-6261.

Dubois, P. and G. Gokkoca (2023): "Antibiotic Demand in the Presence of Antimicrobial Resistance," *Working Paper*.

Falagas, M. E., I. I. Siempos, and K. Z. Vardakas (2008): "Linezolid versus glycopeptide or beta-lactam for treatment of Gram-positive bacterial infections: meta-analysis of randomised controlled trials," *The Lancet. Infectious Diseases*, 8, 53–66.

Fankhauser, S. and R. S.J. Tol (2005): "On climate change and economic growth," *Resource and Energy Economics*, 27, 1–17.

Finkelstein, A. (2004): "Static and dynamic effects of health policy: Evidence from the vaccine industry," *Quarterly Journal of Economics*, 119, 527–564.

Gaessler, F. and S. Wagner (2018): "Patents, Data Exclusivity, and the Development of New Drugs," *Academy of Management Proceedings*, 2018, 17812.

Glavina, A., K. Mišić, J. Baleta, J. Wang, and H. Mikulčić (2025): "Economic development and climate change: Achieving a sustainable balance," *Cleaner Engineering and Technology*, 26, 100939.

Goldsmith-Pinkham, P., P. Hull, and M. Kolesár (2025): "Leniency Designs: An Operator's Manual," .

Gregory, E. and C. Martin (2022): "The Intersection of Antimicrobial Stewardship, the Pharmaceutical Industry, and the Federal Legislature," *Open Forum Infectious Diseases*, 9, ofac404.

Hendren, N. and B. Sprung-Keyser (2020): "A Unified Welfare Analysis of Government Policies," *The Quarterly Journal of Economics*, 135, 1209–1318.

Herrmann, M. and R. Laxminarayan (2010): "Antibiotic Effectiveness: New Challenges in Natural Resource Management," *https://doi-org.ezp-prod1.hul.harvard.edu/10.1146/annurev.resource.050708.144125*, 2, 125–138, publisher: Annual Reviews.

Ho, K. and A. Pakes (2024): "Policy Options for the Drug Pricing Conundrum," *NBER Working Paper 32606*.

Horowitz, J. B. and H. B. Moehring (2004): "How property rights and patents affect antibiotic resistance," *Health Economics*, 13, 575–583.

Igami, M. (2018): "Industry Dynamics of Offshoring: The Case of Hard Disk Drives," *American Economic Journal: Microeconomics*, 10, 67–101.

Kearsley, A. (2024): "HHS Standard Values for Regulatory Analysis," Tech. rep., Assistant Secretary for Planning and Evaluation.

Kling, J. R. (2006): "Incarceration Length, Employment, and Earnings," *American Economic Review*, 96, 863–876.

Kohno, S., K. Yamaguchi, N. Aikawa, Y. Sumiyama, S. Odagiri, N. Aoki, Y. Niki, S. Watanabe, M. Furue, T. Ito, R. Croos-Dabrera, and K. J. Tack (2007): "Linezolid versus vancomycin for the treatment of infections caused by methicillin-resistant Staphylococcus aureus in Japan," *The Journal of Antimicrobial Chemotherapy*, 60, 1361–1369.

Kong, E. and O. Zhao (2025): "Market Incentives and the Drug Development Pipeline: Evidence from Antibiotics," .

Kremer, M., J. Levin, and C. M. Snyder (2020): "Advance Market Commitments: Insights from Theory and Experience," *AEA Papers and Proceedings*, 110, 269–273.

Laxminarayan, R. and M. L. Weitzman (2002): "On the implications of endogenous resistance to medications," *Journal of Health Economics*, 21, 709–718, publisher: North-Holland.

Lewis, K. (2013): "Platforms for antibiotic discovery," *Nature Reviews Drug Discovery*, 12, 371–387, publisher: Nature Publishing Group.

Lindsey, P., P. Roulet, and S. Romañach (2007): "Economic and conservation significance of the trophy hunting industry in sub-Saharan Africa," *Biological Conservation*, 134, 455–469.

Lipsitch, M. (2001): "The rise and fall of antimicrobial resistance," *Trends in Microbiology*, 9, 438–444, publisher: Elsevier Current Trends.

Majewska, G. (2022): "Incentivizing novelty in antibiotic development," .

McFadden, D. L. (1974): "Conditional logit analysis of qualitative choice behavior," in *Frontiers in Econometrics*, ed. by P. Zarembka, New York: Academic Press, 105–142.

Murray, B. E. (2000): "Vancomycin-Resistant Enterococcal Infections," *New England Journal of Medicine*, 342, 710–721, publisher: Massachusetts Medical Society.

Murray, C. J., K. S. Ikuta, F. Sharara, et al. (2022): "Global burden of bacterial antimicrobial resistance in 2019: a systematic analysis," *The Lancet*, 399, 629–655, publisher: Elsevier B.V.

Newey, W. K. (1987): "Efficient estimation of limited dependent variable models with endogenous explanatory variables," *Journal of Econometrics*, 36, 231–250.

Niewiadomska, A. M., B. Jayabalasingham, J. C. Seidman, L. Willem, B. Grenfell, D. Spiro, and C. Viboud (2019): "Population-level mathematical modeling of antimicrobial resistance: a systematic review," *BMC Medicine 2019 17:1*, 17, 1–20, publisher: BioMed Central.

Nordhaus, W. (2018): "Projections and Uncertainties about Climate Change in an Era of Minimal Climate Policies," *American Economic Journal: Economic Policy*, 10, 333–360.

Outterson, K., H. W. Boucher, A. Engel, R. M. Conti, and G. Gracia (2022): "Antimicrobial Resistance: What's At Stake And What Are We Doing About It?" .

Outterson, K., J. H. Powers, G. W. Daniel, and M. B. McClellan (2015): "Repairing the broken market for antibiotic innovation," *Health Affairs*, 34, 277–285, publisher: Project HOPE.

O'Brien, K. S., A. M. Arzika, A. Amza, R. Maliki, B. Aichatou, I. M. Bello, D. Beidi, N. Galo, N. Harouna, A. M. Karamba, S. Mahamadou, M. Abarchi, A. Ibrahim, E. Lebas, B. Peterson, Z. Liu, V. Le, E. Colby, T. Doan, J. D. Keenan, C. E. Oldenburg, T. C. Porco, B. F. Arnold, and T. M. Lietman (2024): "Azithromycin to Reduce Mortality — An Adaptive Cluster-Randomized Trial," *New England Journal of Medicine*, 391, 699–709, publisher: Massachusetts Medical Society _eprint: https://www.nejm.org/doi/pdf/10.1056/NEJMoa2312093.

Pakes, A., M. Ostrovsky, and S. Berry (2007): "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," *RAND Journal of Economics*, 38, 373–399.

Roodman, D. (2009): "How to do Xtabond2: An Introduction to Difference and System GMM in Stata," *The Stata Journal*, 9, 86–136, publisher: SAGE Publications.

Roope, L. S., R. D. Smith, K. B. Pouwels, J. Buchanan, L. Abel, P. Eibich, C. C. Butler, P. S. Tan, A. Sarah Walker, J. V. Robotham, and S. Wordsworth (2019): "The challenge of antimicrobial resistance: What economics can contribute," *Science*, 364.

Schuts, E. C., M. E. Hulscher, J. W. Mouton, C. M. Verduin, J. W. Stuart, H. W. Overdiek, P. D. van der Linden, S. Natsch, C. M. Hertogh, T. F. Wolfs, J. A. Schouten, B. J. Kullberg, and J. M. Prins (2016): "Current evidence on hospital antimicrobial stewardship objectives: A systematic review and meta-analysis," Pages: 847-856 Volume: 16 Issue: 7 ISSN: 14744457.

Stamm, L. V. (2010): "Global Challenge of Antibiotic-Resistant Treponema pallidum," *Antimicrobial Agents and Chemotherapy*, 54, 583–589.

Stern, N. and J. E. Stiglitz (2023): "Climate change and growth," *Industrial and Corporate Change*, 32, 277–303.

Stosor, V., L. R. Peterson, M. Postelnick, and G. A. Noskin (1998): "Enterococcus faecium Bacteremia: Does Vancomycin Resistance Make a Difference?" *Archives of Internal Medicine*, 158, 522–527, publisher: American Medical Association.

Tommasi, R., D. G. Brown, G. K. Walkup, J. I. Manchester, and A. A. Miller (2015): "ESKAPEing the labyrinth of antibacterial discovery," *Nature Reviews Drug Discovery*, 14, 529–542, publisher: Nature Publishing Group.

Viscusi, W. K. and J. E. Aldy (2003): "The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World," *Journal of Risk and Uncertainty*, 27, 5–76.

Walsh, C. T. and T. A. Wencewicz (2014): "Prospects for new antibiotics: a molecule-centered perspective," *The Journal of Antibiotics*, 67, 7–22, publisher: Nature Publishing Group.

Weigelt, J., K. Itani, D. Stevens, W. Lau, M. Dryden, and C. Knirsch (2005): "Linezolid versus Vancomycin in Treatment of Complicated Skin and Soft Tissue Infections," *Antimicrobial Agents and Chemotherapy*, 49, 2260–2266.

Weyl, E. G. and J. Tirole (2012): "Market Power Screens Willingness-to-Pay," *The Quarterly Journal of Economics*, 127, 1971–2003.

Wilcox, M. H. (2003): "Efficacy of linezolid versus comparator therapies in Gram-positive infections," *The Journal of Antimicrobial Chemotherapy*, 51 Suppl 2, ii27–35.

Wunderink, R. G., M. S. Niederman, M. H. Kollef, A. F. Shorr, M. J. Kunkel, A. Baruch, W. T. McGee, A. Reisman, and J. Chastre (2012): "Linezolid in Methicillin-Resistant Staphylococcus aureus Nosocomial Pneumonia: A Randomized, Controlled Study," *Clinical Infectious Diseases*, 54, 621–629.

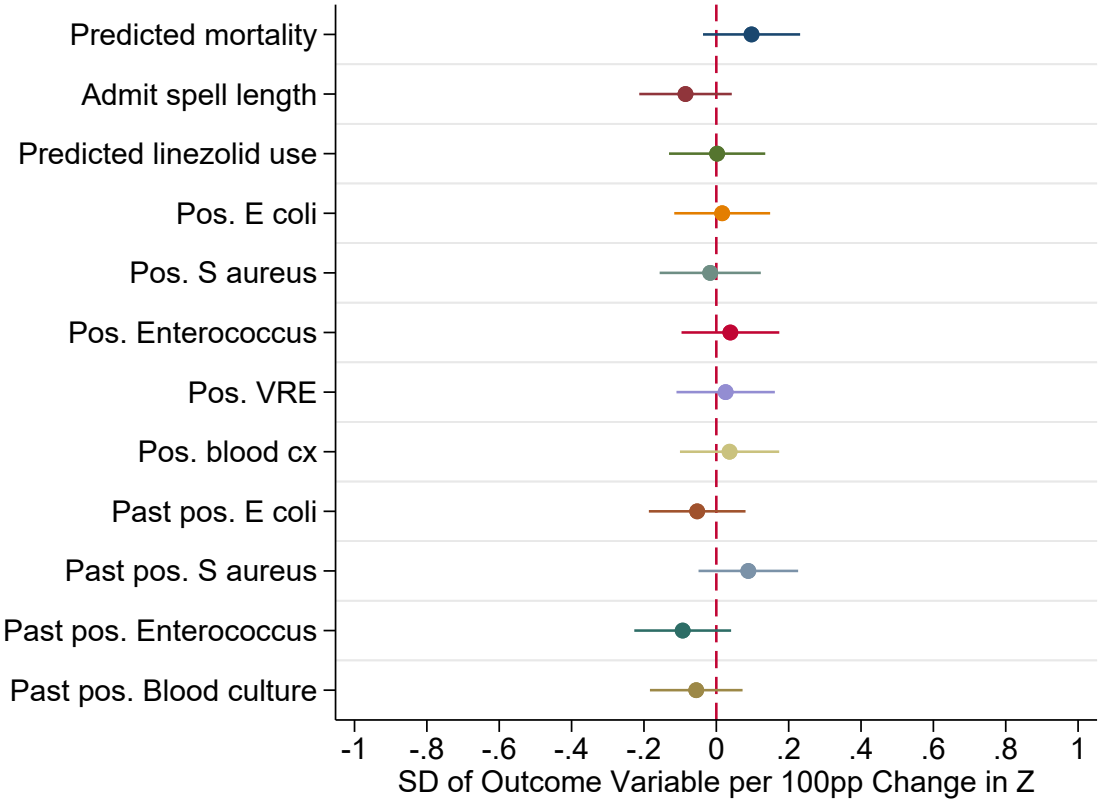
Årdal, C., M. Balasegaram, R. Laxminarayan, D. McAdams, K. Outterson, J. H. Rex, and N. Sumpradit (2019): "Antibiotic development — economic, regulatory and societal challenges," *Nature Reviews Microbiology 2019 18:5*, 18, 267–274, publisher: Nature Publishing Group.

