Optimal Conditional Drug Approval

Giuseppe Lopomo, David Ridley, Peng Sun, & Chenxi Xu*

December 2, 2024

Abstract

Regulators often allow firms to sell new drugs based on limited testing, with final approval dependent on confirmatory tests. We model optimal approval policies where firms have private information about testing costs, and the regulator's payoff depends on expected effectiveness. We find it may be optimal to conditionally approve drugs with low expected effectiveness, limit conditional approval to a share of patients, expand conditional approval share when the monopoly pricing period is fixed, and commit to a lower efficacy threshold for final approval. Optimal policies balance early access, innovation incentives, and early price reductions against the costs of ineffective drugs.

*Lopomo: glopomo@duke.edu. Ridley: david.ridley@duke.edu. Sun: psun@duke.edu. Xu: chenxi.xu@duke.edu. The authors are affiliated with Duke University's Fuqua School of Business, Durham, NC 27708-0120. The authors appreciate feedback from Angelique Acquatella, Pierre Dubois, Jonathan Skinner, Margaret Kyle, Paul Novosad, seminar participants at the Congressional Budget Office, Dartmouth College, Duke University, and the Toulouse School of Economics, and conference participants at the CEPR-TSE Health Economics Conference and the INFORMS annual meeting. The authors received no external funding for this research.

1 Introduction

Drug regulators frequently grant conditional approvals for new drug indications (uses) based on limited evidence, with final approval contingent on additional confirmatory testing conducted after market entry. Conditional approval programs exist in jurisdictions including the United States, the European Union, and Japan.[1] Over the past decade, the United States Food and Drug Administration (FDA) has granted conditional approval to 80 new drugs (Figure 9).

Critics of conditional approval argue that it promotes spending on therapies with insufficient evidence. For instance, the 2021 conditional approval of an Alzheimer's disease indication prompted several FDA advisory committee members to resign, citing inadequate efficacy data and warning that billions could be wasted (Belluck and Robbins, 2021). A similar controversy surrounded the 2016 approval of an indication for Duchenne Muscular Dystrophy with limited supporting evidence (Florko, 2021). However, proponents emphasize that conditional approval can expedite access to potentially life-saving treatments. For example, an FDA leader responsible for gene therapies has argued that the risk of early market entry is justified for patients with severe, life-threatening conditions (Feuerstein, 2024).

Conditional approvals not only enable earlier access for patients but also encourage investment in drug development, particularly for diseases with limited commercial appeal. Indications for slowly progressing diseases often involve lengthy and costly testing, deterring investment (Budish, Roin, and Williams, 2015). More broadly, stringent regulatory requirements and high testing costs can reduce innovation and social welfare (Peltzman, 1973; Philipson and Sun, 2008). While confirmatory testing generates critical evidence, it also imposes financial burdens on firms, discouraging market entry (Grennan and Town, 2020).

We present a mechanism design framework to optimize conditional approval policies. The scale of conditional approval determines the firm's incentive to conduct confirmatory testing. While regulators such as the FDA cannot provide direct subsidies to firms, they can enable early market access. However, the extent of the subsidy is constrained by the market size of the indication.

If testing costs were publicly observable, regulators could offer just enough conditional approval to make testing profitable. However, because these costs are privately known to

[1]Examples include the FDA's Accelerated Approval program, the EU's Conditional Marketing Authorization, and Japan's Sakigake Designation System. The FDA's Emergency Use Authorization for the first COVID-19 vaccines is similar.

firms, optimal policies must be incentive-compatible. Our model shows that the regulator's policy resembles a take-it-or-leave-it offer: the firm receives conditional approval conditional on conducting further testing, and final approval depends on the test outcomes.

The analysis yields five key recommendations. Regulators should: (1) not be deterred from approving indications with low efficacy probabilities, provided safety standards are met; (2) lower the bar for conditional approval but limit market entry through narrower indications or less generous insurance coverage; (3) grant higher conditional approval shares for drugs with fixed-duration monopolies, because these drugs earlier price reductions; (4) commit to lowering the bar for *final* approval (in addition to conditional approval) to promote testing; and (5) not necessarily grant a higher approval share for indications narrowly missing efficacy probability thresholds, as these are likely to be tested regardless.

Conditional approval provides the regulator with option value, as it enables testing of less promising indications that could later prove beneficial. Far from being wasteful, conditional approval can accelerate the end of monopoly pricing for fixed-duration exclusivities such as orphan drug (7 years) and biologic exclusivities (12 years), or for top-selling drugs subject to Medicare price controls.[2] However, firms will resist earlier monopoly expiration unless conditional approval shares are sufficient. Thus, optimal policies are more generous with conditional approval shares for drugs with fixed-duration monopolies.

This paper contributes to the literature on financing drug testing. Conditional approval, which provides indirect financing for drug development, complements other strategies. One approach to financing is direct funding for early-stage testing, known as "push funding." Agencies like the NIH typically fund low-cost, early-stage testing. Another approach is direct funding contingent on approval, known as "pull funding." During the COVID-19 pandemic, governments used both push and pull mechanisms to support vaccine testing. Snyder, Hoyt, and Gouglas (2023) identified optimal push and pull funding when testing costs are privately known.

