Optimal Policy with Opposing Externalities: Evidence from Antibiotics

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

December 20, 2025

Preliminary draft — Comments welcome; please do not cite or distribute.

Please click here for the most recent version.

Abstract

Antibiotic use generates two opposing externalities: it increases antimicrobial resistance, reducing the effectiveness of existing drugs, but also supports the discovery of new ones. Policies that restrict antibiotic use may slow the emergence of resistance but weaken incentives for pharmaceutical innovation. Direct innovation payments have been proposed to decouple revenues from use, yet their optimal magnitude and welfare implications remain unclear. These opposing forces complicate optimal policy design, which depends on the effects of antibiotic demand on resistance, mortality, and innovation. Using electronic health records from two large medical centers and quasi-experimental variation in antibiotic restrictions, we find that antibiotic use reduces mortality for treated patients but also generates hospital-wide resistance externalities. Embedding these estimates in a structural model of antibiotic demand, resistance, mortality, and new antibiotic entry, we quantify the welfare effects of optimal demand incentives and innovation prizes. In the U.S. alone, a new antibiotic reduces annual infection-related mortality by 2,700, and increases the present-discounted value (PDV) of welfare by $54 billion. An optimal innovation prize of $1.04 billion per year reduces annual mortality by 759, and raises the PDV of welfare, net of prize costs, by $6.5 billion. Innovation prizes are complementary with greater utilization: jointly optimal prizes and utilization avert over 4,500 deaths per year and increase the PDV of welfare by $100 billion. More generally, opposing externalities imply that multiple policy levers improve welfare.

*Harvard University; edwardkong@hms.harvard.edu. We thank Mark Shepard, David Cutler, Tim Layton, Ariel Pakes, Jeffrey Pearson, and David Kubiak for valuable feedback. I thank 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. In response, antibiotic stewardship programs (ASPs), which restrict and guide antibiotic prescribing within hospitals, have become central to infection control policy 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. 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 (Chapagain and Poudyal, 2020; Belant et al., 2025; Lindsey et al., 2007). Rent control may improve short-run housing affordability but decrease long-run investments in the quantity and quality of housing (Diamond et al., 2019; Autor et al., 2014). A particularly salient example is that of climate change, where rising global temperatures parallel rising bacterial resistance and regulations on economic production mirror restrictions on antibiotic use. Reducing economic output and other decarbonization efforts aim to avert large future damages (Bilal and Känzig, 2024; Nordhaus, 2018; Fankhauser and S.J. Tol, 2005), but involve immediate costs (Dasgupta et al., 2002; Glavina et al., 2025), and thus continue to be debated (Stern and Stiglitz, 2023).

While it is common for economic policies to incur tradeoffs, these examples bear the common trait that expanding the space of available policies can improve welfare by decoupling revenues from sales volume. Also, in all of these cases, determining the optimal combination of policies requires quantifying the key opposing externalities and their joint contributions to long-run welfare.

In the case of antibiotics, the three key elasticities describe how antibiotic use affects resistance, how use and resistance affect mortality, and how innovation responds to market size. To estimate the impact of antibiotics on resistance and mortality, 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. We use two sources of identifying variation: for resistance, we

use hospital-specific changes in antibiotic restriction policies. For mortality, we leverage variation in antibiotic approver staffing.

First, we identify the causal effects of stewardship policies on antibiotic demand and resistance using difference-in-difference and triple-difference frameworks leveraging changes in antibiotic restriction policies. Our reduced-form estimates demonstrate that antibiotic restriction policies reduce use by about 20%, which translates into a 60% decline in resistance, with at least 76% of this effect operating via hospital-wide resistance externalities rather than effects on individual patients.

Next, we recover the causal effect of antibiotic coverage (defined as antibiotic use among non-resistant individuals) on mortality. In our setting, the identity of the provider in charge of approving or denying requests for restricted antibiotic changes every week. We use provider staffing data and variation in antibiotic approval rates across different antibiotic approvers to generate quasi-random variation in antibiotic use, in what is commonly known as a judge or examiner design (Kling, 2006; Goldsmith-Pinkham et al., 2025; Chyn et al., 2025). For patients in the top decile of predicted use, receiving antibiotic coverage 10 percentage points earlier reduces 30-day mortality by 6.4 percentage points.

To build intuition for how hospital objectives may differ from optimal policy, we introduce a simple graphical framework of optimal antibiotic use. Absent innovation prizes, optimal use equates the marginal costs of increased antibiotic resistance to the marginal benefits of treating current infections, plus the benefits of increased innovation. Innovation prizes decouple demand and innovation but also lower overall resistance; hence, whether innovation prizes imply lower optimal antibiotic demand is theoretically ambiguous.

To solve for optimal-based policies and innovation prizes, we construct a structural model of antibiotic use, resistance, innovation, and mortality. In the model, hospitals set restrictions on physicians' antibiotic demand taking into account patient characteristics, treatment effects, antibiotic costs, and hospital-wide resistance. Resistance evolves endogenously as a function of antibiotic use. Pharmaceutical firms observe expected future revenues, net of resistance-driven obsolescence, and decide whether to enter the market. The social planner minimizes welfare losses from long-run mortality and innovation prize costs. We identify the key parameters of the model using the identification strategies from our reduced-form analyses. To incorporate endogenous antibiotic entry, we evaluate counterfactuals across the range of published elasticities, spanning 0.25 to 2.75 (Dubois et al., 2015). In our preferred specification, we combine FDA approval data with revenues for antibiotics between 1986 and 2024 to estimate an innovation elasticity of 0.89, following Pakes et al. (2007). Varying the innovation elasticity scales our welfare estimates but does not change our qualitative conclusions.

We show that the entry of a single new compound yields a 2.8% reduction in mortality and increases long-run welfare by $54 billion in the U.S. alone. To induce entry, the optimal innovation prize awards the innovator $1.04 billion per year for 10 years and raises net long-run welfare by $6.5 billion. Demand-side policies, such as changes to stewardship restrictions, taxes, or subsidies, have even larger effects. Optimal demand incentives *increase* antibiotic use and raise resistance, but

resistance increases are outweighed by reductions in mortality and increased innovation, such that long-run welfare increases by $88.6 billion. Allowing for both policy instruments yields the greatest welfare increases due to complementarity between them: the optimal policy combination raises long-run welfare by $99.6 billion. Efforts to suppress resistance too aggressively can potentially raise mortality and dampen innovation, reinforcing the need for policies that balance conservation with access and discovery. Our main conclusion that demand and innovation are complements is robust to sensitivity analyses across a wide range of resistance, mortality, and innovation elasticities.

Our paper makes several contributions. By combining two separate identification strategies, we provide, to our knowledge, the first causal estimates of how antibiotic stewardship practices affect both hospital-wide resistance externalities and short-run health outcomes for individual patients. Together, these estimates clarify the tradeoffs inherent in regulating antibiotic use and allow us to quantify the welfare benefits of a novel antibiotic in terms of lives saved. We further contribute to the theoretical literature on optimal antibiotic policy by jointly estimating and modeling resistance and innovation externalities alongside direct mortality effects. Embedding these estimates in a structural framework, we deliver new evidence of complementarities between demand-based policies and innovation prizes—two widely used or proposed policy tools.

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

Empirical work estimating these factors is rare. Recent examples include Adda (2020), who identifies the causal effect of human antibiotic use on resistance using temperature-driven variation in prescribing. Dubois and Gokkoca (2023) investigates the opposite causal direction, finding that

resistance reduces use, controlling for endogeneity using antibiotic sales in veterinary medicine. Using a structural model, they find that decreasing either veterinary or human antibiotic use reduce resistance, restricting human antibiotic use reduces consumer surplus, and improving rapid antibiotic susceptibility testing is beneficial. We add to this work by providing the first estimates of resistance externalities in the hospital setting, where most novel broad-spectrum antibiotics are first introduced. We combine our resistance model with mortality estimates to compute welfare effects in terms of lives saved, and additionally assess innovation prizes. Outside economics, a large epidemiological literature has developed calibrated models of resistance dynamics under alternative utilization regimes (Niewiadomska et al., 2019), but typically abstracts from the economic incentives governing antibiotic development and use.

Estimates of the effect of antibiotic restrictions on health outcomes 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. 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 decreased mortality by 14% (O'Brien et al., 2024).

To help specify 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).

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 describes the electronic health records data and other datasets. Section 4 shows reduced-form evidence for the important elasticities in the model. Section 5 describes our conceptual framework and gives a graphical illustration of the relationships between antibiotic use, resistance, innovation, and mortality. Section 6 specifies the full structural model, estimation procedures, and model estimates. Section 7 uses the estimated structural model to simulate welfare implications for alternative demand-targeted policies and innovation prizes. Section 8 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 additional healthcare costs and lost productivity in the U.S. alone (Centers for Disease Control and Prevention, 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 will tend to decline in the absence of continued use. In accordance with the literature on hospital-acquired infections (Lipsitch, 2001), 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 by CMS (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 typically govern use for 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 (Majewska, 2022; Kong and Zhao, 2025) 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 demand, which is often 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 tradeoff: by reducing consumption they slow resistance, but they may also 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. That approach has yielded diminishing returns; as rediscovery of known compounds became common and the rate of discovery of new scaffolds

declined, recent innovation has shifted toward synthetic compounds and combination therapies that pair antibiotics with adjuvants targeting resistance mechanisms (Brown and Wright (2016); Lewis (2013), and 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 (Walsh and Wencewicz, 2014; Tommasi et al., 2015). 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, we focus on 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. This creates markedly different incentives compared to outpatient drugs, which are reimbursed separately by payors. The fixed bundled payment amount effectively serves as an upper bound for any input costs (including drugs) incurred by the hospital. For example, the average Medicare payment for pneumonia (DRG 194) was about $5,084 in FY2021.[1] 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 demand; in our counterfactual simulations, we explore how changes in demand could affect long-run welfare. 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.

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 via distinct economic mechanisms. The DISARM Act

[1] From the Optum 2021 DRG National Average Payment Table.

would enhance Medicare reimbursement for qualifying antibiotics—effectively acting as a demand-side incentive by increasing hospitals' willingness to use novel agents—whereas the PASTEUR Act would create subscription-style pull contracts worth up to $6 billion in total funding to reward developers of critical antimicrobials 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 treatment effectiveness accounting for resistance dynamics as well as demand for novel agents. 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* species, 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 *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; and (iii) linezolid resistance is the least rare 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 Data

3.1 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.[2] 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. The EHR data also allow us to control for rich sets of covariates, including age, sex, demographics, and admission-level characteristics (diagnoses,[3] prior admissions,

[2]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.

[3]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,

prior antibiotics, prior positive cultures,[4] etc.) in the preceding 12 months.

Antibiotic resistance. Similarly, individual-level data on antibiotic resistance are not available in standard claims datasets.[5] 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.

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.

3.2 Other Datasets

Antibiotic prices and restriction status. Because the EHR data do not include antibiotic costs or restriction status, we obtained data on the cost-per-day and restriction status of each antibiotic at Hospital 1. New entrant prices were 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 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

and transplant.

[4]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.

[5]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.

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).

4 Reduced Form Evidence

This section presents reduced-form evidence for two key elasticities: (1) the relationship between antibiotic resistance and use and resistance and (2) the relationship between mortality and antibiotic coverage, defined as administering a broad-spectrum antibiotic to a non-resistant individual. We identify the use-resistance relationship using 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. To identify the effects of antibiotic treatment on mortality, we use a judges design that leverages heterogeneity in antibiotic approval propensities across different providers.

4.1 Antibiotic Restrictions Lower Use and Resistance.

Following prior reduced-form work (Kong et al., 2025), 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. During our sample period, two policy changes were implemented at the treatment hospital (Hospital 1), both aimed at reducing inappropriate antibiotic use. The first policy, implemented in 2011Q4, formalized the prior authorization process for linezolid and other restricted antibiotics shown in Figure 3. 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. The second policy, implemented in 2014Q2, issued new linezolid-specific prescribing guidelines that defined narrower clinical indications for its use. The control hospital (Hospital 2) did not adopt comparable ASP restrictions targeting linezolid during this period, allowing us to isolate the effect of these policies using quasi-experimental methods.

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 difference-in-difference-in-differences (DDD) models, paired with event-study specifications that assess the timing and dynamics of policy effects.

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:

$$Y_{ht} = \alpha_h + \delta_t + \beta_1(\text{Hospital1}_h \times \text{Post}_t) + \varepsilon_{ht}, \tag{1}$$

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.[6] For resistance, the comparison group consists of the subset of drugs for which non-susceptibility is routinely measured across hospitals.[7] 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:

$$\log(Y_{hdt}) = \beta_1(\text{Linezolid}_d \times \text{Hospital } 1_h \times \text{Post}_t) + \text{FE} + \varepsilon_{hdt}, \tag{2}$$

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.

[6] 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.

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

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 (2):

$$\log(Y_{hdt}) = \sum_{k \neq k_0} \beta_k \left(\text{Linezolid}_d \times \text{Hospital } 1_h \times \mathbf{1}\{t = k\}\right) + \text{FE} + \varepsilon_{hdt}, \tag{3}$$

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.

Antibiotic Restrictions Lower Use. Figure 1 summarizes the difference-in-difference and triple-difference estimates. Panel (a) plots the mean share of admissions with any linezolid use in each hospital, aggregated to the half-year level.[8] The vertical dashed lines denote the introduction of each policy change. Prior to the first policy change, trends in use are similar across hospitals. Immediately following policy implementation, 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 roughly a 20 percent decline relative to baseline use.

Panel (b) of Figure 1 reports event-study estimates from a triple-difference (DDD) specification in which 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 contraction in linezolid use, above and beyond contemporaneous trends in control antibiotics or hospitals.