A Appendix

This appendix contains supplementary materials.

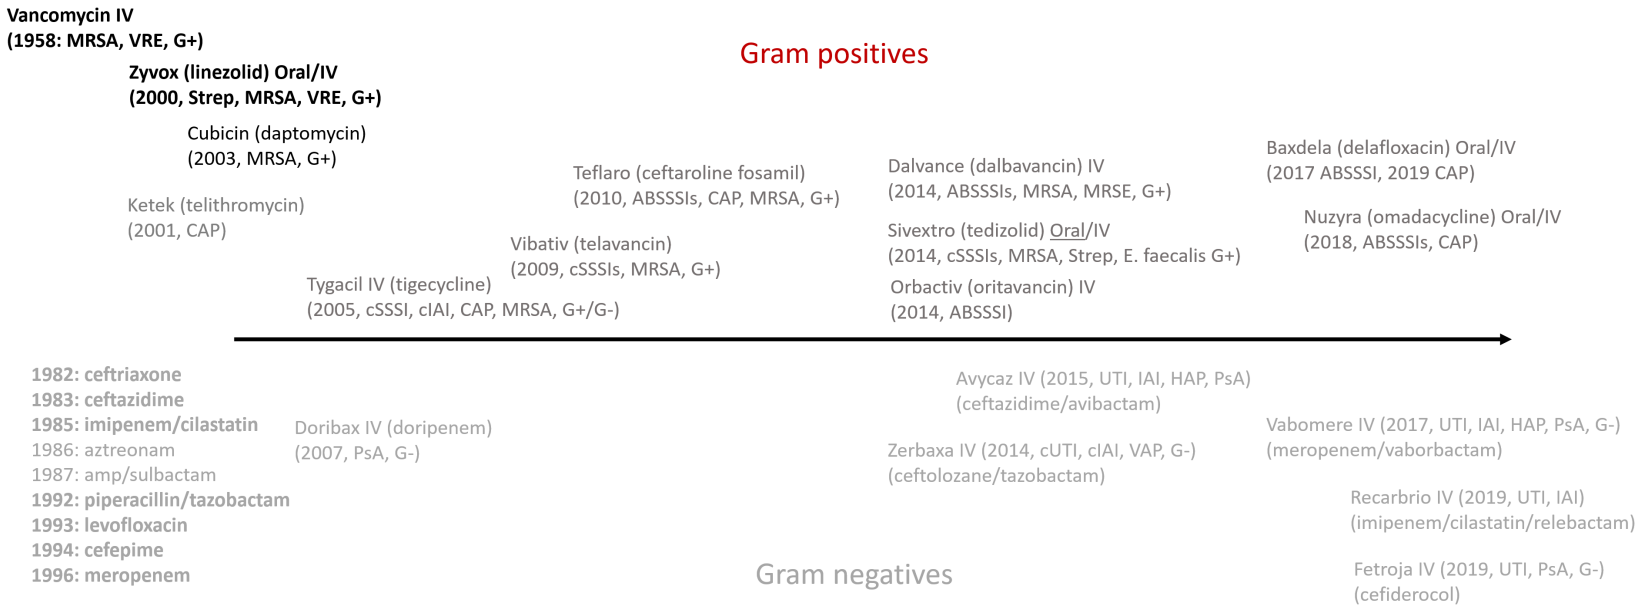
B Additional Exhibits

Figure A2: Mortality Model: Placebo Outcomes and LOO Linezolid Approval Rates



Notes: Figure shows OLS regressions of various placebo outcomes on the leave-one-out (LOO) linezolid approval rate. Each estimate represents the change in that outcome (in standard deviation units) for a 100 percentage point increase in the LOO approval rate (from 0 to 1). Predicted mortality and linezolid use are derived from random forest predictors trained on exogenous covariates. Admit spell length measures the length of each admission (in days), including readmissions within 30 days. Other placebo outcomes include whether the admission had a positive culture for *E. coli, S. aureus, Enterococcus spp.*, positive VRE screen, or any positive blood culture. We also tested placebo outcomes occurring within the prior 12 months.

Figure A1: Timeline of Antibiotic Innovation



Notes: Figure shows approval dates and indications for antibiotics treating primarily gram-positive bacteria (above the horizontal line) and gram-negative bacteria (below the horizontal line). Bolded antibiotics are those most commonly used in our sample.

C Model Appendix

C.1 Derivations for Graphical Model

This section presents derivations for the marginal benefit and cost curves shown in Figure 2. In this simplified version, D represents overall antibiotic demand in steady state across all antibiotics. $R(D, I)$ represents an index of overall steady-state resistance, which depends on demand and innovation, with $\frac{\partial R}{\partial D} > 0$ and $\frac{\partial R}{\partial I} < 0$. Higher innovation decreases resistance by increasing the flow of new antibiotics. $I(D, P)$ represents steady-state investment in antibiotic R&D as an increasing function of demand and direct payments P from the social planner. Mortality $M(D, R)$ is decreasing in demand but increasing in resistance.

The social planner's welfare function $W(D, P)$ considers mortality as a cost (scaled by λ, the value of statistical life) plus innovation costs cI, where the marginal cost c is taken to be constant for simplicity:

$$W(D, P) \equiv -\lambda M(D, R(D, I)) - cI(D, P). \tag{16}$$

Note that we have made the reasonable assumption that social marginal costs of production are zero for antibiotics. We also assume that innovation payments P are considered to be transfers from the social standpoint.

Private healthcare entities such as hospitals do not bear innovation costs, but do pay unit prices for antibiotics p. Assuming they use the same value of λ, the hospital objective function $HW(D)$ is:

$$HW(D) \equiv -\lambda M(D, R(D, I)) - pD. \tag{17}$$

Next, we derive optimal policies, which equate marginal benefits and costs, for several cases corresponding to the panels of Figure 2.

Social Planner, Demand-based Policy Only: Panel A shows the case where the social planner minimizes welfare costs, given access only to demand-based policies such as antibiotic taxes/subsidies, restriction policies, or fixed antibiotic allotments. Taking the derivative of (16) with respect to demand yields:

$$\underbrace{-\lambda \frac{\partial M}{\partial D}}_{\text{Direct Mortality Effect (+)}} + \underbrace{\left(\lambda \frac{\partial M}{\partial R}\left|\frac{\partial R}{\partial I}\right| - c\right)\frac{\partial I}{\partial D}}_{\text{Innovation Externality (sign ambiguous)}} = \underbrace{\lambda \frac{\partial M}{\partial R}\frac{\partial R}{\partial D}}_{\text{Resistance Externality (+)}} . \tag{18}$$

The left hand side represents the marginal benefits of increasing antibiotic use, which include the direct mortality effect as well as the innovation externality. Since mortality is decreasing in demand ($\frac{\partial M}{\partial D} < 0$), we write the direct mortality effect, $-\frac{\partial M}{\partial D}$, as $\left|\frac{\partial M}{\partial D}\right| > 0$. The innovation term, $\left(\lambda \frac{\partial M}{\partial R}\left|\frac{\partial R}{\partial I}\right| - c\right)\frac{\partial I}{\partial D}$, has ambiguous sign: higher demand stimulates innovation ($\partial I/\partial D > 0$), which lowers resistance ($\partial R/\partial I < 0$) and reduces mortality, but this benefit must be weighed against the marginal social cost of additional R&D, c.

The right hand side represents social marginal costs, which comprise the resistance externality. Because higher antibiotic use increases resistance ($\frac{\partial R}{\partial D} > 0$) and higher resistance raises mortality ($\frac{\partial M}{\partial R} > 0$), the resistance term $\lambda \frac{\partial M}{\partial R} \frac{\partial R}{\partial D}$ is strictly positive. This highlights the opposing externalities that arise when a single policy instrument—antibiotic use—jointly determines current health outcomes and innovation incentives.

Hospital, Demand-based Policy Only: The first-order condition (FOC) below shows the case where a hospital minimizes its objective function, given access to demand-based policies such as restriction policies, corresponding to Panel B of Figure 2.

$$\underbrace{\lambda \left| \frac{\partial M}{\partial D} \right|}_{\text{Direct Mortality Effect}} = \underbrace{\lambda \frac{\partial M}{\partial R} \frac{\partial R}{\partial D}}_{\text{Resistance Externality}} + \underbrace{p}_{\text{Antibiotic Costs}} . \tag{19}$$

The left hand side represents marginal benefits of increasing antibiotic use, which includes the direct mortality effects but *excludes* the innovation externality. The right hand side represents hospital marginal costs, which include the resistance externality but also includes the price paid for antibiotics. Compared to the social planner's solution in (16), hospital demand is lower for two reasons: first, it does not take into account innovation and second, it takes into account antibiotic prices.

Social Planner with Demand Shifters and Innovation Payments: Figure 3 shows the case where the social planner has access to both demand-based policies and direct innovation payments. Combining these two policies results in the following first-order conditions for demand D and payments P respectively:

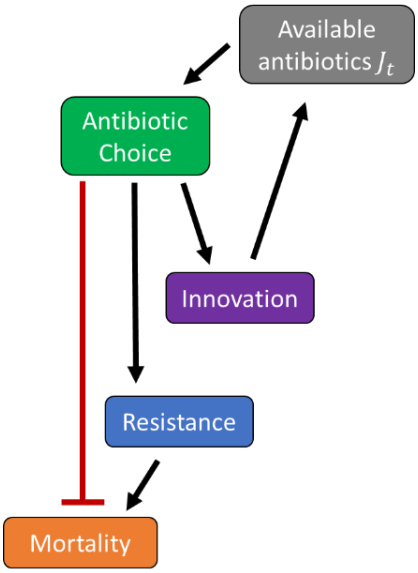
$$\text{FOC } D: \quad \underbrace{\lambda \left| \frac{\partial M}{\partial D} \right|}_{\text{Direct Mortality Effect}} + \underbrace{\left(\lambda \frac{\partial M}{\partial R} \left| \frac{\partial R}{\partial I} \right| - c \right) \frac{\partial I}{\partial D}}_{\text{Innovation Externality}} \overset{0}{=} \underbrace{\lambda \frac{\partial M}{\partial R} \frac{\partial R}{\partial D}}_{\text{Resistance Externality}} , \tag{20}$$

$$\text{FOC } P: \quad \underbrace{\left(\lambda \frac{\partial M}{\partial R} \left| \frac{\partial R}{\partial I} \right| - c \right) \frac{\partial I}{\partial P}}_{\text{Innovation Externality}} = 0$$

where we assume an interior solution in which innovation responds to payments, $\partial I / \partial P > 0$. Under this condition, the first-order condition for P implies $\lambda \frac{\partial M}{\partial R} \left| \frac{\partial R}{\partial I} \right| = c$, which eliminates the innovation externality from the first-order condition for D.