For diseases with limited commercial appeal, such as tropical diseases in lower-income countries, advance market commitments can create markets by guaranteeing the purchase of approved products at predetermined prices (Berndt et al., 2007; Kremer, Levin, and Snyder, 2022). However, such commitments require substantial upfront investments, often exceeding a billion dollars. Other indirect incentives include transferable vouchers, such as those in the US priority review voucher program, which allow faster regulatory review of another product and can sell for around $100 million (Gans and Ridley, 2013). While these incentives can

[2]See stylized fact 9 in the appendix for details on exclusivity and pricing controls.

stimulate development, their effectiveness diminishes as more vouchers are issued (Ridley and Regnier, 2016). Similarly, vouchers for exclusivity extensions have been proposed but would delay generic competition (Dubois, Moisson, and Tirole, 2022).

Second, this paper contributes to the literature on optimal regulatory approval. Carpenter and Ting (2007) analyze a regulator-firm interaction where approval decisions rely on imperfect information. They examine the trade-offs between false positives and negatives in regulatory decisions. Conditional approval often depends on surrogate endpoints, which must be carefully validated to ensure reliability for regulators and payers (Bognar et al., 2017). Beyond conditional approval, regulators employ pathways like the FDA's "breakthrough" designation to expedite market entry (Chandra et al., 2024).

Permissive conditional approval policies may lead to withdrawals of drug indications if further testing reveals insufficient efficacy. Orlov, Skrzypacz, and Zryumov (2020) show that firms might delay reporting negative information to postpone withdrawals. Xu, Zhao, and Petruzzi (2021) demonstrate that firms balance regulatory costs by choosing between extended and shorter testing periods. To address delays, regulators like the FDA now require firms to initiate confirmatory testing before granting conditional approval. The dynamics differ for most medical devices, especially those similar to previous devices. Testing is less rigorous, markets see frequent entry, and competition is dynamic (Grennan and Town, 2020; Collard-Wexler, Grennan, and Steck, 2024).

Henry and Ottaviani (2019) model a continuous-time game where firms (informers) conduct costly experiments, and regulators (evaluators) decide on final approval based on outcomes. They show that optimal outcomes arise when the firm can make a single take-it-or-leave-it offer. Henry, Loseto, and Ottaviani (2022) allow for approval revocation when negative evidence emerges. The results highlight the importance of flexible regulatory policies.

McClellan (2022) considers a scenario where the informer has private information about quality. He shows that optimal policies are history-dependent. The initial threshold for final approval decreases if the evaluator's belief drops below a certain level.

While continuous-time models suit settings with adaptable protocols, late-stage human testing typically involves fixed endpoints, with outcomes revealed simultaneously. Accordingly, our analysis employs discrete time. Our model also differs by incorporating partial conditional approval and examining how the timing of price reductions influences optimal regulatory decisions. These features enable calibration to a specific case. We examine avelumab, a conditionally approved cancer treatment (Appendix C).

2 Basic model

Consider the following two-period game between a regulator and a pharmaceutical firm. At the beginning of the game, the firm has already conducted preliminary clinical testing for a new drug. Based on the outcome of these tests, both parties believe that the drug is efficacious with probability $\gamma \in (0,1)$. To ascertain the true efficacy of the drug $\eta \in \{0,1\}$, the firm must conduct a final round of confirmatory tests which cost c. The firm knows the value of c privately, while the regulator believes that c is drawn from a distribution with cumulative distribution function F and positive density f on its support $[c_L, c_H]$.

The game begins with the regulator committing publicly to a regulatory policy. Following the revelation principle (Myerson, 1981), it is without loss of generality to consider direct revelation mechanisms, that is, mechanisms under which the regulator asks the firm to report their private cost c, while making sure that the mechanism provides the incentive for the firm to report truthfully. In our setting, such a mechanism specifies i) a probability of requesting confirmatory testing $\tau(c) \in [0,1]$, and ii) a scale of conditional approval $a_1(c) \in [0,1]$ contingent on the firm conducting confirmatory testing, both of which depend on the firm's reported cost value $c \in [c_L, c_H]$. If the regulator does not request testing, or if the firm declines the request, the game ends immediately with both parties earning zero payoffs. Otherwise, the firm enters the market in period one at scale a_1 while testing is taking place. In practice, the regulator does not need to ask the firm to directly reveal its testing cost c. As we will show later in the paper, the optimal mechanism can be implemented as a "take-it-or-leave-it" offer, such that the firm tests the drug if and only if its cost c is lower than a threshold. In period two, the regulator observes the outcome of the final test and grants final approval if and only if $\eta = 1$.

If the firm is granted conditional approval at scale a_1, its payoff in period one is $a_1 \cdot \pi \cdot q$, where π is the net revenue per dose, and q is the total quantity sold during the conditional approval period.