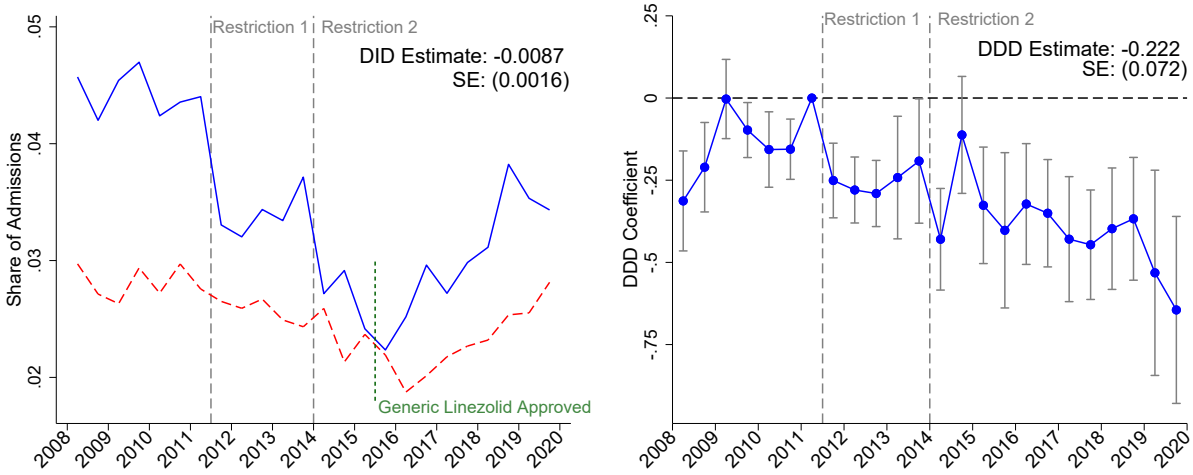
Antibiotic Restrictions Lower Resistance. Figure 2 examines whether these reductions in use translated into measurable decreases in resistance. Panel (a) plots the share of *Enterococcus spp.* isolates resistant to linezolid in each hospital over time. Prior to 2013, resistance rates trended similarly 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

[8] 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 1: 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) show 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 (3) 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. Graphs omit the singleton quarters 2008Q1 and 2019Q4, but these are included in the reported aggregate regressions.

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.

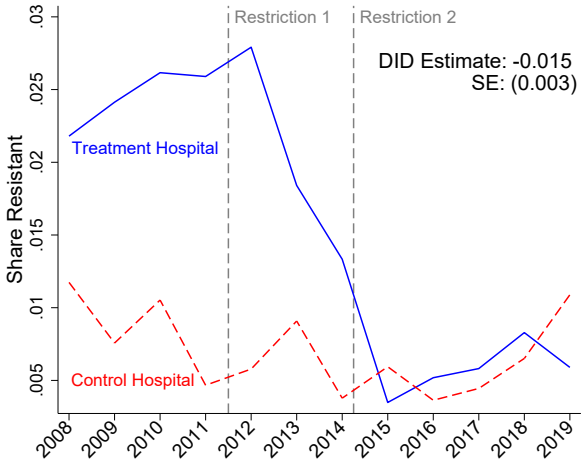
Together, these results demonstrate that reducing antibiotic use can lower hospital-level resistance. The timing of the decline in resistance—lagging use by roughly one year—is consistent with biological mechanisms where resistant bacterial strains are out-competed by non-resistant strains in the absence 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.

4.2 Identifying effects on mortality using provider staffing schedules.

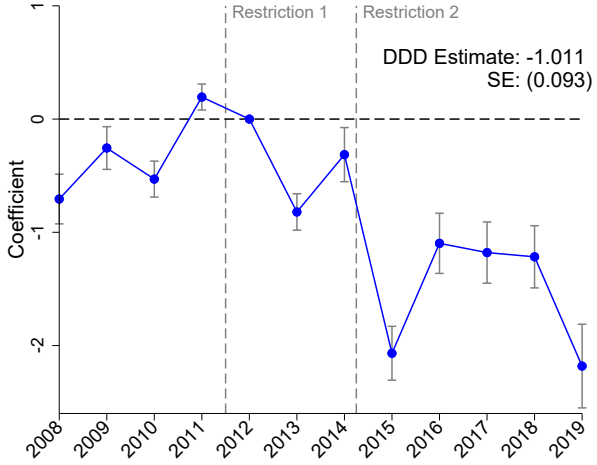
The causal effect of antibiotics and resistance on health outcomes, including mortality, is challenging to identify due to endogenous selection of patients into treatment. This is especially true for

Figure 2: 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.

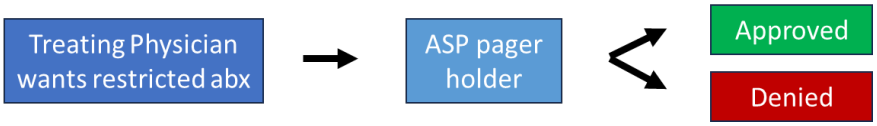
restricted broad-spectrum antibiotics, which are only given to patients satisfying strict criteria. To estimate the mortality equation of our structural model, we leverage exogenous features of the ASP prior-authorization process: physicians who wish to administer restricted antibiotics must obtain prior approval by submitting a request to an antibiotic stewardship pager, and the individual holding the pager (a physician, fellow, or pharmacist specializing in infectious disease) changes from week to week. The diagram in Figure 3 visualizes the approval process for restricted antibiotics.[9] We observe approver staffing and request data for the treatment hospital from 2011Q4-2016Q4.

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; this yields nontrivial week-to-week variation in approval rates. Second, many approvers are active in a given year, allowing us to control for year fixed effects. Without year fixed effects, a portion of our identifying variation comes from intertemporal variation in the approval rate due in part to the restriction policies described in Section 4.1. Our estimates are robust to whether or not year fixed effects are included. Lastly, this pager variation complements the policy variation by allowing

[9]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.

us to shift individual-level antibiotic use independently from hospital-wide antibiotic use. The individual-level variation is better suited for identifying short-run mortality effects on treated individuals, whereas the hospital-wide variation is necessary for estimating hospital-wide resistance changes, a large portion of which may be due to externalities that operate between individuals.

Figure 3: ASP Pager Approval Process



Notes: The figure illustrates the prior-authorization process for restricted antibiotics at Hospital 1. 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. These approval decisions generate quasi-random variation across pager holders, which we use to construct the leave-one-out approval rate instrument for linezolid utilization.

Instrument Construction and Identification. 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 on the first day of each ASP pager shift.[10] We exploit variation in different providers' propensity to approve linezolid as an instrument for linezolid use. 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 the ASP pager approval rate Z_i should only affect outcomes for patients for whom linezolid might make sense (e.g., those with more serious infections with gram-positive organisms). 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 10% of predicted linezolid use. Intuitively, covariates associated with high predicted linezolid use reflect physicians' assessments of which patients stand to benefit most from linezolid.[11]

Appendix Figure A2 shows that the LOO approval rate Z_i is uncorrelated with a range of placebo outcomes, including the admission length, predicted mortality, predicted linezolid use, and bacterial culture results, lending strong support to the validity of our identification strategy.

[10] 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.

[11] 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.

Estimation. We estimate the following two-stage least squares (2SLS) regression specification.

$$M_i = \beta_0 + \beta_1 X_i + \mathbf{W}_i'\boldsymbol{\gamma} + \varepsilon_i \tag{4}$$

$$X_i = \pi_0 + \pi_1 \cdot (T_i \cdot Z_i) + \mathbf{W}_i'\boldsymbol{\theta} + \eta_i \tag{5}$$

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,[12] 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.

Antibiotics lower mortality. Figure 4 shows the first stage and reduced form of our instrumental variables analysis of the effect of antibiotic use on mortality. Results are shown separately for patients with high predicted linezolid use $T_i = 1$ and low predicted linezolid use $T_i = 0$. Panel (a) shows the relationship between linezolid use and the leave-one-out pager approval rate, with linezolid use measured as the post-linezolid share of the admission. Panel (b) shows the relationship between within-hospital 30-day mortality and the same instrument. The LOO approval rate varies about 30 percentage points across approvers in the data. [13]

Consistent with our identification strategy, the LOO approval rate instrument strongly predicts linezolid use and mortality among the $T_i = 1$ group but not the $T_i = 0$ group. Our estimates imply that for the high predicted use group $T_i = 1$, each 10 percentage point increase in the LOO approval rate raises linezolid use by 1.32 percentage points and reduces 30-day mortality by 0.8 percentage points.Taken together, these results imply that for the marginal individual, being approved for linezolid earlier in their admission reduces their 30-day mortality risk.

Our two-stage least squares 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 worried 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. The goal of our reduced-form mortality results is to illustrate our identification strategy and provide evidence that antibiotic use is important for patient health outcomes. In our analyses of counterfactual policies, we test sensitivity to a range of mortality effects.

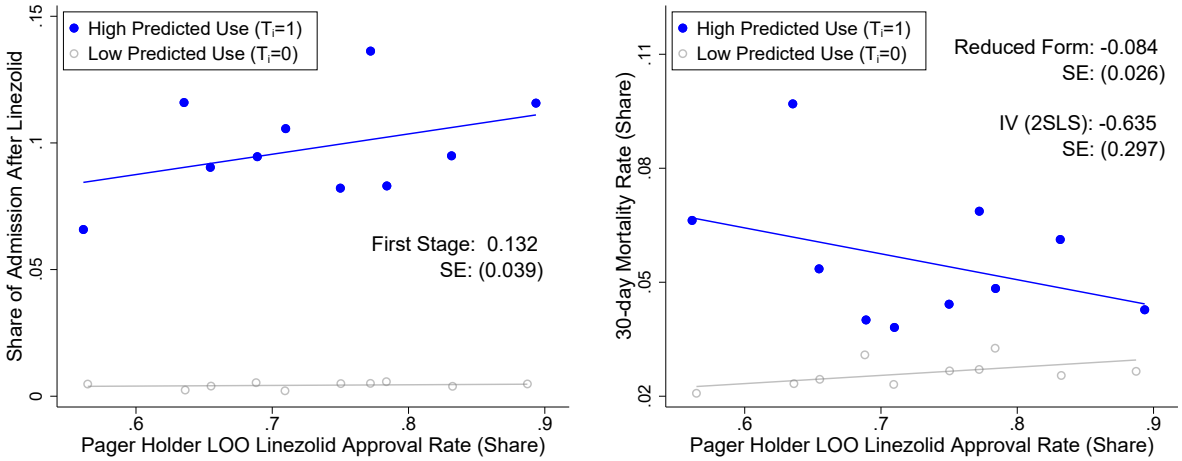
[12] Age bins were defined as 0-17, 18-40, 41-64, 65-80, and $\geq$80 years.

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

Figure 4: 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 ($T_i = 0$), as defined by the machine-learning predictor described in Section 3. Panel (a) shows that linezolid utilization increases with the leave-one-out (LOO) linezolid approval rate, only for patients with high predicted use. Linezolid utilization is measured as the post-linezolid share of each admission (taking value 0 for admissions with no linezolid use and 1 for admissions where linezolid is given on the first day). 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. Estimates corresponding to Equation (4) are shown on each figure; standard errors are clustered at the approver level, with 36 unique approvers.

The following section introduces our conceptual framework that relates these elasticities and incorporates innovation externalities. We start by providing a graphical intuition for optimal antibiotic use and innovation prizes, and then proceed to specify the full structural model. The same identification strategies presented in this section will be leveraged to identify the structural equations.

5 Model Overview

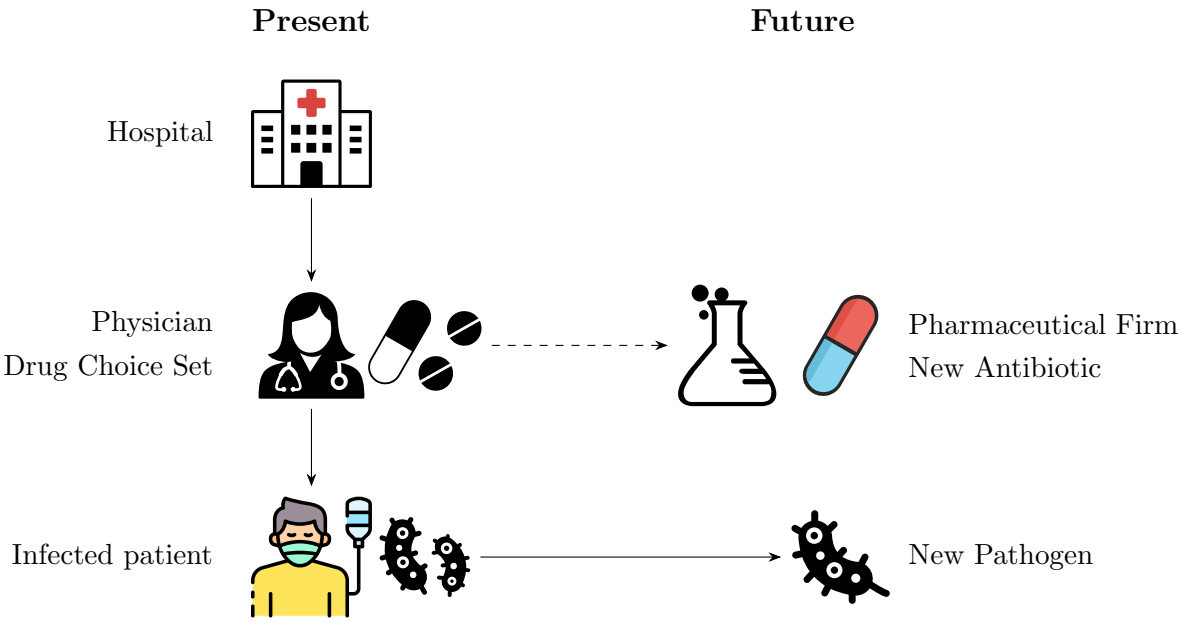
The following model relates antibiotic demand, resistance, innovation, and mortality to provide a framework for understanding the trade-offs involved in antibiotic regulation. A core result of the model is that antibiotic use generates multiple opposing externalities: higher use raises resistance, but also increases the chance that a new antibiotic enters the market.