D Structural Model Appendix

Figure A3: Relationship Between Model Equations



NOTES: Figure shows the relationship between the four model equations. Black arrows indicate positive effects (e.g., increased antibiotic use increases resistance). Red capped line segments indicate decreases (e.g., increased antibiotic use decreases mortality).

D.1 Demand Model - Supplementary Methods and Exhibits

D.1.1 Difference-in-differences and Triple Difference Specifications

This section describes our methodology for estimating the reduced form relationships between antibiotic restriction policies, antibiotic use (Figure 4), and resistance (Figure 5). Our empirical approach employs a difference-in-differences design leveraging two exogenous policy changes that only occurred in the treatment hospital. In addition, we run triple-difference specifications that compare our restricted antibiotic of interest (linezolid) to unrestricted antibiotics that were unaffected by the policy changes.

Difference-in-Differences (DiD). To document reduced-form shifts in utilization and resistance following the antibiotic restrictions, we first estimate DiD models comparing Hospital 1 to Hospital 2 before versus after policy implementation. For utilization, the outcome is the share of admissions with any linezolid use, aggregated to the hospital-by-half-year level. For resistance, the outcome is the share of *Enterococcus* isolates resistant to linezolid, aggregated to the hospital-by-year level.

For each outcome, we estimate:

$$Y_{ht} = \alpha_h + \delta_t + \beta_1(\text{Hospital1}_h \times \text{Post}_t) + \varepsilon_{ht}, \tag{21}$$

where Y_{ht} is either the utilization share or the resistance share in hospital h and period t. Hospital1_h is an indicator for the treated hospital, and Post_t indicates post-policy periods, where the first post-policy period is 2011Q4-2012Q1 for utilization and 2013 for resistance. For utilization, Post_t takes an intermediate value of 0.5 between the first and second policies to capture the sequential drops in use after each policy. We include hospital and time fixed effects and use robust standard errors.

Triple Differences (DDD). To further refine our identification strategy, we estimate triple-difference specifications that compare between-hospital changes for linezolid to that of a set of control drugs, which were not affected by the ASP policies. For utilization, the comparison group consists of gram-positive agents not subject to restriction.[52] For resistance, the comparison group consists of the subset of drugs for which non-susceptibility is routinely measured across hospitals.[53] We log-transform the outcome variables to account for large differences in baseline use and resistance levels between drugs. The triple-difference specifications ensure that the hospital-specific changes in use and resistance we find in our DiD results are specific to linezolid, and do not affect other drugs.

We estimate the following DDD specification:

$$\log(Y_{hdt}) = \beta_1(\text{Linezolid}_d \times \text{Hospital } 1_h \times \text{Post}_t) + \text{FE} + \varepsilon_{hdt}, \tag{22}$$

where Y_{hdt} is either the utilization share at the hospital–drug–half-year level or the resistance share at the hospital–drug–pathogen–year level. We include a saturated set of fixed effects, including hospital-by-drug fixed effects, hospital-by-time fixed effects, drug-by-time fixed effects. For resistance, we additionally include hospital-drug-pathogen and pathogen-by-year fixed effects. These absorb all cross-sectional and secular differences in drug use and resistance patterns. Standard errors are clustered at the hospital-drug level for utilization and hospital-drug-organism level for resistance.

Event Study DDD Specification. To assess dynamic treatment effects and evaluate the parallel trends assumption, we also estimate and plot coefficients from event-study versions (22):

$$\log(Y_{hdt}) = \sum_{k \neq k_0} \beta_k \left(\text{Linezolid}_d \times \text{Hospital } 1_h \times \mathbf{1}\{t = k\}\right) + \text{FE} + \varepsilon_{hdt}, \tag{23}$$

where k_0 denotes the pre-policy reference period. The resulting coefficients trace the evolution of linezolid utilization and resistance at Hospital 1 relative to Hospital 2 and relative to comparison drugs.

[52] Comparison antibiotics for utilization comprise amoxicillin, ampicillin, amoxicillin–clavulanate, ampicillin–sulbactam, azithromycin, cefazolin, cephalexin, clarithromycin, clindamycin, dicloxacillin, doxycycline, erythromycin, minocycline, nafcillin, penicillin G, penicillin V, rifampin, trimethoprim–sulfamethoxazole, and vancomycin.

[53] Comparison antibiotics for resistance comprise cephalosporins, clindamycin, trimethoprim/sulfamethoxazole, and vancomycin.

Table A1: Descriptive Statistics (Demand)

Drug	Tier	Chosen	Past use	Predicted Resistance	N
No Antibiotic of Choice Set	0	0.42	0.000	0.000	386,498
Cephalosporin	1	0.111	0.148	0.475	386,498
Clindamycin	2	0.025	0.045	0.348	386,498
Trimethoprim/Sulfamethoxazole	2	0.053	0.098	0.189	386,498
Vancomycin	3	0.364	0.149	0.146	386,498
Linezolid	4	0.028	0.018	0.007	386,498

Notes: Table shows descriptive statistics from the drug-level choice dataset. "Chosen" denotes the fraction of admissions where each drug represents the highest-tier antibiotic selected from the hospital choice set. "Past use" corresponds to the share of patients with any use of the antibiotic in the past 12 months. "Predicted Resistance" denotes model-predicted resistance to each drug using Equation (7), averaging across pathogens.

Table A2: Patient Characteristics by Hospital

Panel (a): Patient Characteristics by Sample

Variable	Hospital 1	Hospital 2	Total
Demographics			
Age (years)	56	53	55
Age standard deviation (years)	23	25	24
Male sex	0.52	0.48	0.51
Non-English speaking	0.11	0.14	0.12
Hispanic	0.030	0.041	0.035
Non-white	0.21	0.28	0.24
ICD diagnoses on admission			
Cancer	0.32	0.40	0.35
Catheter	0.26	0.21	0.24
Pneumonia	0.19	0.17	0.18
Fever	0.21	0.26	0.23
Abdominal pain	0.23	0.20	0.22
Post-operative infection	0.043	0.048	0.045
Urinary tract infection	0.14	0.12	0.13
Bacteremia/sepsis	0.12	0.13	0.13
Cellulitis/abscess	0.13	0.11	0.13
Heart failure	0.22	0.19	0.21
End-stage renal disease	0.042	0.035	0.039
Transplant	0.064	0.084	0.072
Prior diagnoses (past 12 months)			
Cancer	0.13	0.18	0.15
Catheter	0.094	0.084	0.090
Pneumonia	0.081	0.071	0.077
Fever	0.084	0.11	0.093
Abdominal pain	0.099	0.092	0.096
Post-operative infection	0.019	0.022	0.020
Urinary tract infection	0.058	0.050	0.055
Bacteremia/sepsis	0.050	0.051	0.050
Cellulitis/abscess	0.054	0.051	0.053
Heart failure	0.10	0.086	0.094

Continued on next page

Table A2 (continued)

Variable	Hospital 1	Hospital 2	Total
End-stage renal disease	0.024	0.021	0.022
Transplant	0.034	0.041	0.037
Culture result by pathogen (past 12 months)			
Coag. Negative Staph.	0.029	0.031	0.030
E. coli	0.061	0.033	0.049
Enterococcus	0.054	0.033	0.046
K. pneumoniae	0.027	0.020	0.024
P. aeruginosa	0.028	0.020	0.025
S. aureus	0.051	0.040	0.047
Streptococcus spp.	0.020	0.011	0.016
Other gram-negative	0.051	0.038	0.045
Other gram-positive	0.0023	0.00065	0.0016
Culture results by site (past 12 months)			
Abdominal	0.0088	0.0059	0.0076
Blood	0.029	0.037	0.033
JP drain	0.0038	0.0052	0.0044
Lung	0.036	0.024	0.031
Nasal	0.0025	0.0039	0.0031
Ortho	0.009	0.0056	0.0076
Rectal	0.020	0.0074	0.015
Skin/soft tissue	0.027	0.027	0.027
Urine	0.095	0.052	0.077
Wound	0.027	0.023	0.025
Prior hospital admission	0.36	0.38	0.37
N admissions	224,780	161,718	386,498

Notes: Table reports patient-level means for demographics, diagnoses, and microbiologic results by hospital and overall. Variables are defined as indicators unless otherwise noted.

Table A3: Demand Model - Full Estimates

	Cephalosporin	**Clindamycin**	**Linezolid**	**Trim/Sulfa**	**Vancomycin**
Linezolid restriction (self)	–	–	-0.374^{**}	–	–
Linezolid restricted (Vancomycin)	–	–	–	–	0.167^{**}
Linezolid restricted (Lower-tier)	0.089^{**}	0.089^{**}	–	0.089^{**}	–
Predicted resistance probabilities					
Predicted resistance (self)	-0.700^{**}	0.427^{*}	-1.291	0.778^{**}	-0.748^{**}
Predicted resistance (higher-tier)	2.956^{**}	-3.914^{**}	–	-3.527^{**}	-2.746^{**}
Predicted resistance (lower-tier)	–	0.790^{*}	3.209^{**}	-0.535^{**}	1.445^{**}
Lagged antibiotic use (past 12 months)					
Any use (self)	0.633^{**}	1.371^{**}	1.745^{**}	1.743^{**}	0.624^{**}
Any use (higher-tier)	-1.344^{**}	0.213	–	0.393^{**}	0.014
Any use (lower-tier)	–	-0.369^{**}	1.234^{**}	0.044	0.498^{**}
Demographics					
Hospital 1 (MGH)	0.191^{**}	-0.402^{**}	0.703^{**}	-0.731^{**}	-0.613^{**}
Prior admission	-0.192^{**}	-0.367^{**}	0.178^{**}	-0.081^{**}	-0.056^{**}
Female, age 0–17	-1.530^{**}	-1.014^{**}	-1.161^{**}	-1.525^{**}	-2.221^{**}
Female, age 18–40	0.191^{**}	1.289^{**}	0.064	-0.461^{**}	-0.431^{**}
Female, age 41–64	-0.042^{*}	0.369^{**}	0.129^{**}	-0.214^{**}	-0.110^{**}
Female, age 65–79	-0.213^{**}	0.170^{**}	0.013	-0.316^{**}	-0.168^{**}
Female, age 80+	-0.524^{**}	-0.207^{**}	-0.078	-0.536^{**}	-0.310^{**}
Male, age 0–17	-1.486^{**}	-0.996^{**}	-0.928^{**}	-1.590^{**}	-2.155^{**}
Male, age 18–40	-0.029	0.259^{**}	0.009	-0.096^{*}	-0.011
Male, age 41–64	—	—	—	—	—
Male, age 65–79	-0.033	-0.180^{**}	0.072	-0.250^{**}	-0.040^{*}
Male, age 80+	-0.316^{**}	-0.475^{**}	-0.130^{*}	-0.648^{**}	-0.197^{**}
Non-English	-0.172^{**}	-0.321^{**}	-0.227^{**}	-0.079^{**}	-0.078^{**}
Hispanic	0.216^{**}	0.055	0.038	-0.061	-0.002
Non-white	-0.177^{**}	-0.048	-0.399^{**}	-0.162^{**}	-0.222^{**}
Culture Draw Sites (current admission)					
Abdominal site	-0.256^{**}	-0.206^{**}	0.636^{**}	-0.040	0.049
Blood	-0.614^{**}	0.282^{**}	0.954^{**}	-0.130^{**}	0.800^{**}
JP drain	0.690^{**}	0.843^{**}	1.649^{**}	1.197^{**}	1.093^{**}
Lung	0.006	-0.030	0.792^{**}	0.220^{**}	0.527^{**}
Nasal	0.331^{**}	0.474^{**}	0.684^{**}	0.121^{**}	0.650^{**}
Orthopedic	1.337^{**}	0.749^{**}	2.322^{**}	0.587^{**}	2.238^{**}
Other	-0.351^{**}	-0.332^{**}	0.880^{**}	0.393^{**}	0.785^{**}
Rectal	0.151^{**}	-0.181^{**}	0.171^{**}	0.371^{**}	0.193^{**}
Skin/soft tissue	0.479^{**}	0.357^{**}	1.285^{**}	0.231^{**}	0.892^{**}
Urine	0.256^{**}	0.281^{**}	0.360^{**}	0.391^{**}	0.372^{**}
Wound	0.983^{**}	1.036^{**}	2.225^{**}	0.763^{**}	1.877^{**}
Diagnosis Codes (current admission)					
Neoplasm / leukemia	0.533^{**}	0.185^{**}	0.152^{**}	0.480^{**}	0.167^{**}