5.1 Agents and Timing

We begin with an overview of the conceptual framework, summarized in Figure 5. 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 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 5: Framework linking use, resistance, and innovation



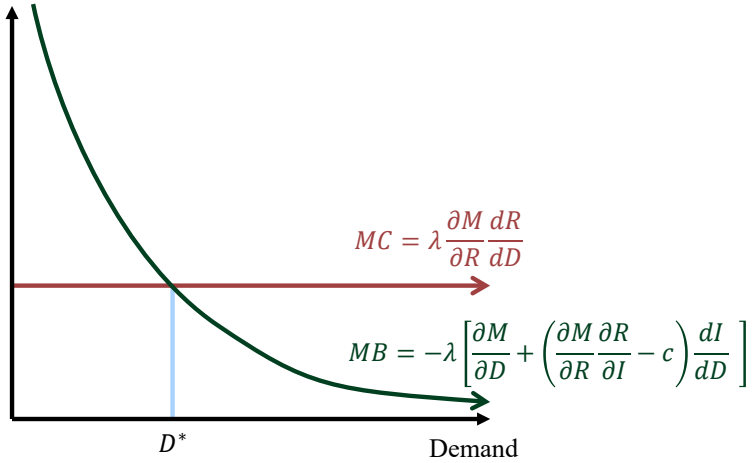
Notes: Figure shows a conceptual schematic of antibiotic use, innovation, and bacterial resistance. The left column ("Present") represents how antibiotic demand and patient outcomes are determined. From top to bottom, hospital policies affect physician antibiotic choice, which in turn affects patient 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.

Graphical Intuition For Optimal Antibiotic Use and Innovation Prizes. Figure 6 illustrates the key tradeoffs that determine optimal antibiotic use and the role of policy in aligning private and social objectives.[14]

[14]These figures adopt a set of simplifying assumptions to obtain illustrative closed-form solutions for the marginal benefit and marginal cost curves under various sets of available policies (see Appendix Section C for derivations).

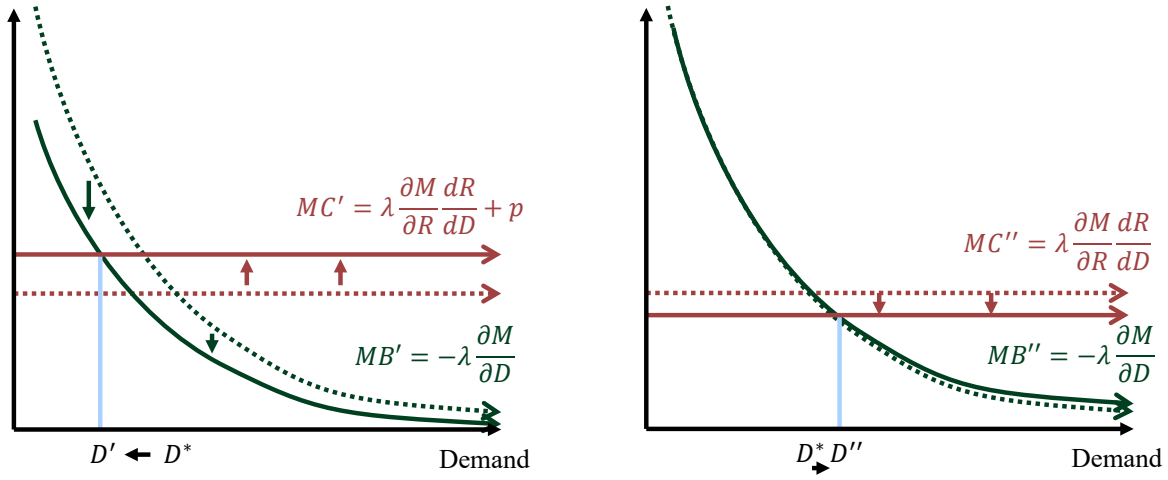
Figure 6: Optimal Antibiotic Use under Alternative Policy Regimes

(a) Optimal antibiotic use without innovation prizes (social planner)



(b) Optimal planner demand, without prize

(c) Optimal planner demand, with prize



Notes: Figure shows marginal benefit and cost curves relevant for regulating antibiotic demand. Panel A shows the social planner's objectives when they only have access to policies that shift demand. Here, D^* equates marginal benefits from mortality reductions and net innovation incentives with marginal costs from resistance externalities. In Panels B and C, dashed lines repeat the equilibrium in Panel A, whereas arrows and solid lines indicate new equilibria. Panel B shows private hospital objectives, which omit innovation effects but include antibiotic prices p. These changes lower the marginal benefit curve and raise the marginal cost curve, resulting in lower demand D'. Panel C shows the social planner's objective with access to both demand shifters and innovation prizes. Here, optimal use only balances mortality and resistance. The net effect moving from MB to MB'' is ambiguous, because higher innovation also reduces the level of resistance, which increases dM/dD. Marginal costs may fall if resistance is less responsive to demand at lower baseline resistance levels. Thus, whether D'' exceeds D^* is theoretically ambiguous.

Optimal use equates marginal costs with marginal benefits. From the social perspective, marginal costs include only the resistance externalities. If innovation prizes are not available, social marginal benefits include both short-run mortality benefits and future innovation incentives (Panel A). From the perspective of hospitals, marginal costs include antibiotic prices whereas marginal benefits exclude the innovation externality (Panel B). If the social planner is allowed to use innovation prizes, this de-couples antibiotic use and innovation incentives. However, because higher innovation reduces resistance, the sign of the overall change in demand, relative to the social optimum without prizes, is theoretically ambiguous (Panel C).

6 Structural Model Setup and Estimation

6.1 Motivation for Estimating a Structural Model

The conceptual framework presented in Section 5 provides intuition for understanding the relationships between antibiotic use, resistance, mortality, and innovation. To analyze the impacts of counterfactual policies, we specify a structural model of demand, resistance, and mortality that combines the identifying variation from Section 4 with individual-level data.

The structural model uses the same timing and equations as introduced in Section 5, but we do not assume a single representative antibiotic. Instead, we model resistance and demand for the full choice set of antibiotics for gram-positive infections (comprising 5 antibiotics and 1 outside option). The resistance and demand equations are identified using stewardship policy changes implemented in the treatment hospital only. 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.[15]

We quantify welfare in terms of mortality, measured as the number of admissions resulting in death within 30 days, scaled to the level of the U.S. The effect of antibiotic use on mortality is identified via our instrumental variables design based on the identity of the provider approving restricted antibiotic requests each week.

The estimated model will allow us to simulate demand, resistance, and mortality over time as a function of antibiotic stewardship restrictions on linezolid. We assume that a single potential entrant decides whether to enter based on expected future revenues over a 10-year exclusivity period.[16] Entrants are added to the choice set with initial resistance set to zero.

The subsections below describe the setup and estimation for each equation, starting with demand, then resistance, then mortality, and finally innovation.

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

[16]We assume that the entrant has the same drug characteristics as linezolid, including drug price, but no use or resistance history.

6.2 Model Equations

Estimating the model amounts to estimating four separate equations describing antibiotic demand, resistance, mortality, and innovation. In the following sections, we report our methods and estimates for each equation, in the following order. First, drug demand D_{ijt} models whether drug j is chosen for admission i at time t. Demand depends on antibiotic restriction policies and patients' medical history (including resistance). Second, resistance R_{ijt} depends on individual-level medical histories (including prior antibiotic use and resistance), as well as hospital-level demand and resistance over the previous year. The inclusion of past demand captures the externality of antibiotic use on population-level resistance, and the inclusion of past resistance accounts for the potential for resistance to self-persist over time in the absence of antibiotic use. Mortality M_{it} is determined by patient characteristics, the antibiotics administered, and the patient's resistance profile. Lastly, 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.[17] Figure A4 summarizes the relationships between these equations.

6.3 Demand Model

Setup and Estimation. Patients indexed by i are hospitalized in year-month t and choose antibiotic alternatives j from a choice set of six antibiotics, plus an outside option of no antibiotic.[18] Treating physicians choose antibiotics based on patients' observable characteristics, including past antibiotic use and any ASP policies in effect that month. We assume physicians do not directly consider antibiotic prices or future resistance when making treatment decisions. For a given alternative j, past use is considered separately for higher- and lower-tier antibiotics, coded as averages across the relevant set.[19]

We model only the highest-tier antibiotic administered during each admission. Linezolid, the highest-tier antibiotic overall, 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 cultures return susceptible to lower-tier antibiotics. Modeling the highest-tier antibiotic administered accounts for all of these scenarios. Physicians have beliefs about likely pathogens and resistance profiles before results are official, and may avoid prescribing antibiotics that are likely to be ineffective. To capture this, we include the resistance probabilities predicted using the estimated individual-level resistance model from Equation (7).

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

[18] The choice set reflects antibiotics used to treat gram-positive infections and is identical to the resistance model (i.e., first-generation cephalosporins, clindamycin, daptomycin, linezolid, trim/sulfa, or vancomycin). Admissions where the outside option is chosen includes cases where a different antibiotic outside of the choice set was administered. Our choice set is limited to antibiotics for gram-positive infections; patients may still receive antibiotics to treat other infections from gram-negative or anaerobic pathogens.

[19] For example, for j = vancomycin, the set of higher-tier antibiotics comprises linezolid and daptomycin. Thus, the higher-tier past use variable for a patient who previously used linezolid but not daptomycin would take value 0.5.

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

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

Choice utility U_{ijt} depends on the current ASP policy A_{ht} 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 also control for past antibiotic utilization through the $\tilde{u}_{ijt}$ term, which includes past use of higher- and lower-tier antibiotics. The coefficient $\beta^D_{u,j}$ is also allowed to vary by drug. The demand model further includes a rich set of controls $\mathbf{W'_{it}}$, 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.[20] The error term ε^u_{ijt} is assumed to follow the standard Type I extreme value distribution.

Descriptive Statistics for Antibiotics and Patients. The demand estimation sample includes 386,498 unique hospital admissions with 6 choice alternatives each. The choice set includes 5 antibiotics used for infections with gram-positive organisms. These include first-generation cephalosporins (cefazolin and cephalexin), clindamycin, the combination trimethoprim/sulfamethoxazole, vancomycin, and linezolid (the antibiotic affected by the ASP policy changes). The choice set also includes an outside option that comprises both no antibiotic use and antibiotics used for non-gram positive pathogens. [21] We divide antibiotics into tiers based on how they are are used clinically, with higher-tier antibiotics reserved for more resistant pathogens. Because a patient may trial multiple antibiotics, we model the highest-tier antibiotic chosen for each admission.[22]

Table A1 shows drug-level characteristics. One of the 5 antibiotics is chosen 58% of the time, with linezolid chosen in 2.8% of admissions. Vancomycin, the most commonly chosen antibiotic at 36.4%, is an IV infusion that is often used as the empiric (initial) treatment when the culprit pathogen not yet known; it covers almost all gram-positive organisms except for some strains of *Enterococcus spp.* Importantly, linezolid covers the same pathogens as vancomycin, plus vancomycin-resistant *Enterococcus.* Prior use of a given antibiotic within the past 12 months is an important predictor for current use. Prior use is highest for vancomycin (at 14.9%) followed by cephalosporins (at 14.8%), which are the the lowest tier of antibiotics. Prior use of linezolid is rare, at 1.8%, reflecting its status as a restricted antibiotic. Predicted resistance to each antibiotic is another important factor in demand.[23]

[20] This term also includes controls for predicted resistance to each antibiotic (and higher- or lower-tier antibiotics).

[21] Other categories of antimicrobial drugs include antibiotics for gram-negative infections, antibiotics for anaerobic infections, antibiotics for "atypical" pathogens, or antifungals. These other categories are often used in conjunction with the anti-gram-positive antibiotics in the choice set, and choices for those categories can be considered independently from the choice of whether to use an anti-gram-positive antibiotic.

[22] When multiple antibiotics occupy the same tier, we take the antibiotic with the latest administration date.

[23] We use the estimated resistance model in (7) to predict, for each admission and antibiotic, the probability of resistance.

Table A2 shows descriptive statistics for admission-level controls including demographics, ICD diagnosis codes, and prior hospital admissions. 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 ($\hat{\beta} = -0.37$, SE = 0.04) and increased demand for other antibiotics in the choice set, including vancomycin ($\hat{\beta} = 0.17$, SE = 0.02), and lower-tier antibiotics ($\hat{\beta} = 0.09$, SE = 0.02), all relative to the outside option. These results confirm that stewardship policies substantially shifted antibiotic choice toward alternatives when linezolid became more restricted.

Table 1: Demand Model Estimates

	Drug j chosen
ASP Policy ×	
Linezolid	-0.374**
	(0.037)
Vancomycin	0.167**
	(0.018)
Lower-tier antibiotic	0.089**
	(0.017)
Predicted resistance for j = linezolid	
Predicted resistance to linezolid (self)	-1.291
	(0.876)
Predicted resistance to lower-tier antibiotic(s)	3.209**
	(0.312)
Prior use for j = linezolid	
Prior use of linezolid (self)	1.745**
	(0.055)
Prior use of lower-tier antibiotic(s)	1.234**
	(0.093)
Number of admissions × choice alternatives	2,318,988
Number of admissions	386,498

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

Notes: Table shows conditional logit estimates for the demand model in Equation (6). Each observation represents a unique patient admission where a bacterial culture was drawn. The model includes controls for admission-level covariates (ICD diagnosis codes, prior antibiotic use, prior cultures, and demographics), which are allowed to vary freely by antibiotic (see Appendix Table A3). Coefficients for higher-tier antibiotics are omitted because linezolid is in the highest tier. Standard errors (in parentheses) are clustered at the patient level.

The next set of coefficients indicates that predicted linezolid resistance lowers 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 and expected resistance, especially to lower-tier antibiotics. Coefficients corresponding to other antibiotics and drug-by-patient characteristic interactions are reported in Appendix Table A3.

Simulations from the demand model. While the ASP policy effects are well identified, 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 baseline simulations, we assume that past use effects reflect true state dependence, while effects of predicted resistance capture unobserved heterogeneity.

6.4 Resistance Model

Setup: Individual- and Hospital-level Determinants of Resistance. Given the set of 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 for the fact that each culture may be resistant to a single antibiotic, multiple antibiotics, or none. Second, for low baseline levels of resistance (as observed for our antibiotic of interest linezolid), the logit functional form encodes convex resistance. Finally, it allows modeling of 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 resistance to higher- and lower-tier antibiotics.[24] Analogously, $\tilde{\mathbf{u}}_{\mathbf{cj}}$ includes indicator variables for whether i used, in the past 12 months,

[24] Antibiotics are divided into 4 tiers as follows. Linezolid and daptomycin are the highest-tier, newest, and most heavily restricted antibiotics in the choice set and resistance data. The next tier includes vancomycin, which treats most gram-positive infections, with the notable exception of vancomycin-resistant *Enterococcus* (VRE), which is only susceptible to linezolid and daptomycin. The next tier comprises trimethoprim/sulfamethoxazole (trim/sulfa) and clindamycin. The final and lowest tier comprises the first-generation cephalosporins cefazolin and cephalexin. These two cephalosporin antibiotics are near equivalents, so we group them together as one alternative in the demand and resistance models.

antibiotic j, any higher-tier antibiotics, and any lower-tier antibiotics. The vector $\mathbf{H_c'}$ 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 3. The coefficients on $\mathbf{\tilde{r}_{cj}}$, $\mathbf{\tilde{u}_{cj}}$, and $\mathbf{H_c'}$ are allowed to vary flexibly by antibiotic; for example, the persistence of resistance within an individual may differ by antibiotic. The term ε_{cj} denotes the logit error, which is assumed to be independent across isolates c and antibiotics j, conditional on the included observables.

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

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

where the $\delta\alpha_{hyj-1}$ term captures the dependency of this year's resistance on last year's resistance, with the decay factor δ. This functional form assumption allows shocks to resistance and utilization to persist, reflecting the biological reality that resistant strains of bacteria may persist in the hospital environment. The hospital-level utilization term $\tilde{u}_{hjy}$ denotes the average lagged utilization of drug j in hospital h in year y, with γ representing the resistance externality of antibiotic use.[25] The a_{hj} term represents hospital-drug fixed effects and adjusts for fixed differences in resistance across different drug-hospital pairs.[26] Lastly, η_{hjy} is the error term.

Identification. 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 between-patient variation in 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 a set of hospital-specific ASP policy changes, denoted by $\tilde{A}_{hjt}$ and defined in Section 3.

Estimation. We use a two-step estimation approach to estimate Equation 8. In step 1, we estimate the patient-level model (Equation 7) as separate logistic regressions for each antibiotic j, where the α_{hyj} terms are concentrated out as fixed effects. The estimated hospital-drug-year fixed effects $\hat{\alpha}_{hyj}$ are treated as data in step 2 below.

[25] Specifically, for each hospital-month, we calculate the share of admissions using linezolid in the past 12 months; the $\tilde{u}_{hjy}$ variable is the yearly average of this monthly measure. We construct the lagged utilization variable in this fashion (averaging over monthly observations), because the individual-level resistance and demand models will be simulated at the monthly level.

[26] 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.

In step 2, we estimate the hospital-level model in Equation 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, especially in panels with few observations in the time dimension (Roodman, 2009; Blundell and Bond, 1998; Arellano and Bover, 1995).[27] 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 (akin to a two-by-two difference-in-differences design), we use between-hospital differences as inputs to the estimators. 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 first replace $\tilde{u}_{jy}$ in Equation 9 by its predicted value from the first stage OLS regression given by Equation 10:

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

In our preferred specification, we estimate Equation 9 using the first-difference GMM estimator (Arellano and Bond, 1991). We also confirm that our results are robust to instead using the system GMM estimator of Blundell and Bond (1998) and Arellano and Bover (1995).

Benchmark 2SLS Estimates. An alternative solution that avoids using dynamic panel estimators is to re-specify 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. One benefit of this is approach is that we can apply a standard two-stage least squares estimator (2SLS) to estimate the resistance externality. This 2SLS specification and accompanying first stage are:

$$\alpha_{hjy} = \tilde{\gamma}\tilde{u}_{hjy} + b_{hjy} + b_{hj} + \nu_{hjy} \tag{11}$$

$$\tilde{u}_{hjy} = \theta\tilde{A}_{hjy} + c_{hjy} + c_{hj} + \varepsilon_{hjy} \tag{12}$$

Although this 2SLS specification lacks the dynamic dependency of current resistance on past resistance, we show in the following section that our dynamic-panel estimates of δ do not differ significantly from zero.

Descriptive Statistics for Bacterial Isolates. Table A4 shows descriptive statistics for the resistance dataset, where underlying observations are independent positive bacterial cultures.[28]

[27] In our case, we have $T = 12$ observations in the time dimension.

[28] Often, cultures may be repeated for the same patient across multiple days. We group cultures within the same patient, specimen type, and pathogen that occur less than 7 days apart.

Each isolate is tested for resistance to a panel of antibiotics that potentially cover that pathogen. 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.

Resistance rates for linezolid, the drug affected by our exogenous ASP policies, 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, patients previously used linezolid in 11% of cases, and used lower-tier antibiotics in 23% of cases. By comparison, only 4.5% of isolates tested for vancomycin resistance were associated with prior use of a higher-tier antibiotic (linezolid or daptomycin), whereas 37% had prior own-antibiotic use.

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. For individuals, 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, $P < 0.01$).

Across other antibiotics, the average own-antibiotic effects are smaller in magnitude but remain statistically significant: prior resistance and prior use of the same antibiotic have estimated coefficients of 1.66 and 0.54, respectively (both $P < 0.01$). These results indicate that while persistence of individual-level resistance is generalizable across antibiotics, magnitudes are particularly large for linezolid. This suggests that the internalities of antibiotic use on a patient's own future resistance are stronger for linezolid, potentially because it is a more rarely used antibiotic. For more commonly used antibiotics, resistant bacterial strains may already be circulating widely, weakening the relationship between one's own use and resistance.

In addition, resistance is positively associated with prior resistance to or use of higher-tier antibiotics (average estimated coefficients 0.17 and 0.18, respectively), although the latter is not statistically significant. We do not find significant associations with prior resistance or use of lower-tier antibiotics. 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. Appendix Figure A5 shows that predicted resistance probabilities closely match observed resistance, both overall and by antibiotic, indicating that the included covariates capture substantial individual-level heterogeneity in resistance.

Hospital-level Resistance Model Estimates Table 3 shows estimates of the hospital level of the resistance model for linezolid. Column 1 shows a strong first stage; our estimates show that the lagged ASP policy instrument led to a 1.0 percentage point decline in linezolid use (SE=0.001,

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.

$P < 0.001$). Column 2 shows 2SLS results that imply a large resistance externality, with a 1pp increase in the share of admissions using linezolid resulting in a 1.37 increase in the hospital-level fixed effect, which 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, hospital-level resistance externalities are nearly as important as an individual's own prior use of the antibiotic. Columns 3 and 4 show results from dynamic panel estimators that obtain consistent estimates of the AR(1) decay coefficient δ. Both the first-differences GMM (Arellano and Bond, 1991) and system GMM estimators (Blundell and Bond, 1998; Arellano and Bover, 1995) yield estimates of 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 persistent from year to year on its own. That said, our 95% confidence intervals include values of δ as large as 0.33.

Decomposing Resistance Into Externalities and Individual-level Effects. To better illustrate the relative magnitude of the hospital-level externality, Appendix Table A5 shows how the

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^{**}	159.7^{**}	142.9^{**}
		(26.49)	(13.94)	(21.67)
AR1 Coefficient δ			-0.07	0.03
			(0.16)	(0.15)
Fixed Effects	Drug × Hosp	Drug × Hosp	Drug	Drug
	Drug × Year	Drug × Year		
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.

effects of the two ASP policy changes implemented in Hospital 1 decompose 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 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 1. 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.[29] We find that 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.

Our estimates suggest that most of the effect of utilization on resistance operates via hospital-level externalities. 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

[29]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.

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%.

6.5 Mortality Model

Setup. We model mortality as a function of linezolid utilization as follows:

$$m_i^* = \psi_0 + \psi_1 u_{i,\text{linezolid}} + \mathbf{X_i'}\psi_{\mathbf{2}} + \mu_i, \quad M_i = \begin{cases} 1 & \text{if } m_i^* \geq 0 \\ 0 & \text{if } m_i^* < 0. \end{cases} \tag{13}$$

Equation (13) specifies a latent-index model of mortality. The dependent variable M_i is an indicator for whether admission i resulted in in-hospital mortality, 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 coverage on mortality, where $u_{i,\text{linezolid}}$ is an indicator for whether admission i ever received linezolid. 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, [30] predicted mortality squared, and Z_i and T_i as defined in Section 4. The associated parameter vector $\psi_{\mathbf{2}}$ captures the effects of these covariates on mortality risk. Finally, the error term μ_i represents unobserved factors affecting mortality, which is allowed to be correlated with unobservables affecting linezolid use.

Identification. The above model illustrates the challenge with identifying the causal effect of coverage on mortality: selection into antibiotic treatment means that antibiotic coverage (specifically, u_{ij}) is likely to be highly correlated with unobserved factors μ_i that affect mortality. For example, patients that are sicker in unobservable ways are more likely to receive higher-tier antibiotics, leading to a strong upward bias. In our sample, this bias leads to a positive correlation between observed linezolid use and mortality. As introduced in Section 4, we employ our ASP pager instrument as a source of exogenous variation in linezolid use.

We instrument $u_{i,\text{linezolid}}$ using the interaction $Z_i \times T_i$, where Z_i denotes the leave-one-out approval rate for linezolid, and T_i is an indicator variable that takes value 1 for admissions in the top decile of predicted linezolid use. We include the main effects of T_i and Z_i in the control vector $\mathbf{X_i}$. The intuition for this instrument is that the leave-one-out pager approval rate Z_i affects linezolid use, but only for patients with $T_i = 1$, as shown in Figure 4. The first stage regression therefore takes the form:

$$u_{i,\text{linezolid}} = \delta_0 + \delta_1 Z_i \times T_i + \mathbf{X_i'}\delta_{\mathbf{2}} + \varepsilon_i. \tag{14}$$

[30] 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.

Estimation. We estimate Equations 13 and 14 as an IV probit using maximum likelihood following Newey (1987).[31] Accordingly, we assume that 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 translate changes in linezolid use and resistance into effects on mortality risk, because patients must both use linezolid and not be resistant to linezolid in order to benefit. This points to a key intuition: linezolid resistance only harms mortality outcomes for linezolid users who require it for coverage.

Limitations. One simplification of our model is that we have assumed exogenous arrival of infections, whereas in reality, treating infections successfully may reduce their spread. This is an additional benefit of antibiotic treatment that is excluded from our model, and would tend to push in favor of expanded antibiotic use. Given modern infection control practices in hospitals and the lower transmissibility of bacteria (compared to, for example, airborne viruses), we expect the benefits of treatment on decreased spread to be modest.

Descriptive Statistics on Likely Users of Linezolid. Table A7 shows actual linezolid use as a function of predicted linezolid use. Importantly for our analysis, linezolid use is highly concentrated among admissions in the top decile of predicted use ($T_i = 1$). Table A6 reports descriptive statistics for the mortality sample, separately by treated ($T_i = 1$) and control ($T_i = 0$) admissions. Recall that the $T_i = 1$ subset is defined as the top decile of predicted linezolid use. Row 1 shows that the top decile contains admissions with predicted linezolid use between 8-96%, consistent with the small share of linezolid users in the sample. Row 2 shows that 17% of the $T_i = 1$ group used linezolid during the admission, compared to only 1% of the $T_i = 0$ group. Baseline mortality is higher in the $T_i = 1$ group relative to the $T_i = 0$ group. We compute our instrument, the leave-one-out (LOO) approval rate for linezolid, for 343 weeks and for 36 unique pager holders. Values of the instrument did not differ between the two groups, supporting our assumption that ASP pager shifts are quasi-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 pivotal for ensuring coverage in this group, including a much higher VRE positivity rate, longer lengths of stay, and a higher rate of bacteremia or sepsis ICD codes for the given admission.

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. 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 non-coverage on

[31] We use the ivprobit function of Stata 17.

Table 4: Mortality Model Estimates

	(1) Baseline	(2) First Stage	(3) No Year FE	(4) No Pred. Mortality
	M_i	$u_{i,\text{linezolid}}$	M_i	M_i
Linezolid Coverage	-0.902**		-0.815**	-0.737**
($u_{i,\text{linezolid}}$)	(0.233)		(0.262)	(0.237)
Instrument		1.168*		
($Z_i \times T_i$)		(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	0.51		0.45	0.49
(μ_i, ε_i)				
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 non-coverage on 30-day mortality at the admission level. The sample is limited to admissions beginning at the start of each ASP pager shift. 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).

mortality of 0.077 within the $T_i = 1$ group. This implies that, for marginal patients in the $T_i = 1$ group, receiving linezolid coverage reduces 30-day mortality by 7.7 percentage points relative to not receiving coverage. This estimated effect is smaller in magnitude than the corresponding linear 2SLS estimates in Section 4, 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.

The results in Column 1 also show that 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. This is consistent with targeting of linezolid to individuals with a high mortality risk not fully captured by observables. Our model estimates imply that absent linezolid coverage, mortality rates for patients who received linezolid would have increased from 6% to 27%. Figure A3 visualizes treatment effect heterogeneity for individuals in the $T_i = 1$ group. As predicted linezolid use increases, so do the predicted mortality benefits from linezolid use.

Finally, we emphasize that 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.

6.6 Innovation Model

A key component of the structural model is the response of innovation to expected profits. A wide range of empirical estimates of this elasticity exist in the literature (for example, (Dubois et al., 2015; Kong and Zhao, 2025; Acemoglu and Linn, 2004; Finkelstein, 2004)). Our goal is not to estimate a detailed model of all steps involved in bringing a drug to market, but rather to capture the well-established relationship between market entry and the expected profits, which we are able to simulate in detail using the demand and resistance equations.

Accordingly, we estimate 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. To measure profits, we use net-of-rebate U.S. revenue data from marketed antibiotics between 1986 and 2024. We use drug approval data from the FDA Orange Book to measure entry. We estimate the distribution of unobserved entry costs using GMM, following Pakes et al. (2007). The estimated innovation model matches observed entry rates, imply average entry costs that are very similar to prior estimates in the literature, and yield innovation elasticities that fall within the range in the literature. The estimated model serves as our preferred specification; however, we test the sensitivity of our results 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 year, we assume that there are a fixed number of potential entrants K_g 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, with up to 6 new antibiotics or antifungals of any type entering each year. For each year and group, K_g potential entrants each draw entry costs $c_t \sim F(\mu)$, which we assume include all costs associated with failed candidates necessary to generate one new FDA-approved antibiotic. For each potential entrant, the value of entering in year t is the discounted present value of flow profits:

$$VE(g,t) = \sum_{\tau=1}^{T} \beta^{\tau} \pi(g,t), \tag{15}$$

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.[32] 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

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

year t, firms do not have private information about their profits before deciding whether to pay the entry cost and enter the market. We estimate $\pi(g,t)$ by computing average net sales across all antibiotics entering in group g and year t, denoted $\bar{\pi}_{gt}$. 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 represent $\pi(g,t)$, which are then used to construct $VE(g,t)$.