	Cephalosporin	**Clindamycin**	**Linezolid**	**Trim/Sulfa**	**Vancomycin**
Cystic fibrosis	-0.888^{**}	-0.721^{**}	1.877^{**}	0.754^{**}	0.174^{*}
Catheter-related infection	1.027^{**}	0.611^{**}	1.063^{**}	0.360^{**}	0.988^{**}
Pneumonia	-0.437^{**}	-0.105^{*}	0.869^{**}	-0.003	0.740^{**}
Fever	0.171^{**}	-0.105^{**}	0.639^{**}	0.091^{**}	0.498^{**}
Abdominal pain / peritonitis	-0.298^{**}	-0.461^{**}	-0.068^{*}	-0.375^{**}	-0.347^{**}
Postoperative infection	0.684^{**}	0.794^{**}	1.394^{**}	0.551^{**}	1.107^{**}
Urinary tract infection	-0.238^{**}	-0.307^{**}	0.767^{**}	0.873^{**}	0.096^{**}
Bacteremia / sepsis	-0.389^{**}	0.075	2.040^{**}	-0.351^{**}	1.348^{**}
Cellulitis / abscess	0.346^{**}	0.899^{**}	1.696^{**}	0.470^{**}	1.296^{**}
Heart failure	-0.029	-0.344^{**}	0.415^{**}	-0.266^{**}	0.295^{**}
End-stage renal disease	-0.073	-0.227^{*}	0.679^{**}	0.604^{**}	0.288^{**}
Transplant status	-0.408^{**}	0.170^{*}	0.590^{**}	1.660^{**}	0.281^{**}
Bone marrow transplant (BMT)	-0.543^{**}	0.619^{**}	0.660^{**}	0.338^{**}	0.641^{**}
Liver transplant	-0.038	-0.731^{**}	0.779^{**}	0.872^{**}	0.090
Diagnosis Codes (past 12 months)					
Neoplasm / leukemia	-0.326^{**}	-0.155^{**}	-0.113^{**}	-0.097^{**}	-0.181^{**}
Cystic fibrosis	-0.054	-0.199	0.049	-0.395^{**}	-0.121
Catheter-related infection	-0.016	0.044	-0.047	0.172^{**}	0.016
Pneumonia	-0.285^{**}	-0.012	-0.112^{**}	0.182^{**}	-0.041
Fever	-0.067^{*}	-0.008	-0.028	0.109^{**}	-0.110^{**}
Abdominal pain / peritonitis	-0.205^{**}	-0.131^{*}	-0.214^{**}	-0.122^{**}	-0.172^{**}
Postoperative infection	0.141^{*}	0.106	-0.108	-0.124	-0.071
Urinary tract infection	-0.152^{**}	-0.248^{**}	0.042	-0.262^{**}	-0.080^{**}
Bacteremia / sepsis	-0.074	-0.163	-0.111	0.214^{**}	0.171^{**}
Cellulitis / abscess	0.083^{*}	0.072	-0.235^{**}	-0.084^{*}	-0.153^{**}
Heart failure	-0.171^{**}	0.071	-0.304^{**}	-0.027	-0.149^{**}
End-stage renal disease	0.026	0.117	-0.336^{**}	-0.163^{**}	0.114^{**}
Transplant status	-0.032	-0.106	-0.376^{**}	-0.409^{**}	-0.205^{**}
Bone marrow transplant (BMT)	0.019	-0.253	-0.291^{*}	-0.409^{**}	-0.167^{*}
Liver transplant	0.173	0.871^{**}	-0.096	-0.132	0.200
Positive Cultures by Pathogen (past 12 months)					
Coag. neg. *Staphylococcus*	-0.025	0.011	0.389^{**}	-0.022	0.166^{**}
E. coli	0.112^{**}	-0.106	-0.012	-0.261^{**}	-0.158^{**}
Enterococcus spp.	0.004	0.052	0.653^{**}	-0.051	0.142^{**}
K. pneumoniae	-0.007	-0.188	0.155^{**}	-0.279^{**}	-0.122^{**}
P. aeruginosa	-0.139^{*}	-0.139	0.149^{*}	-0.043	-0.044
S. aureus	0.068	-0.031	0.619^{**}	-0.068	0.373^{**}
Streptococcus spp.	0.130^{*}	0.182^{*}	0.107	-0.161^{*}	-0.010
Other gram-negative	-0.121^{**}	-0.313^{**}	0.048	-0.195^{**}	-0.108^{**}
Other gram-positive	-0.039	0.266	0.395^{*}	0.994^{**}	0.495^{**}
Positive Cultures by Specimen Type (past 12 months)					
Abdominal	0.087	-0.208	-0.216^{*}	-0.057	-0.283^{**}
Blood	-0.044	0.004	-0.368^{**}	-0.121	-0.299^{**}
JP drain	-0.083	0.109	-0.073	-0.092	-0.091

	Cephalosporin	Clindamycin	Linezolid	Trim/Sulfa	Vancomycin
Lung	-0.165^{*}	0.236^{*}	-0.369^{**}	0.325^{**}	0.070
Nasal	-0.085	0.008	-0.340^{*}	0.240	-0.048
Orthopedic	-0.082	0.580^{**}	-0.072	0.171	0.172^{**}
Other	0.156	0.097	-0.082	0.127	-0.002
Rectal	-0.234^{**}	0.459^{**}	-0.101	0.284^{**}	0.006
Skin/soft tissue	0.044	0.017	-0.052	-0.128^{*}	-0.048
Urine	-0.180^{**}	-0.058	-0.166^{**}	-0.310^{**}	-0.145^{**}
Wound	-0.001	-0.106	-0.195^{**}	0.027	-0.053
Constant	-1.887^{**}	-2.812^{**}	-7.919^{**}	-1.763^{**}	-1.692^{**}

Notes: Table reports conditional logit estimates of the antibiotic choice model (Equation 6) using 386,498 admissions across two hospitals. The dependent variable equals one if antibiotic j was chosen for a given isolate. "Cephalosporin" includes cefazolin and cephalexin. Covariates include hospital fixed effects, patient demographics, culture draw sites, diagnosis codes during the current admission and in the prior 12 months, and indicators for prior positive cultures by pathogen and specimen type. Significance levels are denoted by $^{*}p < 0.05$, $^{**}p < 0.01$, and $^{***}p < 0.001$. Standard errors are clustered at the patient level.

D.2 Resistance Model – Supplementary Methods and Exhibits

D.2.1 Resistance model – Benchmark 2SLS Estimates.

This section reports the alternative two-stage least squares specification used to estimate the hospital-level resistance model in Section 5. This specification avoids using dynamic panel estimators by re-specifying Equation 8 without the AR(1) term ($\delta\alpha_{hyj-1}$), instead replacing it with drug-year fixed effects to allow for drug-specific time trends. The 2SLS specification and accompanying first stage are:

$$\alpha_{hjy} = \tilde{\gamma}\tilde{u}_{hjy} + b_{hjy} + b_{hj} + \nu_{hjy} \tag{24}$$

$$\tilde{u}_{hjy} = \theta\tilde{A}_{hjy} + c_{hjy} + c_{hj} + \varepsilon_{hjy}, \tag{25}$$

where $\tilde{u}_{hjy}$ denotes the utilization share for antibiotic j in hospital h and year y.

In the first-stage equation, hospital-level antibiotic utilization $\tilde{u}_{hjy}$ is instrumented using the ASP policy changes $\tilde{A}_{hjy}$. The a_{hj} term represents hospital-drug fixed effects and adjusts for fixed differences in resistance across different drug-hospital pairs.[54] Lastly, η_{hjy} is the error term.

D.3 Decomposing Resistance Into Externalities and Individual-level Effects.

This section explains how we use our structural model to decompose the resistance impacts of the two ASP policy changes implemented in Hospital 1 into hospital- and individual-level effects.

At the individual level, the decline in linezolid use affects resistance rates by decreasing both the share of individuals that used linezolid in the past 12 months as well as the share of individuals

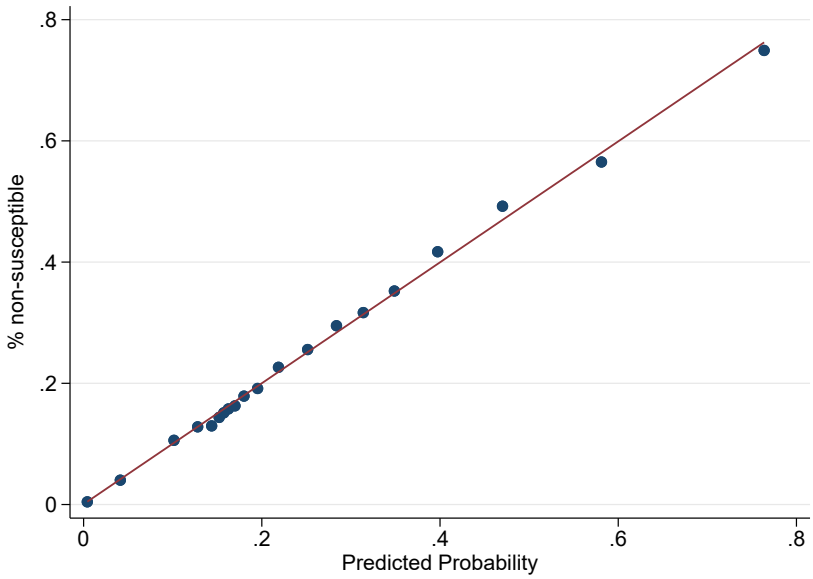
[54]For example, for a given drug, utilization or resistance may be persistently higher in one hospital compared to the other. The hospital-drug fixed effects account for these differences and ensure that the identification of the utilization effect γ is driven only by changes over time occurring within each drug-by-hospital combination.

resistant to linezolid in the past 12 months. Comparing Hospitals 1 and 2 in the pre-policy period (2008-2012) to the period after resistance reaches its new steady state (2015-2019), we find that ASP policy changes caused a 17.5% decline in linezolid use and 69.5% decline in linezolid resistance in Hospital 1 relative to Hospital 2. Applying these changes to the individual-level resistance model in (7) yields an upper bound for the effect of ASP policies on resistance, because it assumes that any changes to use or resistance at Hospital 1 is immediately reflected in all future admissions. In reality, patients discharged from Hospital 1 are unlikely to be re-admitted to the same hospital; in other words, hospitals rarely *directly* internalize the effects of use and resistance changes.

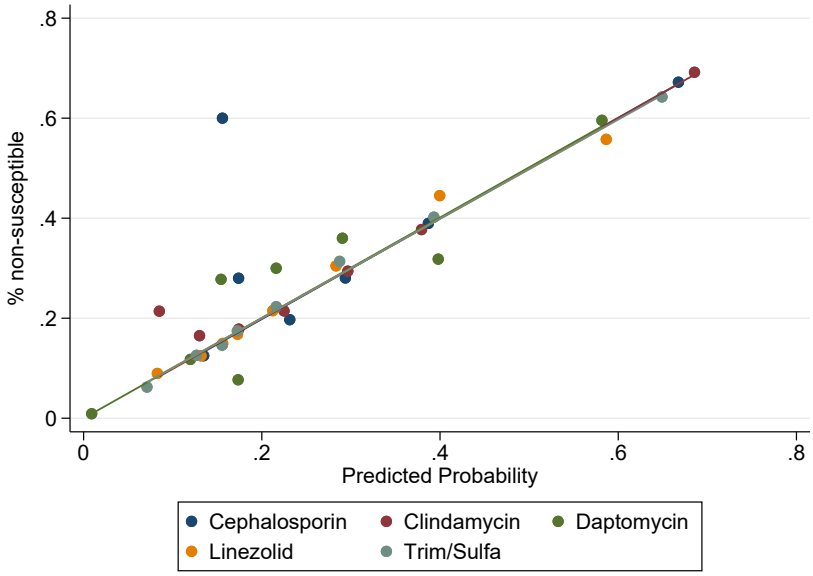
As shown in Table A5, individual-level effects can account for at most 24% of the total observed decline in linezolid resistance caused by the ASP policy, implying that at least 76% of the effect of ASP policies on resistance operates through externalities.

This suggests that most of the effect of utilization on resistance operates via hospital-level externalities. Importantly, this does not mean that individual-level effects are not important, but rather highlights the fact that existing linezolid utilization and resistance are concentrated among a small subset of individuals. For instance, less than 15% of isolates were from individuals who had used linezolid in the past 12 months. However, among the 1.6% of isolates that are actually resistant to linezolid, 61% were from individuals who had used linezolid in the past 12 months. Indeed, for individuals with prior linezolid use and resistance, zeroing out their individual-level effects would reduce their current chances of testing resistant to linezolid by 84%.

Figure A4: Resistance Model Fit



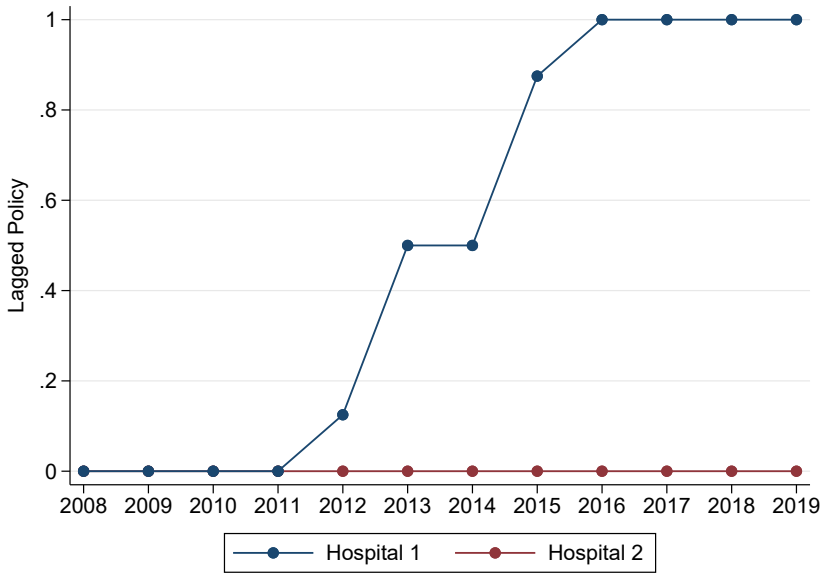
(a) All antibiotics



(b) Resistance by Antibiotic

Notes: Figure shows the fit of the individual-level resistance model (from Step 1 of the resistance model estimation procedure). Panel (a) shows a binned scatterplot of the actual share of non-susceptible isolates compared to the model-predicted probabilities for each isolate. Panel (b) shows the same binned scatterplot, separately by antibiotic.

Figure A5: Lagged ASP Policy Definition



Notes: Figure shows the average lagged ASP policy variable (Y-axis) for the antibiotic linezolid over time, separately for the treatment hospital (Hospital 1) and control hospital (Hospital 2). The lagged ASP policy variable is defined for each hospital-drug-year as the average over the preceding year of the non-lagged ASP policy variable, which in turn takes value 0 from January 2008 through September 2011, 0.5 from October 2011 through March 2014, and 1 from April 2014 through December 2019. These transition points correspond to the implementation of the two ASP policies.

Table A4: Descriptive Statistics (Resistance)

		Current	Past Resistance			Past Use		
Drug	N	Resistance	Own	Higher	Lower	Own	Higher	Lower
Cephalosporin	9,422	0.52	0.15	0.0648		0.26	0.14	
Clindamycin	81,982	0.40	0.15	0.016	0.097	0.083	0.13	0.21
Trim/Sulfa	250,279	0.23	0.16	0.022	0.13	0.19	0.13	0.16
Vancomycin	43,082	0.24	0.11	0.0066	0.15	0.37	0.058	0.19
Linezolid	31,438	0.012	0.0064		0.12	0.11		0.23

Notes: Table shows descriptive statistics for the resistance model dataset. Column 2 shows the number of isolates obtained for each drug. Column 3 shows current own-drug resistance. Columns 4–6 show past resistance for own-drug, higher-tier, and lower-tier antibiotics. Columns 7–9 show past use for own-drug, higher-tier, and lower-tier antibiotics.

Table A5: Effect of Hospital and Individual-level Utilization on Resistance

	Control Hospital	Treatment Hospital
Actual Resistance	0.008	0.025
Model-Predicted Baseline	0.008	0.025
With ASP Policy:	0.008	0.008
Individual-level Effect (Upper bound)		0.004
Hospital-level Effect (Lower bound)		0.013
Total Policy Effect		0.017

Notes: Table shows the decomposition of ASP policy effects on resistance into individual-level and hospital-level effects. Row 1 shows actual linezolid resistance in 2008-2012, the year prior to the effects of the two ASP policies, separately by Treatment and Control hospital (Hospitals 1 and 2 respectively). Row 2 demonstrates that our baseline model predictions match the actual resistance. Row 3 shows resistance after subtracting the estimated effect of the ASP policy (Row 4) from Row 1. This row shows that the ASP policy reduced resistance in the Treatment hospital to match resistance levels in the Control hospital. Row 4 is obtained by comparing observed changes in resistance in the Treatment versus Control hospitals, comparing the pre-policy period of 2008-2012 to the period after resistance reaches its new steady state (2015-2019). Row 5 shows the upper bound of the individual-level effects of the ASP policy, calculated using the individual-level model estimates from (7). Row 6 shows the lower bound of the hospital-level effects of the ASP policy, calculated as Row 4 - Row 5.

D.4 Mortality Model – Supplementary Methods and Exhibits

D.5 Two-stage Least Squares (2SLS) Estimation.

The text in Panels A and B of Figure 7 reflect the following two-stage least squares (2SLS) regression specification.

$$M_i = \beta_0 + \beta_1 X_i + \mathbf{W}_i'\boldsymbol{\gamma} + \varepsilon_i \tag{26}$$

$$X_i = \pi_0 + \pi_1 \cdot (T_i \cdot Z_i) + \mathbf{W}_i'\boldsymbol{\theta} + \eta_i \tag{27}$$

The outcome M_i is a binary indicator for whether an admission resulted in mortality within 30 days. The endogenous regressor X_i is defined as the share of each admission occurring after linezolid is first administered, taking values in [0,1]. This measures whether linezolid was used during the admission as well as how quickly linezolid was used after admission. This measure takes value 1 for admissions where linezolid was given on the first day, 0.5 if linezolid was given halfway through the admission, and 0 if linezolid was never given. The instrument is the LOO linezolid approval rate Z_i interacted with the binary indicator T_i for whether the patient was in the top 10% of predicted linezolid use. The specification includes the vector of controls $\mathbf{W}_i$, which includes the main effects

of Z_i and T_i, age-by-sex bins,[55] and demographics (non-English speaking, non-White, Hispanic), and year fixed effects. We cluster standard errors at the level of the ASP pager holder on shift during each admission, resulting in 36 unique clusters.

Table A6: Descriptive Statistics (Mortality)

	(1) $T_i = 1$	(2) $T_i = 0$	(3) Total
Predicted linezolid use (range)	0.08–0.96	0–0.08	0–0.96
Linezolid use	0.17	0.01	0.030
30-day mortality	0.058	0.033	0.036
Linezolid LOO approval rate	0.73	0.73	0.73
Covariates			
Positive for VRE (share)	0.32	0.03	0.06
Length of stay (mean)	32	13	15
Bacteremia/Sepsis ICD code (share)	0.40	0.08	0.12
Age (mean)	59	56	56
Sex (share male)	0.54	0.52	0.52
Non-English (share)	0.13	0.10	0.11
Hispanic (share)	0.04	0.03	0.03
Nonwhite (share)	0.20	0.21	0.20
N admissions	2,353	16,564	18,917

Notes: Table reports descriptive statistics for the mortality model estimation sample, which is restricted to admissions between 2011 and 2016 that Column (1) shows admissions in the treated group ($T_i = 1$), defined as the top decile of predicted linezolid use. Column (2) shows admissions in the control group ($T_i = 0$), defined as the bottom 90% of predicted linezolid use. Column (3) reports the full sample. Variables include actual linezolid use, 30-day mortality, leave-one-out (LOO) linezolid approval rate, and patient covariates (VRE status, length of stay, sepsis/bacteremia diagnosis, demographics, and language/ethnicity indicators).

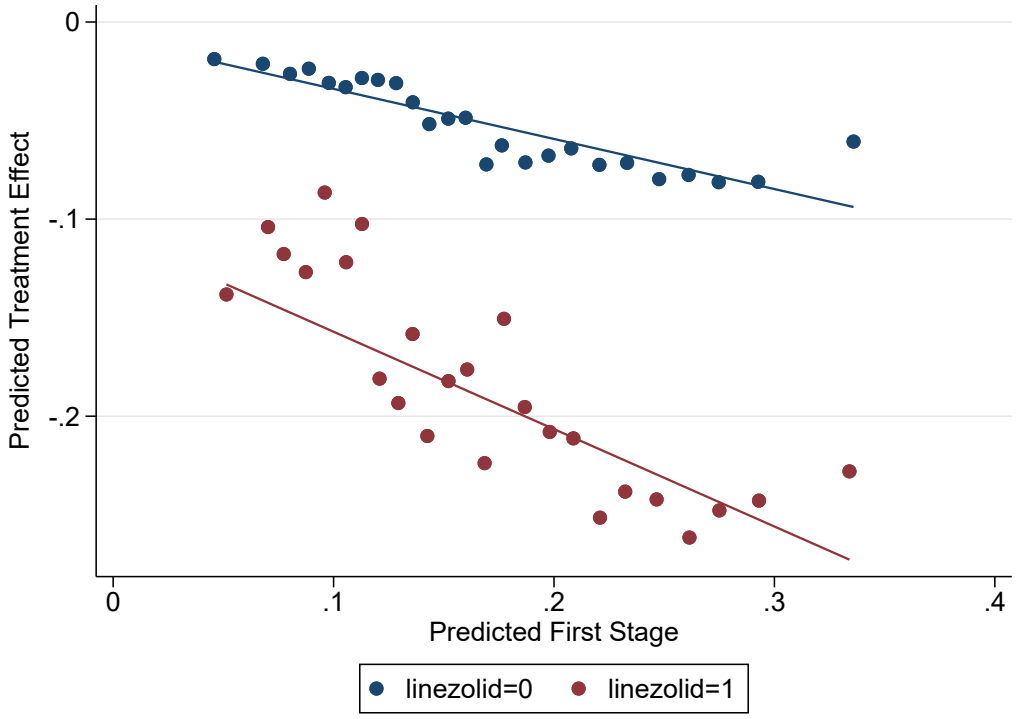
[55] Age bins were defined as 0-17, 18-40, 41-64, 65-80, and $\geq$80 years.