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 present-discounted value of profits is \$2.3B. The observed share of potential entrants that enter the market (15.4%).

Estimation. We assume the entry cost distribution $c_t \sim F(\mu)$ is exponential with mean μ. Given the exponential distribution assumption and the estimated values of $VE(g,t)$, there is only one scalar parameter. Hence, we use the following aggregate moment condition:

$$\frac{1}{N}\sum_{g,t} F(VE(g,t);\mu) = \frac{1}{N}\sum_{g,t}\frac{e(g,t)}{K_g}. \tag{16}$$

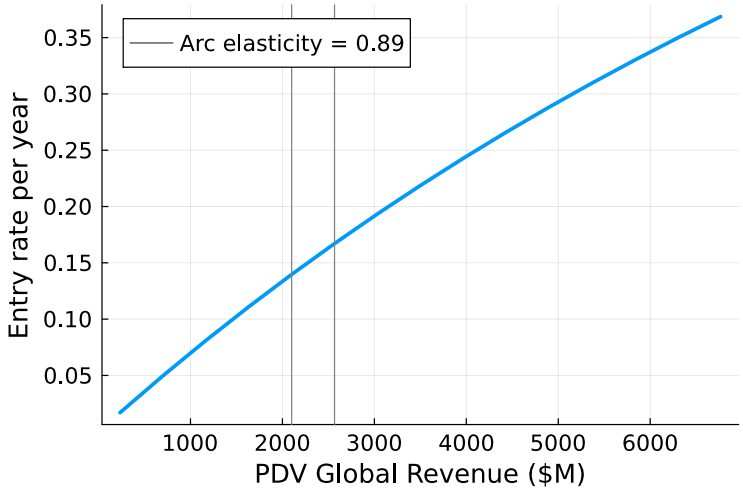
where $e(g,t)$ is the observed number of entrants in group g and year t. Thus, the estimation procedure matches the average model-implied entry rate to the average observed entry rate.

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 7 plots how the entry probability varies with profit in the 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, albiet 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

reasonable. That said, we also test the sensitivity of our counterfactuals to values at the extreme ranges of 0.25 to 2.75.

Figure 7: 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.

Simulating from the Innovation Model. The estimated model allows us to translate changes in expected profits to changes in the antibiotic entry probability. In our simulations, we assume that entrants have similar clinical characteristics to linezolid, except that entrants begin with zero resistance. Higher simulated demand increases expected profits and increases the entry probability. Similarly, innovation prizes are represented by adding the discounted flow of innovation prizes to $VE(g,t)$.

7 Counterfactuals

This section first describes how we run simulations using the estimated structural model. This is followed by baseline simulation results reflecting observed policies and comparative statics. Finally, we solve for optimal counterfactual policies that maximize long-run welfare.

7.1 Simulation Methods

To simulate from the full structural model (comprised of resistance, demand, mortality, and innovation equations), we will adopt a finite-horizon, forward simulation approach. Briefly, we simulate demand and resistance jointly using a forward-simulation approach based on (6) for demand and (7)-(8) for resistance. This captures the lagged structure of the resistance and demand equations,

whereby current values depend on past realizations. Given time series of demand and resistance, we simulate mortality from (13), under the assumption that linezolid use only reduces mortality for non-resistant individuals.[33] We repeat this process $S = 10$ times across different draws of the underlying *iid* shocks. For innovation, we model a single potential entrant deciding whether to enter in January 2009, one year after the start of our sample, using estimates from Section 6.6. We assume that the entering antibiotic has the same observed characteristics as linezolid, and is subject to the same restriction policies. For full details, see Appendix Section D.5.

7.2 Baseline Simulation Results and Comparative Statics

Figure 8 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 A10 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.

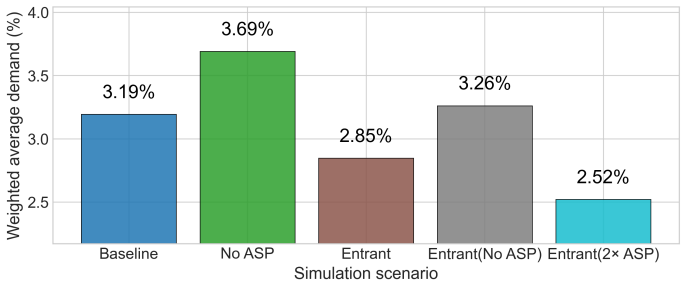
Overall, these results demonstrate how demand, resistance, and mortality respond to demand restrictions and new entry. Specifically, they show that new antibiotics yield greater mortality

[33] 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.

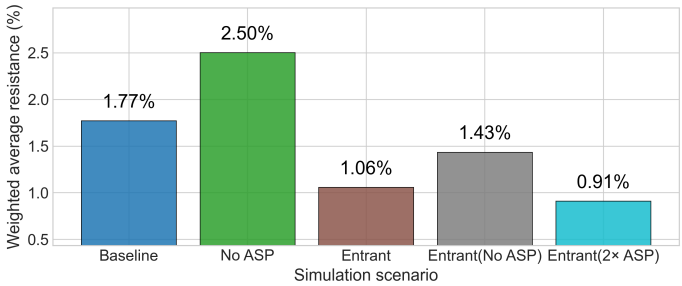
benefits when they are allowed to be used. Moreover, while higher resistance lowers antibiotic effectiveness for all individuals (including inframarginal users), this effect is outweighed by the direct mortality improvements among marginal antibiotic users, at least in a local region around existing 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 8: Simulated Comparative Statics

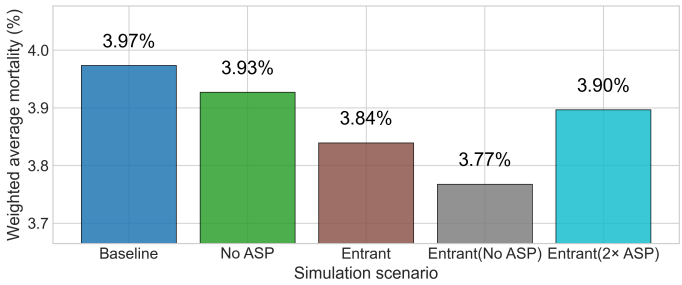
(a) Linezolid Demand



(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).

7.3 Social Welfare and Hospital Objectives

Solving for optimal policies first requires specifying the social planner's objective function. We assume that the social planner aims to minimize the expected sum of long-run mortality and

innovation prize payments. Accordingly, we define the long-run social welfare function as:

$$W(\theta, P) \equiv - \sum_{t \in \{0,\infty\}} \beta^t E \left[\lambda \underbrace{M_t^N(D_t, R_t)}_{\text{Mortality}} + \underbrace{P \times 1\{\text{patent}\}_t}_{\text{Innovation Prize}} \right], \tag{17}$$

where θ and P denote the demand shifter and innovation prize respectively. M_t^N denotes the yearly number of infection-related deaths scaled nationally to the entire U.S.[34] We use estimates of (13) to simulate the initial path of mortality, assuming mortality subsequently reaches a steady state.[35] We assume a value of statistical life (VSL) of λ =\$1 million and an annual discount rate of β =0.95. Finally, $P \geq 0$ denotes the innovation prize that is paid annually conditional on $1\{\text{patent}\}_t$, which takes value 1 if the novel antibiotic enters the market in period $t = 1$, for the duration of market exclusivity, which we assume to be 10 years. We consider P as a social cost rather than a transfer to drug manufacturers, because real costs are incurred when bringing a drug to market (e.g., labor and capital costs of running clinical trials, drug marketing, etc.).[36] The true social cost of prize P may be higher if the marginal cost of public funds is greater than 1 (Browning, 1976), or lower if part of the cost is recouped in the form of future tax revenue. Notably, θ affects mortality (via demand and resistance) as well as innovation (by changing expected profits and thus the probability of entry). 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 17 allows us to directly calculate the social value of a new antibiotic, by calculating (17) in the case without entry and then comparing to the where a new antibiotic enters the market. Of note, this estimate of the social value depends on θ, because the effect of the new antibiotic on mortality depends on the degree to which its use is restricted. In the extreme case, a novel antibiotic that is never used has zero value. While holding new antibiotics in reserve may provide insurance against resistance increases, our findings suggest that the costs of doing so can be significant.

Another caveat is that our estimate will tend to overstate the marginal value of a single entrant. This is not because of growing resistance for the entrant; rather, it is because the value of each entrant declines with subsequent entry. With those limitations in mind, we calculate that the 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 a long-run mortality reduction of 54,700 deaths when aggregated over time and scaled up to the level of the U.S. At a \$1 million VSL, this implies that each new antibiotic generates \$54.7 billion in additional present-discounted social welfare, or \$2.74bn per year, not including any additional deaths averted globally.

[34] 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.

[35] For linezolid, this assumption is supported by our empirical finding in Table 3 that resistance does not significantly persist over time on its own.

[36] We ignore sunk and fixed costs incurred by the drug manufacturer, which may vary with expected drug demand and innovation payments.

7.4 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 could 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. 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.[37]

We find that the optimal innovation prize (holding demand restrictions at status-quo levels) awards the innovator $1.04 billion per year, which increases the entry probability to 85.4% and raises long-run welfare (net of innovation prize payments) by $6.5 billion.

This welfare increase is achieved without increasing expected linezolid demand or resistance. Indeed, entry lowers linezolid resistance by substituting away a portion of demand. Thus, innovation prizes have the favorable quality of improving welfare with fewer information requirements given their modest effects on demand and resistance patterns.

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 rationally consider the acquisition costs of antibiotics. However, because the marginal cost of production of most drugs (after they have been invented) 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 payment. 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 shown in Table 1, Row 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 demand from 3% of admissions to 4.3% of admissions. Resistance also increases substantially, from 1.4% to 7.5%,

[37]This entry probability is significantly higher than the average entry rate per year reported in Section 6.6. This is because the potential entrant we consider has drug characteristics equivalent to linezolid, which had significantly above-average revenues.

but this is outweighed by declines in mortality from treatment expansion, raising overall long-run welfare by $88.6 billion.

This welfare increase is much larger than that of the optimal innovation prize. One reason for this is because the innovation prize involves direct fiscal costs: entering firms are paid a discounted sum of $8.4B, which is directly subtracted from long-run welfare. That said, raising demand 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 the entry probability only increases modestly, accounting for the *possibility* of innovation plays a pivotal role in determining optimal demand. Setting the probability of entry to zero and holding the antibiotic choice set fixed returns a lower optimal policy scaler of 2.0 (compared to 2.6 with endogenous innovation). 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 rationalizes more conservative use. That said, even with an entry probability of zero, the optimal policy scaler of 2.0 suggest 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; thus, the benefits of expanding treatment by 1.3 percentage points outweigh the costs of a 6.1 percentage point rise in resistance. However, these gains are bounded. Beyond a linezolid demand level of 4.6%, resistance increases steeply past 11%, and the associated mortality costs exceed the benefits of additional treatment (Figure 9). In our model, this reflects the nonlinear response of resistance to utilization, as shown in Figure A13. This sharp reversal highlights a challenge with demand-based policies: they require careful modeling of demand responses and resistance dynamics. As we show in our sensitivity analyses in Section 7.5, the use-resistance relationship is the most important elasticity to pin down for determining the optimal policy scaler. Together, these results show that while expanding demand can increase welfare significantly, any demand expansions would require ongoing surveillance of resistance levels to ensure that antibiotic use remains within welfare-maximizing ranges.

Optimal policy with two policy levers. The optimal combination of policies includes a similar policy scaler (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 long-run 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.
Figure 9 breaks down how different combinations of demand and innovation prizes translate into long-run welfare. Figure A11 show effects of policy combinations on mortality, linezolid resistance, entrant demand, entrant profit, and entry probability. Low levels of demand (policy scaler < 0)

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
Demand Shifter	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 (17). 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 (17) (Column 5), average linezolid demand (Column 6) and resistance (Column 7). Data are limited to the treatment hospital (Hospital 1).

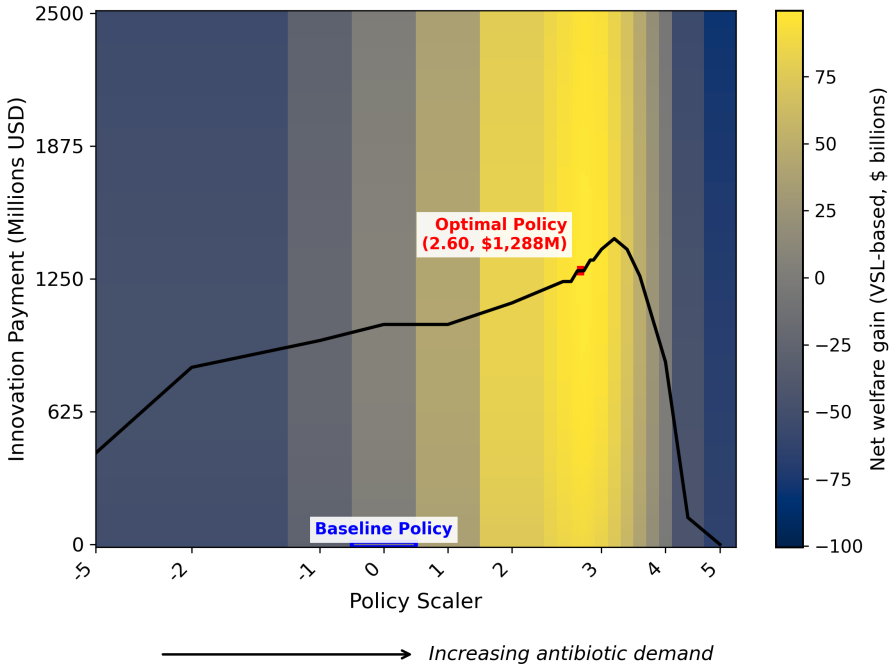
lead to undertreatment and low entry probabilities, resulting in negative welfare gains compared to the status quo. In contrast, very high levels of demand yield higher entry probabilities but also lead to higher resistance and overall mortality. The black curve shows the optimal innovation payment for each level of demand. For most levels of demand, 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, as shown in the upward-sloping optimal prize curve in Figure 9. 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. Conversely, a new antibiotic that is never used confers no value. Additionally, for convex regions of the use-resistance relationship, spreading the same amount of antibiotic utilization over several antibiotics, including entrants, improves welfare.

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 high values of baseline demand, demand and innovation payments are substitutes: if baseline demand were higher than optimal, then reducing demand would increase welfare and suggest higher optimal innovation payments.

Figure 9: Welfare as a function of demand and innovation policies



Notes: Figure shows the simulated national welfare gain relative to the baseline case (*policy scaler* = 0, *innovation payment* = 0), measured as the present discounted value of annual U.S. deaths (scaled by VSL = $1 million) plus innovation prize costs. Each panel shows long-run welfare 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). Lighter shading indicates higher net welfare. The black curve denotes the optimal innovation payment as a function of the policy scaler. The boxed cell denotes the welfare-maximizing policy combination.

7.5 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. 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 A12).

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.

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.

Higher resistance elasticities reduce the optimal level of demand (optimal policy scaler = 1.50) and decrease overall welfare gains, whereas lower resistance elasticities imply higher optimal demand (optimal policy scaler = 4.00) and much higher welfare gains from expanded demand and innovation. This is unsurprising, as antibiotic use directly trades off current mortality and future resistance. Optimal innovation prizes and their welfare gains are stable. 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 optimal innovation prizes. Even without changing antibiotic demand, doubling the mortality effect of antibiotics increases the optimal prize level from \$1.04B to \$1.72B and raises the associated long-run welfare gain 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 $\lambda_I = 0.28$.[38] We test an upper bound innovation elasticity of 2.75 using $\lambda_I = 3.09$.

Moving to a small innovation elasticity of 0.25 does not significantly change the optimal joint policy, but long-run welfare gains are somewhat lower (\$79.1B versus \$99.6B). However, the optimal innovation prize in isolation becomes \$0, as demand complementarities become necessary to justify any level of prize.

Moving to a high innovation elasticity of 2.75 modestly increases all long-run welfare gains. Optimal innovation prizes are lower. Even without expanding demand, 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 demand and innovation are substitutes rather than complements. The reason for this is because innovation becomes saturated as

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

entry probabilities approach 1. The point where demand and innovation policies transition from complements to substitutes occurs at innovation elasticities exceeding 1.4.

We have also imposed structure on entrant and social planner information. We assumed that the potential entrant knows its drug characteristics prior to entry. Thus, promising antibiotics will tend to have high entry probabilities even without innovation prizes, and prizes may have larger effects on antibiotics with lower expected demand. To consider multiple potential entrants or distributions of entrant characteristics, we would need to specify how the planner would allocate prizes, which may involve asymmetric information between the 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) and an important area for future work.

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 7.5.

8 Discussion

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

On the theory side, we introduced a simple conceptual framework that expresses optimal antibiotic policy as a tradeoff between current mortality and future effects on resistance and innovation.

Because hospitals do not internalize innovation externalities, antibiotic demand may fall short of optimal levels. This is further exacerbated by high prices for non-generic antibiotics and hospital budget-constraints.

Empirically, we used electronic health record data to estimate reduced form effects of antibiotic use on resistance and mortality. Surmounting limitations in claims data, the EHR data allowed us to observe details on use and resistance for specific antibiotic and pathogen combinations. Leveraging hospital-specific policy changes facilitated difference-in-difference and triple-difference identification strategies for the use-resistance relationship. Additional administrative data on staffing schedules of antibiotic approvers allowed us to identify effects of antibiotic use and resistance on mortality. We use panel data on antibiotic restriction policies and acquisition costs to show that hospitals are price sensitive.

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 tradeoff between antibiotic use and resistance. If antibiotic coverage did not significantly affect mortality, then the debates surrounding antibiotic use and resistance would have far less social relevance. Instead, our estimates implied 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, and highlight the importance of developing optimal policies for regulating antibiotic demand and resistance.

Consistent with prior findings, our results confirm that revenue 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; Acemoglu and Linn, 2004; Dubois et al., 2015). In our simulations, we found that a large optimal innovation prize of $1.04 billion per year increases the probability of a new antibiotic by 47%, lowers expected resistance among existing antibiotics, and raises net welfare by $6.5 billion.

Our results also underscore the effect of demand regulation 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 is 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 long-run 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.

We also find that demand-focused policies and innovation prizes are complementary over a large region of the policy space, implying that combinations of innovation prizes and higher demand yield the largest welfare gains. 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. 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 is consistent with prior work showing that prizes primarily serve to supplement intellectual property arrangements or follow-on contracts that scale rewards with innovation quality (Weyl and Tirole, 2012; Che et al., 2021). Our findings also highlight that 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; our results show that this will depend crucially on the shape of the resistance externality for each antibiotic-pathogen pair. To induce optimal quantities, potential policies range from carving out antibiotics from bundled hospital reimbursement; to taxes or subsidies for novel antibiotics; to direct, fixed allotments of novel antibiotics.

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 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.

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.

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.

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

Centers for Disease Control and Prevention (2013): "Antibiotic Resistance Threats in the United States," Tech. rep.

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*.

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.

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.

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

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

Kong, E. L., Kubiak, David W., E. S. Shenoy, and A. R. Letourneau (2025): "Trade-offs in Antibiotic Stewardship: Resistance, Screening, and Mortality," *Working Paper*.

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.

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.

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.

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.

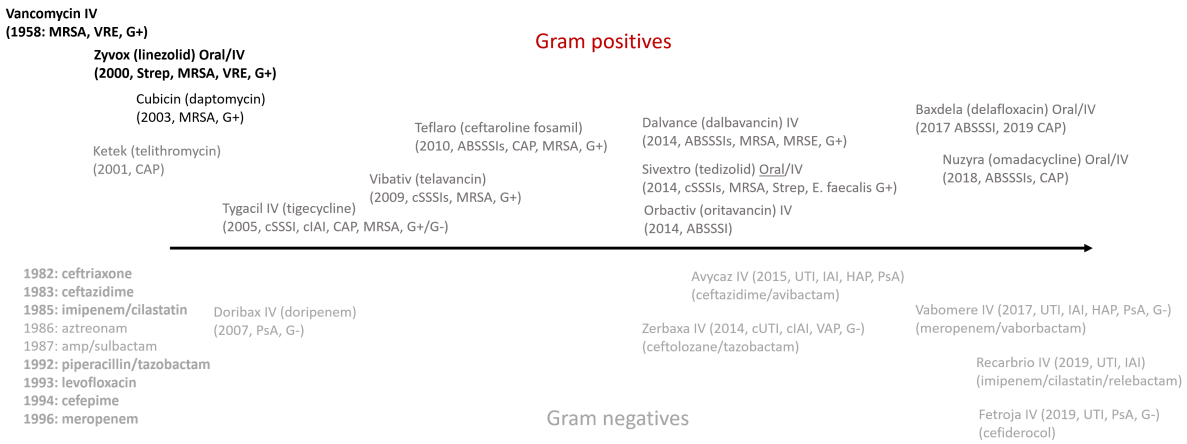
Årdal, C., M. Balasegaram, R. Laxminarayan, D. McAdams, K. Outterson, J. H. Rex, and N. Sumpradit (2019): "Antibiotic development — economic, regulatory and societal challenges," *Nature Reviews Microbiology 2019 18:5*, 18, 267–274, publisher: Nature Publishing Group.

A Appendix

This appendix contains supplementary materials.

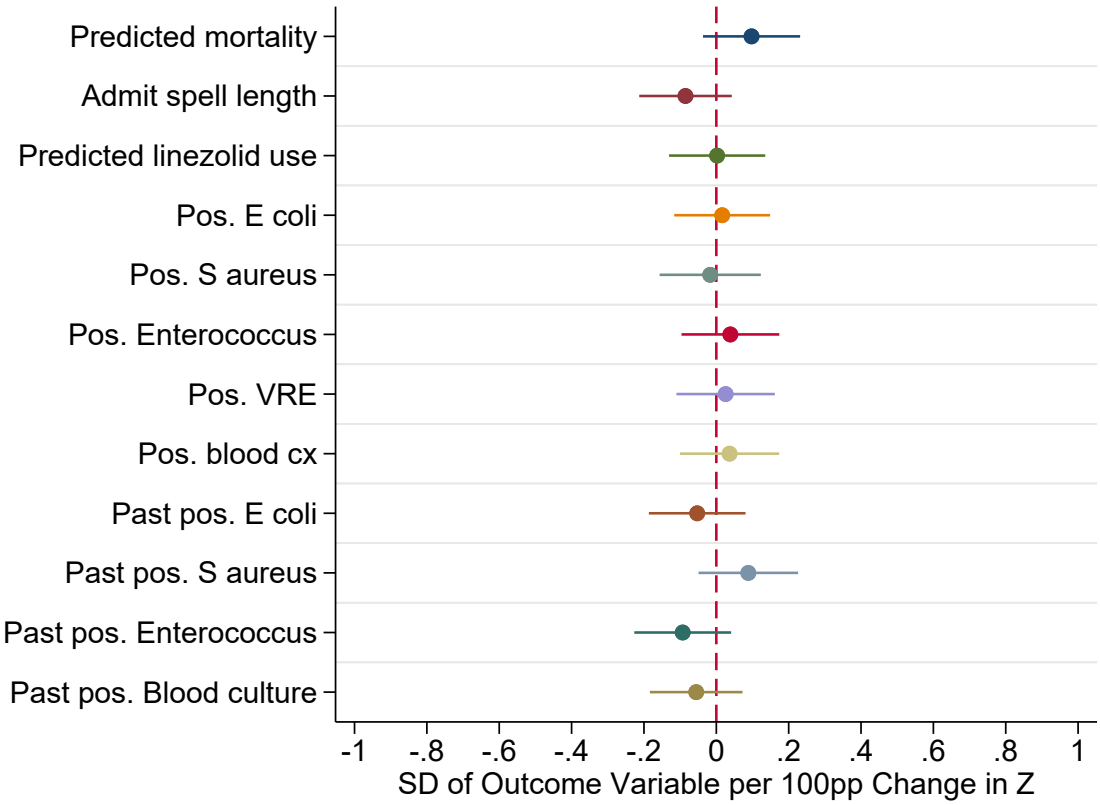
B Additional Exhibits

Figure A1: Timeline of Antibiotic Innovation



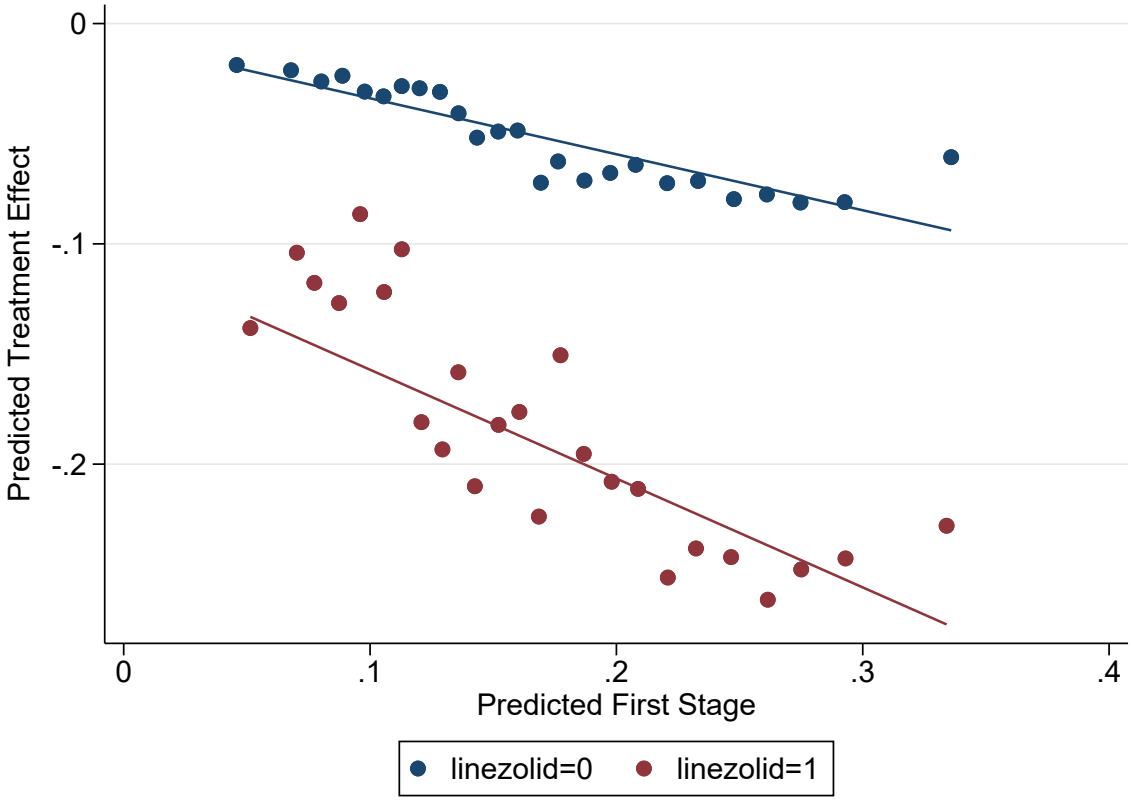
Notes: Figure shows approval dates and indications for antibiotics treating primarily gram-positive bacteria (above the horizontal line) and gram-negative bacteria (below the horizontal line). Bolded antibiotics are those most commonly used in our sample.

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.

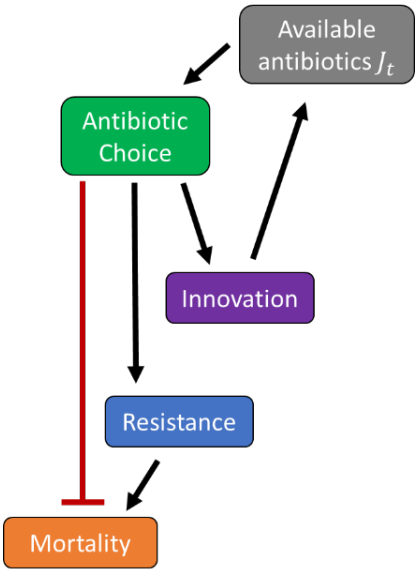
Figure A3: Mortality Model: Treatment Effects by Predicted Linezolid Use



Notes: Figure shows a binned scatterplot of model predictions from Equation 13. 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.

C Model Appendix

Figure A4: 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).