Table A7: Deciles of Predicted Linezolid Use

Decile of Predicted Use	N admissions	Actual Share with linezolid use
1–6	11,199	0.003
7	1,004	0.016
8	2,035	0.020
9	2,326	0.034
10 ($T_i = 1$)	2,353	0.170

Notes: Table shows deciles of predicted linezolid use. Column 1 reports the number of admissions in each group. Column 2 reports the share of admissions in each group with any linezolid use. Groups do not contain exactly 10% of the sample because predicted linezolid use takes on a finite number of unique values, and identical values must be allocated to the same group.

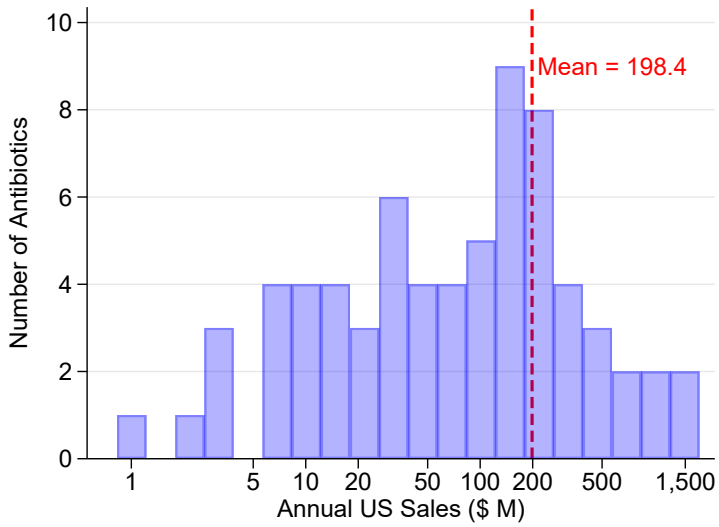
Figure A6: Mortality Model: Treatment Effects by Predicted Linezolid Use



Notes: Figure shows a binned scatterplot of model predictions from Equation 11. The Y-axis shows treatment effects, defined as predicted 30-day mortality with linezolid coverage minus predicted mortality without linezolid coverage, separately by whether linezolid was actually used in the data. The X-axis shows the model-predicted probability of linezolid use. For both groups, patients with higher predicted linezolid use are predicted to benefit more from linezolid. Patients observed to use linezolid in the data have larger predicted benefits.

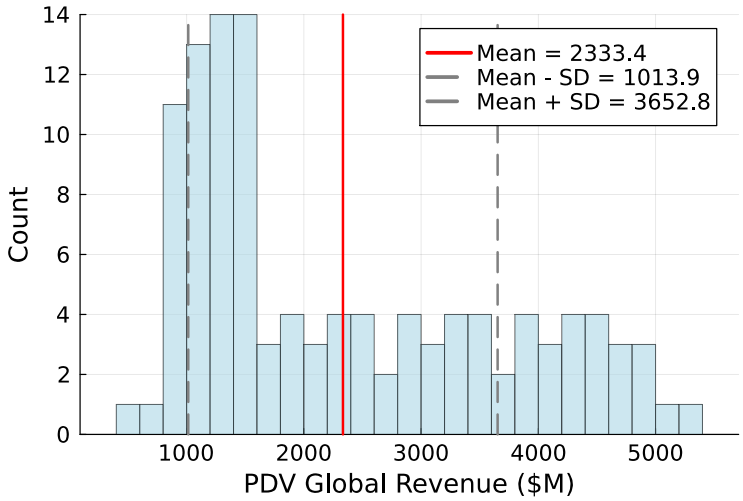
D.6 Innovation Model - Supplementary Exhibits

Figure A7: Distribution of Annual U.S. Antibiotic Sales.



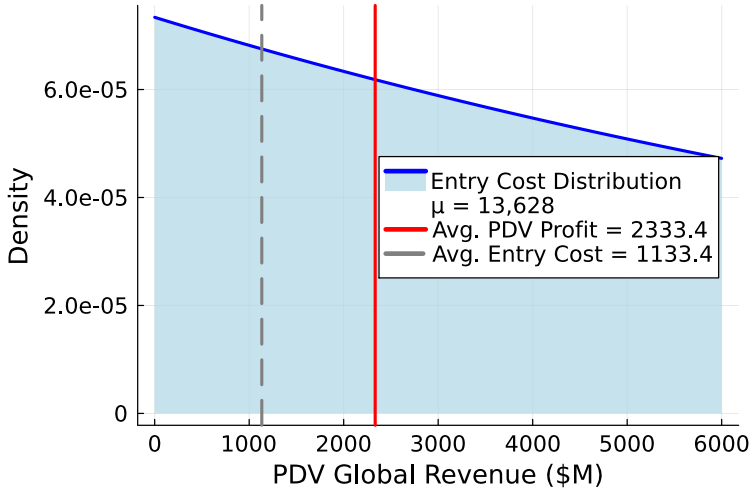
Notes: Histogram shows the distribution of annual branded U.S. sales across all N=69 antibiotics in the estimation sample (in millions of 2024 dollars). The dashed red vertical line marks the sample mean. Sales are plotted on a logarithmic scale for readability.

Figure A8: Innovation Model - Expected Profits



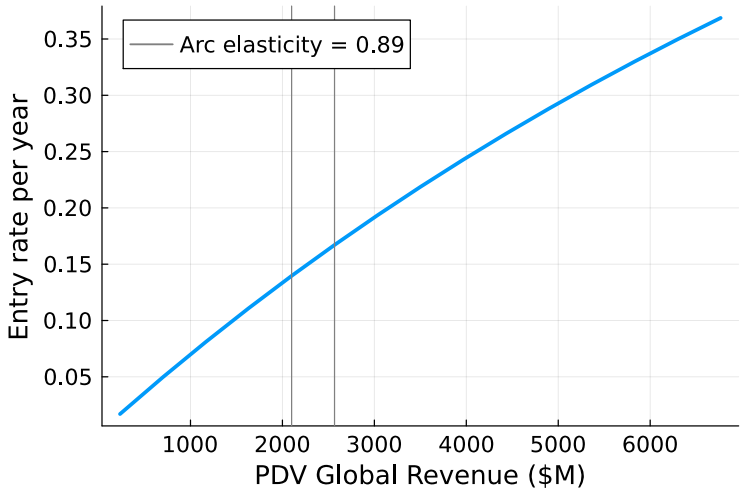
Notes: Figure shows the distribution of the present-discounted value of global revenues over 10 years of patent life, $VE(g,t)$, across antibiotic groups g approval years t, as defined in Section 5.6. The vertical red line denotes the average of $VE(g,t)$ in the sample. The vertical gray dashed lines represent the mean plus or minus the standard deviation. Sales are plotted on a linear scale.

Figure A9: Innovation Model - Entry Cost Distribution Estimates



Notes: Figure shows the exponential entry cost distribution with estimated mean $\mu = 13,628$. The vertical red line denotes the average present-discounted value of global revenues over 10 years of patent life. The vertical gray dashed line represents the average realized entry cost conditional on entry. All dollar values are in millions of 2024 dollars.

Figure A10: Entry Probability as a Function of Total Profit



Notes: Figure relates the entry probability for a single potential entrant in a given year as a function of the present-discounted value (PDV) of total profits. The plotted curve shows how the entry rate varies by the PDV of profits, separately for exponential and log-normal models of entry costs. Arc elasticities represent the local change in the entry rate for a 20% change in profits, calculated using the increase in the entry rate between the vertical gray lines.

D.7 Simulations - Supplementary Methods

Below, we describe in detail how we simulate each of the four parts of the model.

Demand and Resistance. We simulate demand and resistance jointly using a forward-simulation approach in one-month increments starting from January 2008. This setup captures the lagged structure of the resistance and demand equations, whereby current values depend on past (discrete) realizations of utilization and resistance. For each entry path e, we simulate $S = 10$ stochastic draws, indexed by s. This allows us to generate distributions of outcomes, accounting for both worst- and best-case scenarios. The following steps are repeated S times for each entry path $e \in \mathcal{E}$.

1. For the resistance model, we draw a set of uniformly distributed shocks $\varepsilon_{el}^{R} \sim U[0,1]$, one for each of the $N_R = 669{,}811$ culture isolates, indexed by l.

2. For the demand model, we draw Type-I extreme value errors, $\varepsilon_{ij}^{D} \sim \text{Gumbel}(0,1)$, one for each of the $N_D = 2{,}730{,}749$ admission-by-choice-alternative observations.

3. Given these simulation draws and our model estimates, we calculate demand $D_{ij}(t) \in \{0,1\}$ and resistance $R_l(t) \in \{0,1\}$ for all admissions and isolates occurring in January 2008.

4. We then propagate these simulated outcomes forward for the next 12 months (since both demand and resistance depend on prior values of resistance and demand in the past 12 months, excluding the current month).

5. We repeat steps 3-4 for February 2008, and continue month by month through the last month in the sample (December 2019).

For entry paths that include new entrants, we track demand and resistance separately for each entrant over time. New entrants in the market are assumed to have characteristics identical to linezolid, except that they begin with zero resistance. In the end, we obtain simulated resistance and demand values for every admission and isolate, separately for each entry path e and simulation iteration s. These simulated trajectories are used to calculate flow profits for each entry path e, which then feed into the innovation model to determine entry probabilities.

Mortality. Once demand and resistance outcomes have been simulated, we simulate the number of deaths occurring for each entry path e and simulation draw s, taking the set of admissions as fixed. [56] As with profits, we assume that mortality rates stop changing after period T. The simulated mortality rate over the sample window constitutes our measure of welfare: the favorability of counterfactual policies will depend on the degree to which they can lower long-run mortality. We simulate mortality over the admissions used in estimating the mortality model in Section 5.5. For each admission, we use the estimated mortality model to calculate the predicted mortality probability with and without linezolid coverage. An admission is considered covered ($u_{i,linezolid} = 1$) if linezolid was administered and no linezolid-resistant culture was drawn. Thus, the outputs of our demand and resistance models determine coverage for each simulation iteration.

[56] For example, if a patient was admitted twice in the original data, but the patient dies in their first admission for a subset of simulation draws, we do not selectively remove their second admission from the data. Similarly, we do not add new admissions for individuals who died in the data but may have survived in a subset of simulation draws.

Innovation. We assume that one potential entrant decides whether to enter in January 2009. This potential entrant draws an *iid* private entry cost from the estimated exponential entry cost distribution from Section 5.6. We assume that entrant prices are equal to the (inflation-adjusted) price of linezolid in 2007-2008 in Hospital 1 ($184 per day in 2024 dollars), each patient using linezolid is treated for 10 days, and marginal costs are zero. The new antibiotic enters the market if the present discounted value (PDV) of future profits (over a 10-year exclusivity period) exceeds the drawn entry cost.

D.8 Baseline Simulation Results and Comparative Statics

Figure A11 shows how simulated demand, resistance, and mortality change as a function of the degree of demand restrictions and whether a new antibiotic enters in 2009. For each panel, bars show comparative statics across different fixed scenarios. Reported outcomes are weighted averages from the treatment hospital (Hospital 1) across the entire sample period (2008–2019 for demand and resistance; 2011–2016 for mortality). Importantly, these comparative statics do not depend on the innovation model; they only depend on the estimated demand, resistance, and mortality equations.