C.1 Derivations for Graphical Model

This section presents derivations for the marginal benefit and cost curves shown in Figure 6. 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 $dR/dD > 0$ and $dR/dI < 0$. 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 cost function includes mortality (scaled by , the value of statistical life) plus innovation costs cI, which we take to be linear in R&D for simplicity:

$$\lambda M(D, R) + cI. \tag{18}$$

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 is:

$$\lambda M(D,R) + pD. \tag{19}$$

Next, we derive optimal policies, which equate marginal benefits and costs, for several cases corresponding to the panels of Figure 6.

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 (18) with respect to demand yields:

$$-\lambda \left(\underbrace{\frac{\partial M}{\partial D}}_{\text{Direct Mortality Effect}} + \underbrace{\left[\frac{\partial M}{\partial R}\frac{\partial R}{\partial I} + c \right] \frac{dI}{dD}}_{\text{Innovation Externality}} \right) = \underbrace{\lambda \frac{\partial M}{\partial R}\frac{dR}{dD}}_{\text{Resistance Externality}} . \tag{20}$$

The left hand side represents marginal benefits of increasing antibiotic use, which includes the direct mortality effects as well as the innovation externality. The right hand side represents social marginal costs, which comprise the resistance externality. This demonstrates the opposing externalities that exist when only one policy instrument (demand) is available to increase revenues.

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 6.

$$\underbrace{-\lambda \frac{\partial M}{\partial D}}_{\text{Direct Mortality Effect}} = \underbrace{\lambda \frac{\partial M}{\partial R}\frac{dR}{dD}}_{\text{Resistance Externality}} + \underbrace{p}_{\text{Antibiotic Costs}} . \tag{21}$$

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 (18), 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: Panel C of Figure 6 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:

$$\text{FOC } D: \quad \underbrace{-\lambda \frac{\partial M}{\partial D}}_{\text{Direct Mortality Effect}} + \underbrace{\left[\frac{\partial M}{\partial R}\frac{\partial R}{\partial I} + c\right]\cancelto{0}{\frac{dI}{dD}}}_{\text{Innovation Externality}} = \underbrace{\lambda \frac{\partial M}{\partial R}\frac{dR}{dD}}_{\text{Resistance Externality}}$$

$$\text{FOC } P: \quad \underbrace{\left[\frac{\partial M}{\partial R}\frac{\partial R}{\partial I} + c\right]\frac{dI}{dP}}_{\text{Innovation Externality}} = 0 \tag{22}$$

where FOC D is the same as in (20), except we have plugged in FOC P, which cancels out the innovation externality. The opposing externalities of resistance and innovation have been decoupled: demand only needs to trade off mortality and resistance. Optimal innovation payments increase until the benefits of lower resistance equal the marginal cost of further R&D.

D Structural Model Appendix

D.1 Demand Model - Supplementary Exhibits

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.

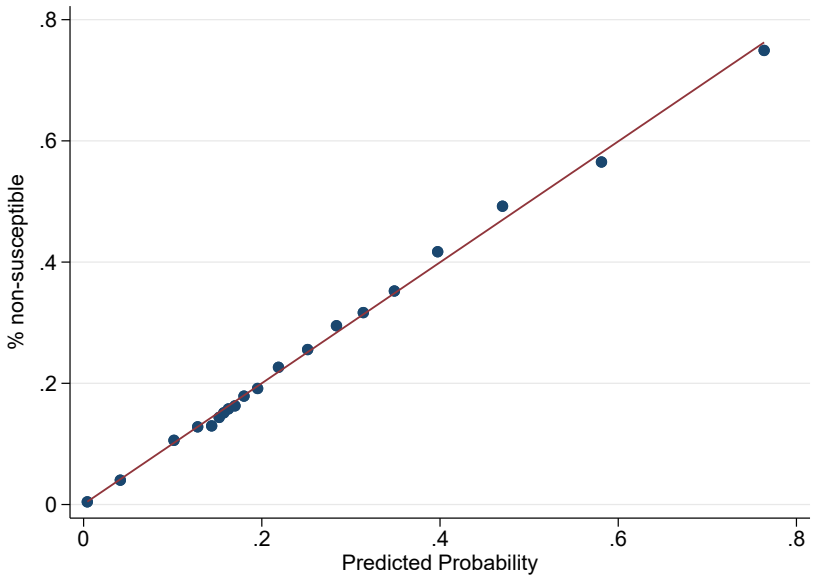
D.2 Resistance Model - Supplementary Exhibits

Table A4: Descriptive Statistics (Resistance)

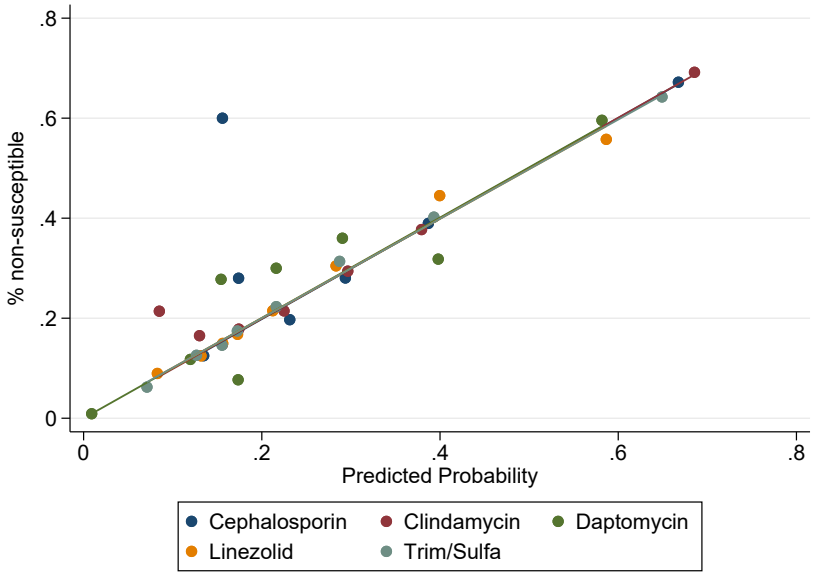
		Current	Past Resistance			Past Use		
Drug	N	Resistance	Own	Higher	Lower	Own	Higher	Lower
Cephalosporin	9,422	0.52	0.15	0.0648		0.26	0.14	
Clindamycin	81,982	0.40	0.15	0.016	0.097	0.083	0.13	0.21
Trim/Sulfa	250,279	0.23	0.16	0.022	0.13	0.19	0.13	0.16
Vancomycin	43,082	0.24	0.11	0.0066	0.15	0.37	0.058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

Figure A5: Resistance Model Fit



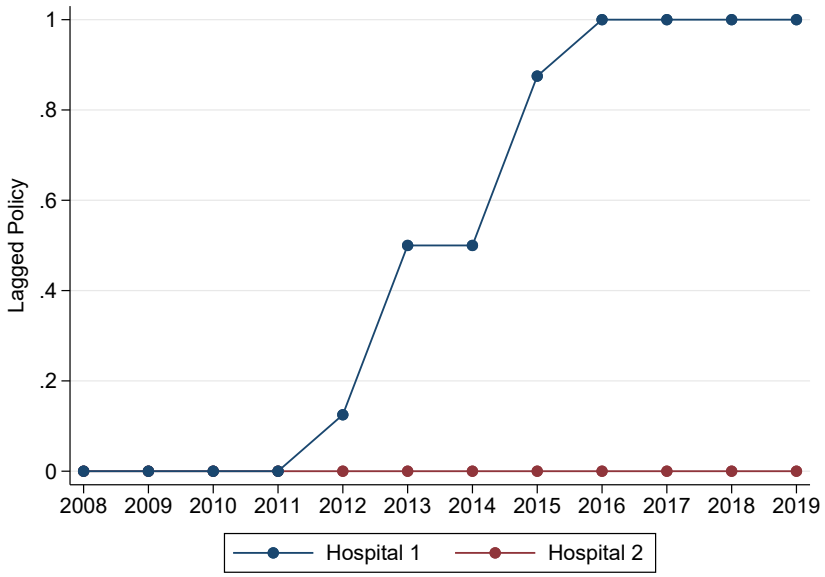
(a) All antibiotics



(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 A6: Lagged ASP Policy Definition



Notes: Figure shows the average lagged ASP policy variable (Y-axis) over time, separately for the treatment hospital (Hospital 1) and control hospital (Hospital 2). The lagged ASP policy variable is defined for each hospital-year-month 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 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.3 Mortality Model - Supplementary Exhibits

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).

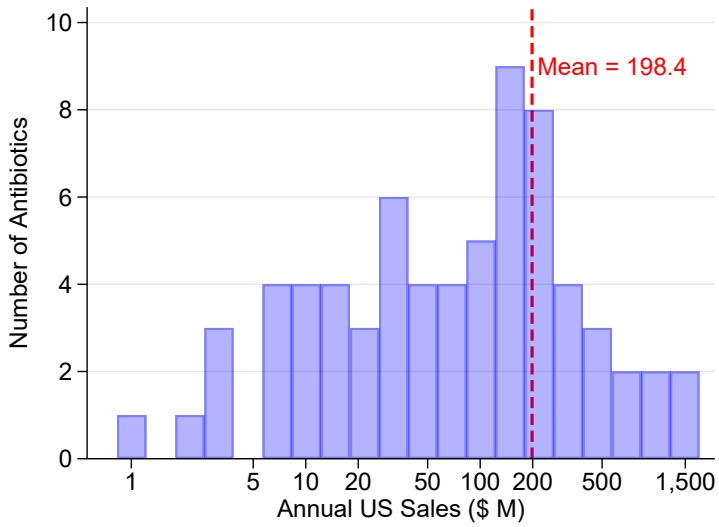
Table A7: Deciles of Predicted Linezolid Use

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.

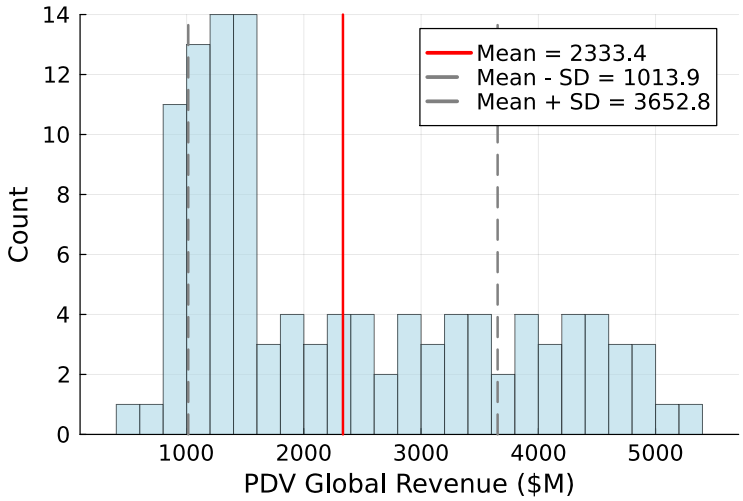
D.4 Innovation Model - Supplementary Exhibits

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



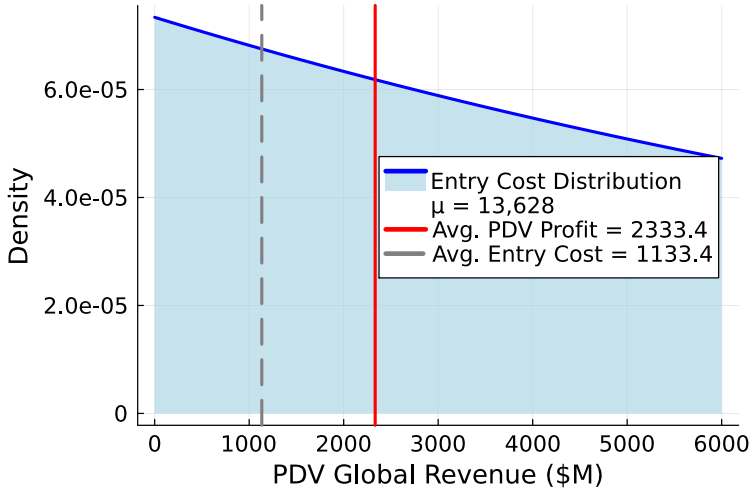
Notes: Histogram shows the distribution of annual branded U.S. sales across all N=69 antibiotics in the estimation sample (in millions of 2024 dollars). The dashed red vertical line marks the sample mean. Sales are plotted on a logarithmic scale for readability.

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 6.6. 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.

D.5 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 $\varepsilon^R_{el} \sim U[0,1]$, 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, $\varepsilon^D_{ij} \sim \text{Gumbel}(0,1)$, 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. [39] 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 6.5. 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.

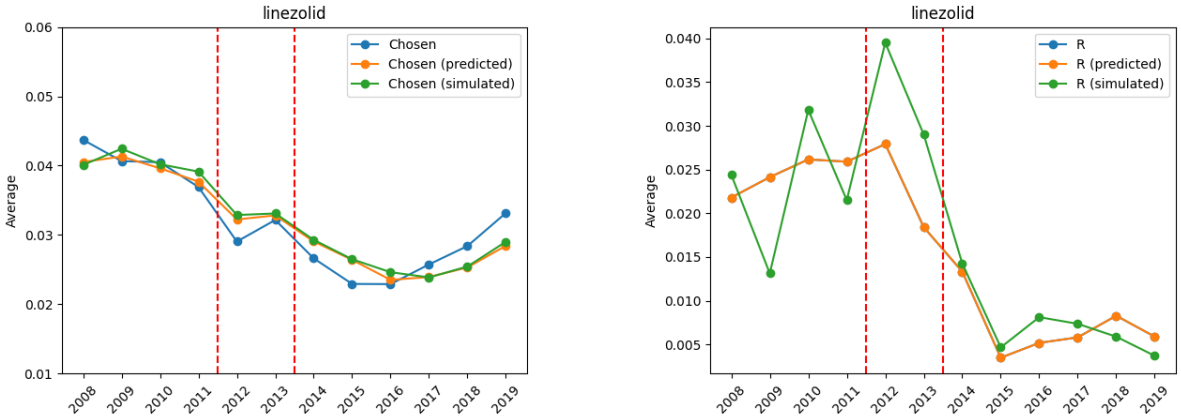
Innovation. We assume that one potential entrant decides whether to enter in January 2009. This potential entrant draws an *iid* private entry cost from the estimated exponential entry cost distribution from Section 6.6. The new antibiotic enters the market if the present discounted value (PDV) of future profits (over a 10-year exclusivity period) exceeds the drawn entry cost.

[39] 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.

D.6 Simulations - Supplementary Exhibits

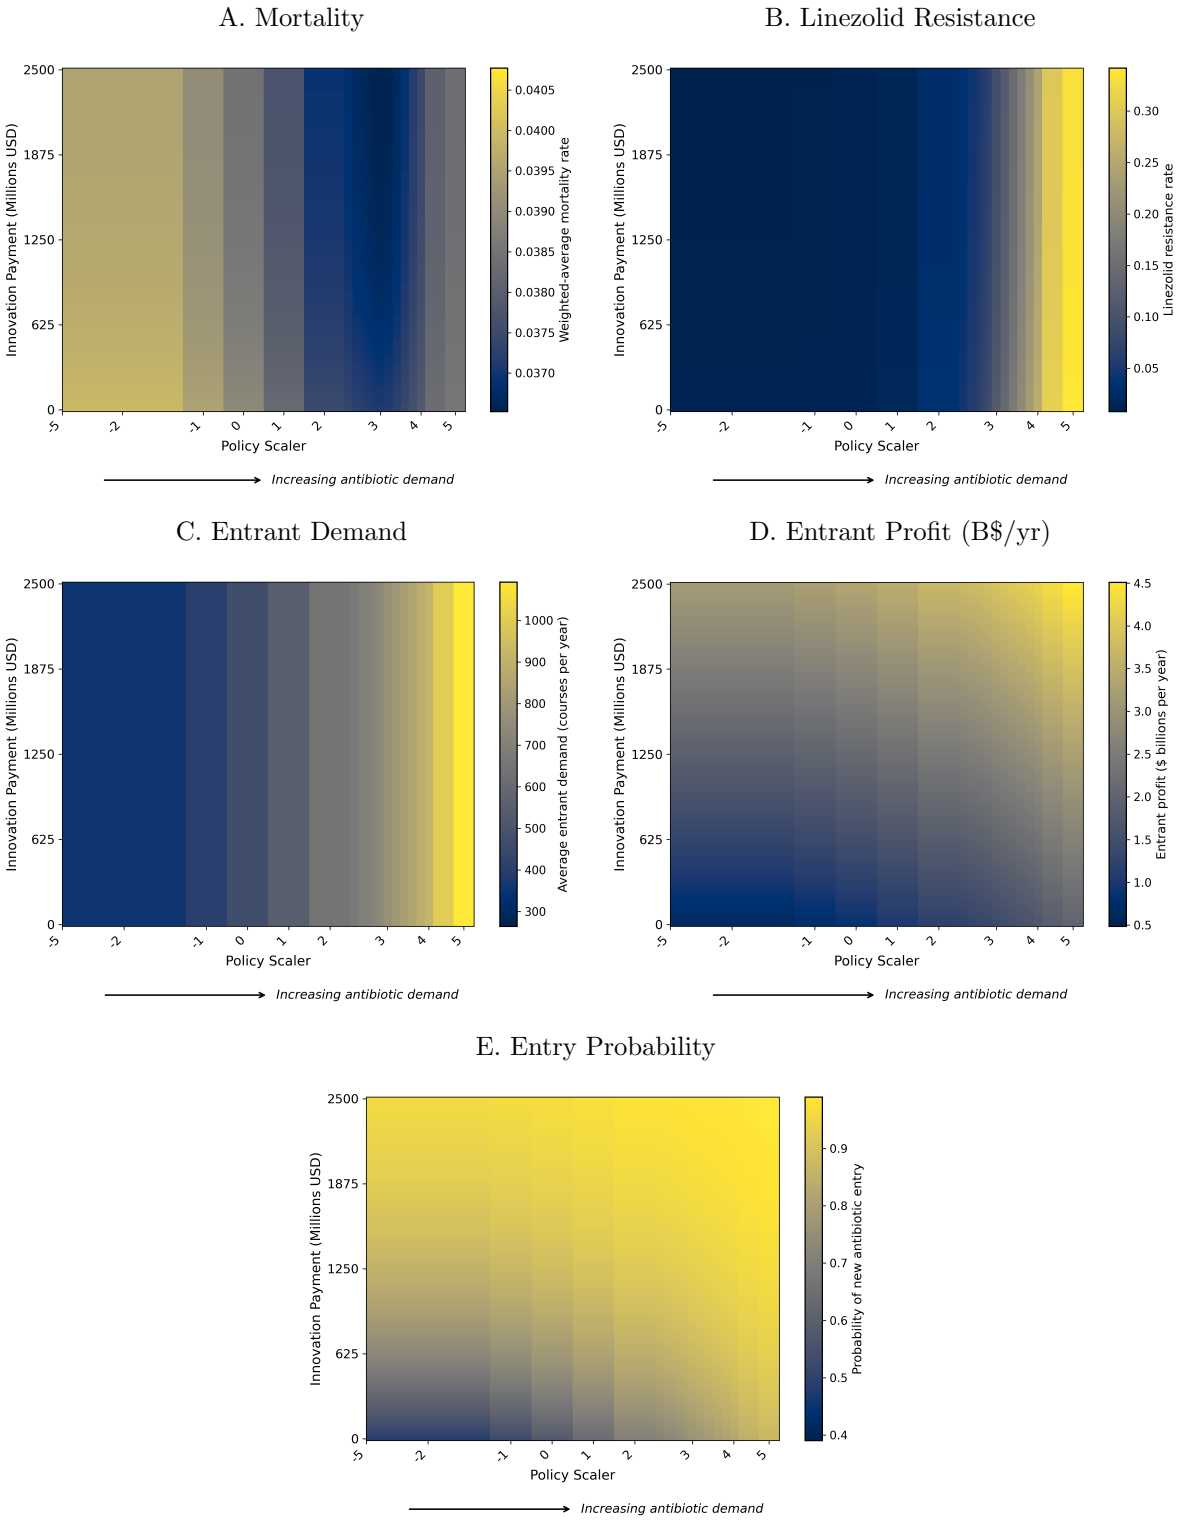
Figure A10: Model Fit: Linezolid Demand and Resistance in Hospital 1

(a) Demand (b) Resistance



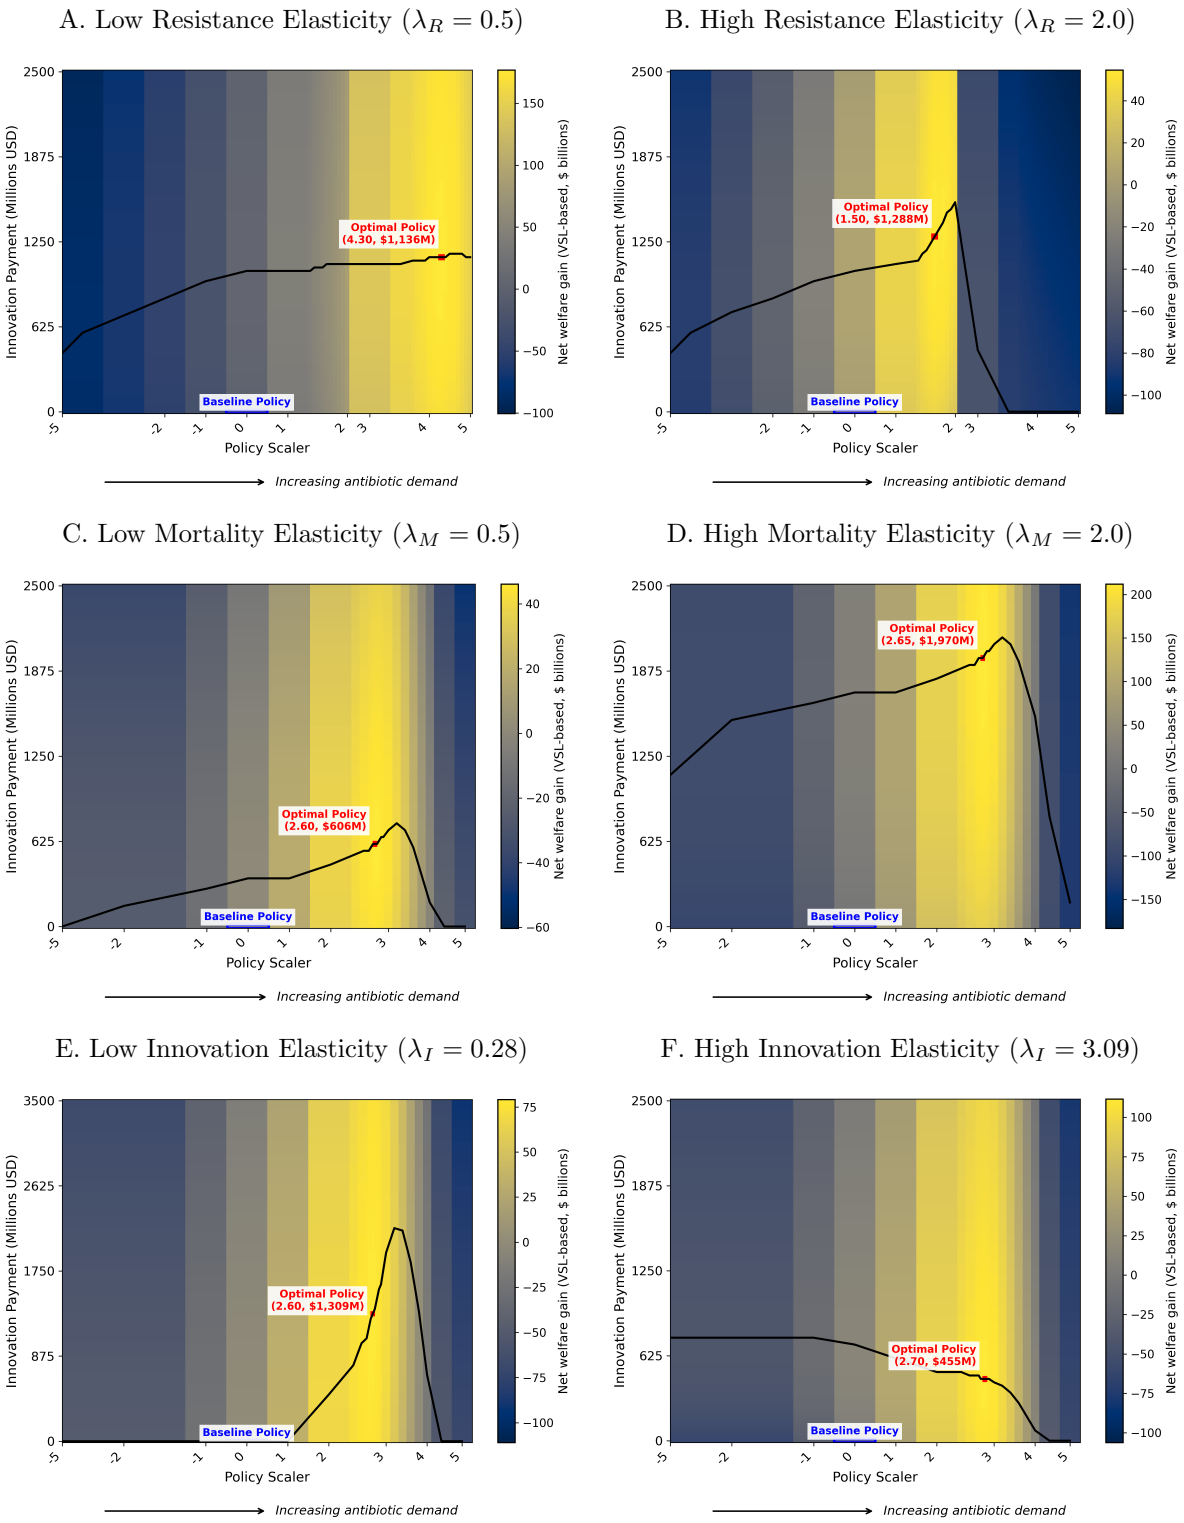
Notes: Figure demonstrates model fits for linezolid demand and resistance in the treated hospital (Hospital 1). For each outcome, blue, orange, and green lines show yearly averages of the actual data, model predictions from (6) and (7), and simulated demand from forward simulations, respectively. Panel (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 A11: Policy Effects on Mortality, Demand, Entry, and Resistance



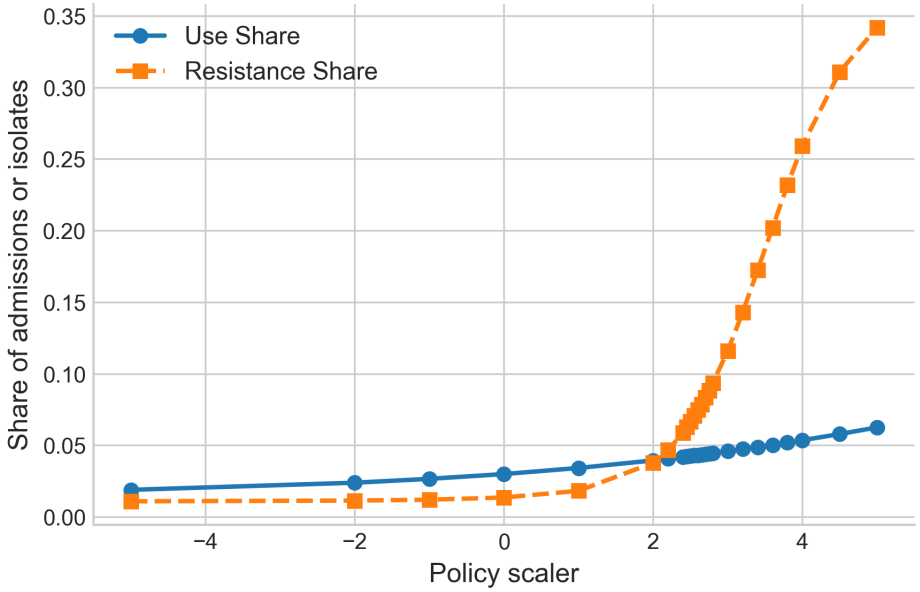
Notes: Each panel shows model-simulated outcomes across the two policy dimensions: the magnitude of the policy scaler (x-axis; higher values correspond to higher antibiotic demand) and the annual innovation payment (y-axis, in millions of U.S. dollars). Color intensity reflects the magnitude of each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

Figure A12: 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 A13: 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.