Bar 1 of each figure panel shows baseline values under status quo policies and no antibiotic entry. Demand averages 3.19% of admissions, resistance averages 1.77% of isolates, and 3.97% of admissions result in mortality. Figure A12 verifies that the baseline forward-simulated demand and resistance dynamics match both the raw data and individual model predictions from (6) and (7).

Bar 2 removes demand restrictions, resulting in a 16% increase in demand to 3.69%, a 41% increase in resistance to 2.50%, and a modest decline in mortality of 0.04 percentage points (0.1%).

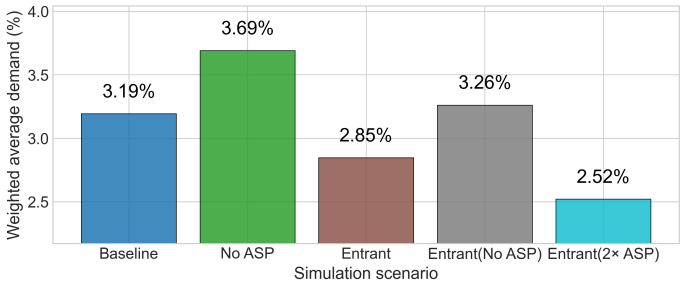
Bar 3 shows baseline policies with the addition of a novel antibiotic. The entrant partly substitutes for linezolid, resulting in a decline in linezolid use to 2.85%. Accordingly, linezolid resistance falls to 1.06%. However, because the entrant expands treatment to more individuals, mortality falls significantly, to 3.84%, a 3.3% relative decline.

Bar 4 shows that combining entry with removal of demand restrictions accentuates mortality benefits at the expense of an increase in resistance. The marginal effect of removing demand restrictions on mortality is larger with entry (-0.20 percentage points) compared to without entry (-0.13 percentage points), because demand restrictions apply to both linezolid and the entrant. While linezolid resistance rises to 1.43%, this remains lower than in the baseline case. That is, new antibiotic entry reduces resistance levels of incumbent antibiotics by substituting away some of their demand.

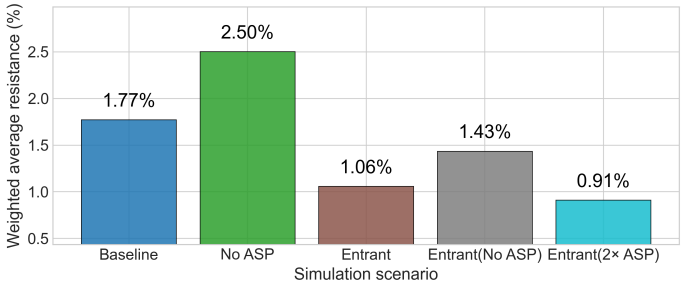
Finally, Bar 5 shows results where a new antibiotic enters but antibiotic demand restrictions are doubled. While this yields the lowest resistance (0.91%), it also reduces utilization of both antibiotics, undoing half of the mortality benefit of the entrant.

Figure A11: Simulated Comparative Statics

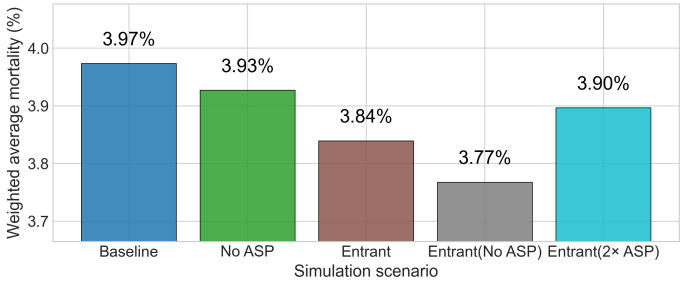
(a) Linezolid Demand



(b) Linezolid Resistance



(c) Mortality

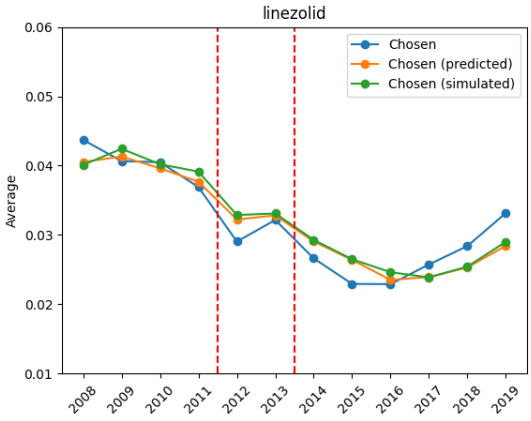


Notes: Figure shows comparative statics across different scenarios: the baseline simulation with status quo policies ("Baseline"), removal of the ASP policies ("No ASP"), the inclusion of a new antibiotic in the choice set ("Entrant"), the inclusion of a new antibiotic and removal of ASP policies ("Entrant (No ASP)"), and the inclusion of a new antibiotic coupled with ASP policies that are twice as strict ("Entrant (2x ASP)"). Panels (a), (b), and (c) show linezolid demand, resistance, and mortality respectively. Data are restricted to the treatment hospital (Hospital 1).

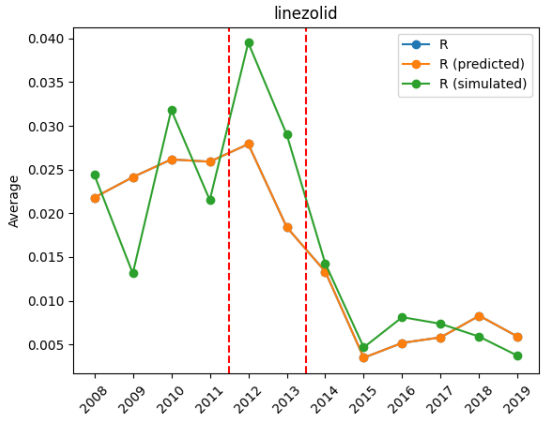
Overall, these results demonstrate how demand, resistance, and mortality respond to demand restrictions and new entry. Specifically, they show that new antibiotics yield greater mortality benefits when they are allowed to be used. Moreover, while higher resistance lowers antibiotic effectiveness for all individuals (including inframarginal users), this effect is outweighed by the direct mortality improvements among marginal antibiotic users, at least in a local region around existing policies. Here, we have examined outcomes conditional on entry decisions. In the following section, we consider the question of optimal policies when innovation is endogenous.

Figure A12: Model Fit: Linezolid Demand and Resistance in Hospital 1

(a) Demand

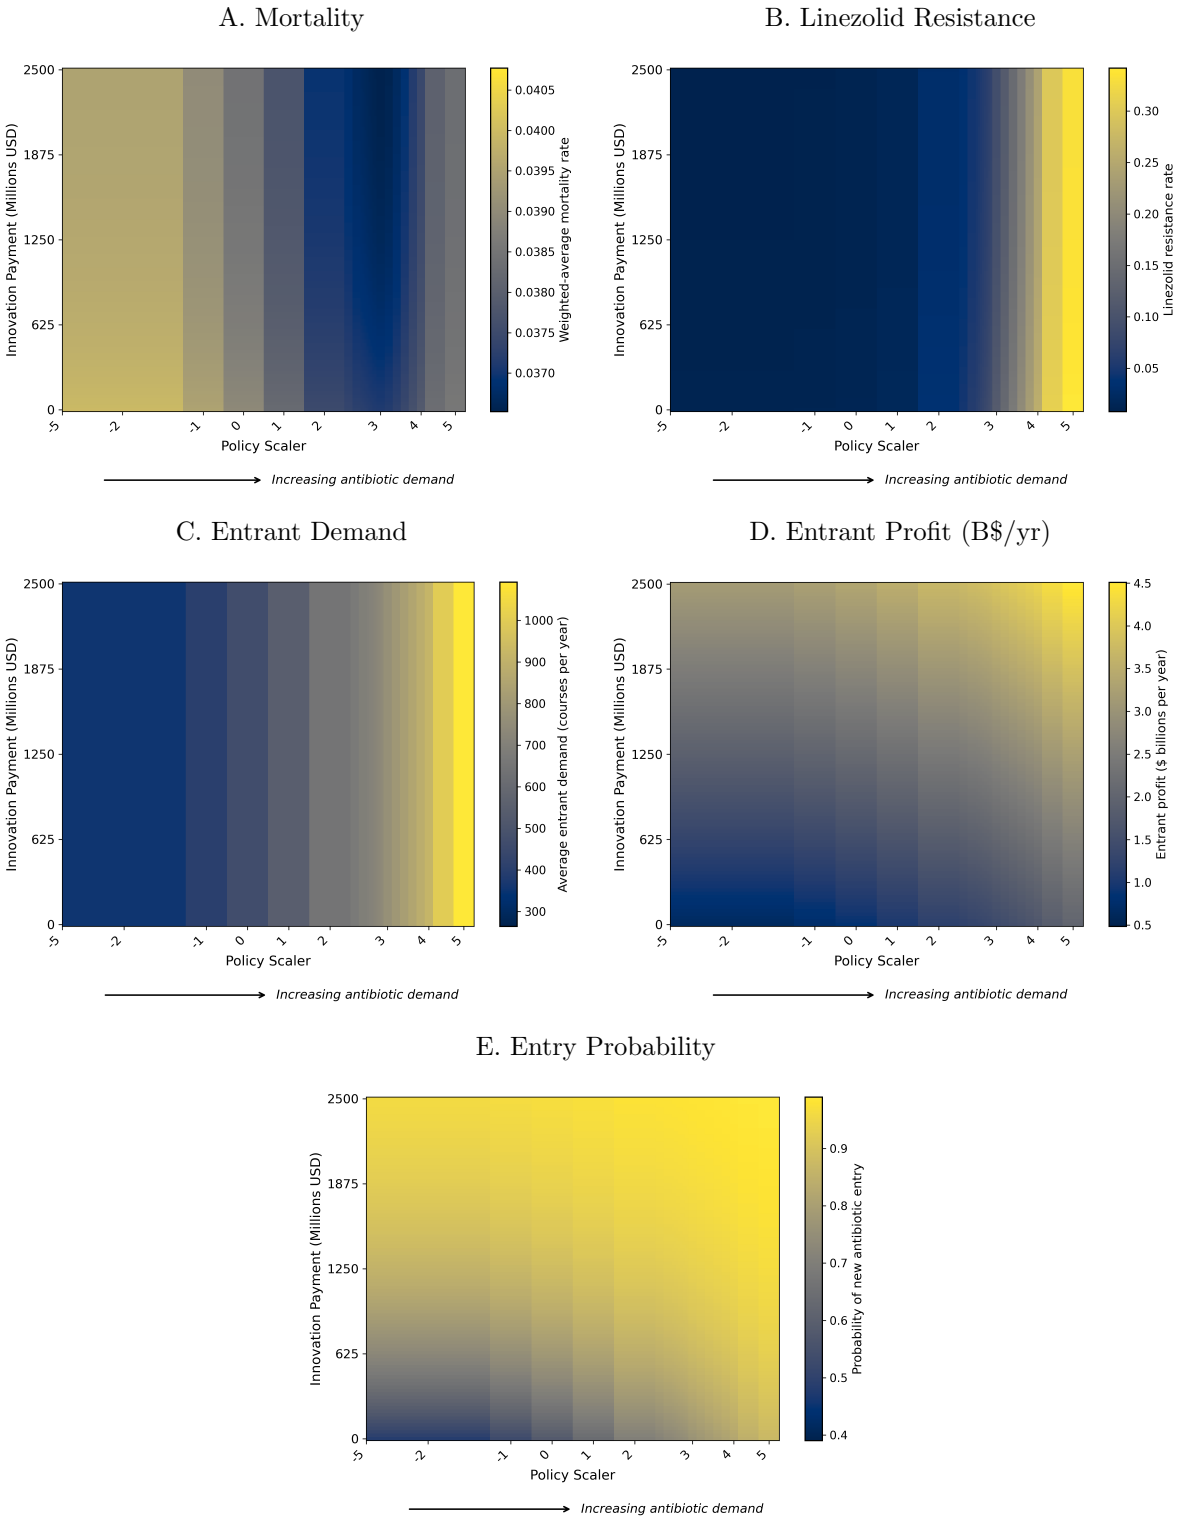


(b) Resistance



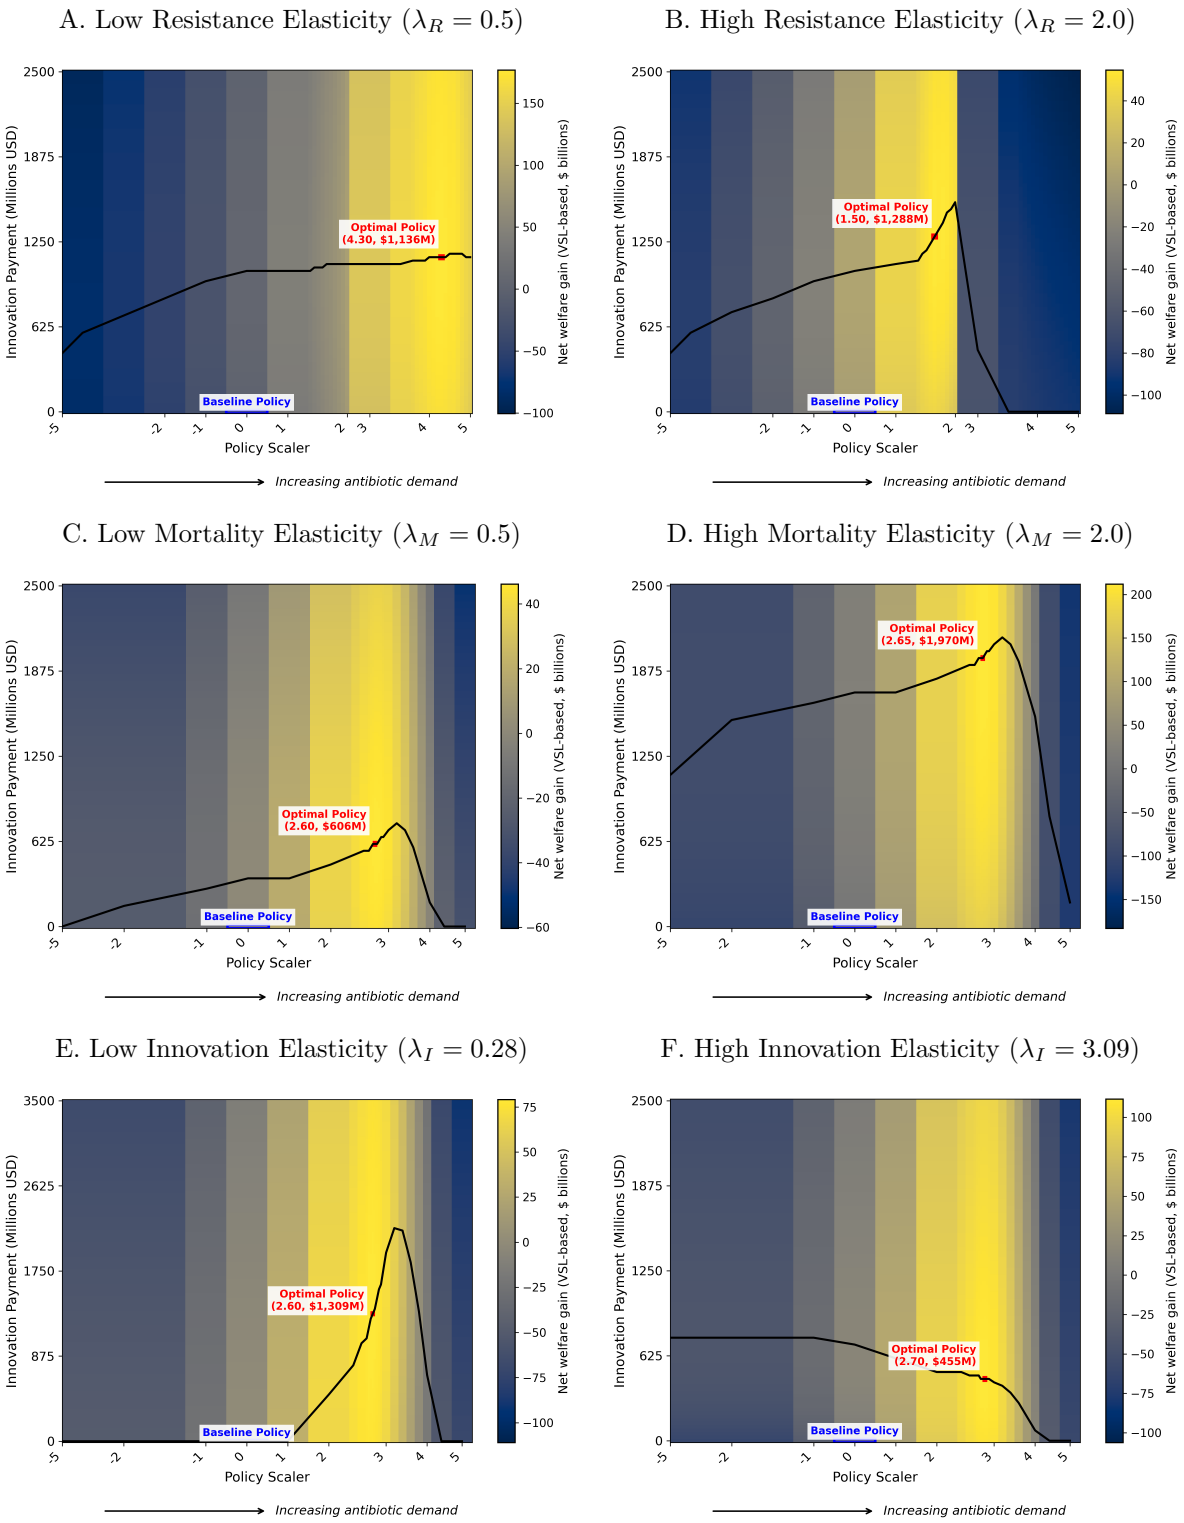
Notes: Figure demonstrates model fits for linezolid demand and resistance in the treated hospital (Hospital 1). For each outcome, blue, orange, and green lines show yearly averages of the actual data, model predictions from (6) and (7), and simulated demand from forward simulations, respectively. Panel (a) shows a close correspondence between all three data series. Panel (b) shows that the actual data and logit model predictions coincide exactly, with simulated resistance exhibiting more year-to-year variation. The red dashed vertical lines represent the implementation of the two ASP policies.

Figure A13: Policy Effects on Mortality, Demand, Entry, and Resistance



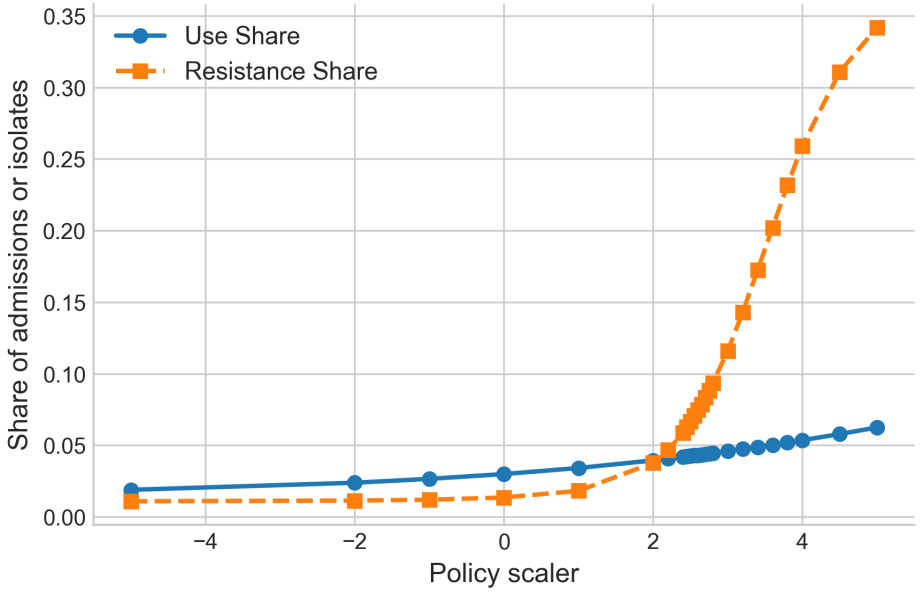
Notes: Each panel shows model-simulated outcomes across the two policy dimensions: the magnitude of the policy scaler (x-axis; higher values correspond to higher antibiotic demand) and the annual innovation payment (y-axis, in millions of U.S. dollars). Color intensity reflects the magnitude of each outcome.

Figure A14: Sensitivity Analyses - Welfare and Optimal Policies



Notes: Figure shows PDV of welfare, relative to the baseline scenario, across two policy dimensions: the magnitude of the policy scaler (x-axis; higher values correspond to higher antibiotic demand) and the annual innovation payment (y-axis, in millions of U.S. dollars). Panels A and B vary the response of resistance to demand. Panels C and D vary the effect of antibiotic coverage on mortality. Panels E and F vary the innovation elasticity to expected profits.

Figure A15: Demand and resistance as a function of the policy scaler



Notes: Figure plots the average share of admissions treated with linezolid ("Use Share") and the share of isolates resistant to linezolid ("Resistance Share") as a function of the policy scaler. The policy scaler indexes the intensity of demand-side interventions, where higher values represent higher demand or lower restrictions. Both series are computed from model simulations without innovation prizes. Shares represent averages across the sample period (2008-2019).

Table A8: Optimal Innovation Payment by Policy Scaler (VSL Gain Relative to Baseline)

Policy Scaler	Innovation Payment (B$/yr)	Annual Profit (B$/yr)	Entry Probability	National VSL Gain (B$)
−5.0	0.429	0.916	0.606	−90.73
−2.0	0.833	1.484	0.779	−47.57
−1.0	0.960	1.701	0.823	−22.40
0.0	1.035	1.893	0.854	6.47
1.0	1.035	2.041	0.875	42.60
2.0	1.136	2.326	0.906	82.20
2.2	1.212	2.443	0.917	90.31
2.4	1.237	2.512	0.922	95.12
2.6	1.288	2.607	0.930	**99.59**
2.8	1.338	2.704	0.936	98.05
3.0	1.389	2.803	0.942	96.67
4.0	0.859	2.543	0.925	0.81
5.0	0.000	2.010	0.871	−64.06

Notes: Table reports simulated outcomes from the forward-simulation model under varying values of the policy scaler. The policy scaler modifies the intensity of antibiotic stewardship restrictions; higher values correspond to looser restrictions. *Innovation Payment* is the annual incentive offered to induce antibiotic market entry, expressed in billions of U.S. dollars per year. *Annual Profit* denotes the expected average annual global profit of entrants, in billions of dollars per year. *Entry Probability* is the probability of a novel antibiotic entering the market. *National VSL Gain* reports the net present value of national welfare relative to the baseline case (policy scaler = 0, innovation payment = 0), combining mortality and innovation costs, in billions of U.S. dollars. The bolded entry denotes the welfare-maximizing policy combination.