The firm's total expected payoff after the drug's final approval depends on its revenue from period one. Drug sales tend to rise over time as doctors and patients learn about the drug through experience and word of mouth.[3] Hence entering the market in period one at a larger scale increases later sales and profit in period two. We assume that the total discounted quantity sold in period two is $\lambda_0 + \lambda_1 \cdot a_1$, where λ_0 and λ_1 are fixed parameters. Therefore the total revenue in period two is given by $\pi \cdot (\lambda_0 + \lambda_1 \cdot a_1)$. When the monopolistic period

[3] A drug's revenue often peaks in year six (Robey and David, 2016).

ends, generic drugs are allowed to enter the market, and the firm's net revenue becomes 0.

The firm's total expected profit is given by the testing probability multiplied by its total revenue across both periods minus the testing cost. Therefore, if a firm with true cost c reports c' to the regulator, its expected total profit is

$$\hat{\Pi}(c, c') \equiv \tau(c') \cdot \Big[\pi \cdot q \cdot a_1(c') - c + \gamma \cdot \pi \cdot [\lambda_0 + \lambda_1 \cdot a_1(c')]\Big]. \tag{1}$$

The value of a drug depends on whether it received conditional approval. Conditionally-approved drugs are more valuable when the monopoly-pricing period has a fixed duration (due to price reductions) and when greater word-of-mouth (λ_1) increases the number of beneficiaries.

To be feasible any regulatory policy must satisfy incentive compatibility, meaning that the firm should not benefit from reporting any cost value other than its true cost c, i.e.,

$$\forall c, c' \in [c_L, c_H]: \qquad \hat{\Pi}(c, c) \geq \hat{\Pi}(c, c'). \tag{IC}$$

Also, because the firm can earn zero profit by declining the request to conduct the final round of tests, the following participation constraints must hold:

$$\forall c \in [c_L, c_H]: \qquad \hat{\Pi}(c, c) \geq 0. \tag{IR}$$

Finally, the probability of testing τ, as well as the scale of conditional approval a_1 must be between 0 and 1:

$$\forall c \in [c_L, c_H]: \qquad 0 \leq \tau(c) \leq 1, \tag{2}$$

$$\forall c \in [c_L, c_H]: \qquad 0 \leq a_1(c) \leq 1. \tag{3}$$

The regulator is risk-neutral, and its payoff per dose $v(\eta)$ depends on the efficacy level. We assume that there exists a threshold $\hat{\eta}$ such that

$$v(\eta) = \begin{cases} -v_L < 0, & \text{if } \eta < \hat{\eta}, \\ v_H > 0, & \text{if } \eta \geq \hat{\eta}. \end{cases} \tag{4}$$

The regulator's payoff reflects the drug's value net of its price. Even an effective drug can have a negative value due to the payment from the regulator to the firm.[4]

[4]As described in the introduction, some doctors opposed conditional approval for a drug for Alzheimer's disease because the drug was a waste of money.

If market access is granted in period one at scale a_1, the regulator's expected payoff is $v_E \cdot q \cdot a_1$, where v_E is the expected value of $v(\eta)$ given by

$$v_E \equiv \gamma \cdot v_H - (1-\gamma) \cdot v_L. \tag{5}$$

Because final approval is granted with probability γ, the regulator's expected payoff in period two is $\gamma \cdot (\lambda_0 + \lambda_1 \cdot a_1) \cdot v_H$.

At the end of period two, the firm loses its monopoly-pricing power due to generic entry or the initiation of government price controls. We assume that the price drops to marginal cost and profit becomes zero. For the firm, this competitive period has no value. However, for the regulator, this competitive period is valuable for the low prices. Let Δ and σ represent the total discounted doses after period two and the price drop per dose, respectively. The regulator's payoff per dose increases to $v_H + \sigma$, and its total payoff during the competitive period can be denoted as $z \equiv \gamma \cdot \Delta \cdot (v_H + \sigma)$.

The regulator maximizes its expected payoff, given by the probability of testing multiplied by its total expected payoff across all three periods:

$$\tilde{U}_R(\tau, a_1) \equiv \int_{c_L}^{c_H} \tau(c) \cdot \Big[v_E \cdot q \cdot a_1(c) + \gamma \cdot [\lambda_0 + \lambda_1 \cdot a_1(c)] \cdot v_H + z \Big] \, \mathrm{d}F(c). \tag{6}$$

The regulator's utility does not account for the firm's profit. If the specification were adjusted to include shareholder welfare in the regulator's utility, the value of conditional approval would rise.

The regulator's problem is

$$\mathcal{U}_R \equiv \max_{\tau, a_1 : \text{(IC)},\text{(IR)},(2),(3)} \tilde{U}_R(\tau, a_1). \tag{7}$$

Without loss of generality, we can normalize the quantity sold in the conditional approval period q to 1. After doing so, the constants π and v_H only appear in each party's payoff function, and thus we can further normalize both π and v_H to 1, without loss of generality. Both parties' payoff functions boil down to

$$\hat{\Pi}(c, c') = \tau(c') \, \Big\{ a_1(c') - c + \gamma \cdot [\lambda_0 + \lambda_1 \cdot a_1(c')] \Big\}$$

and

$$\tilde{U}_R(\tau, a_1) = \int_{c_L}^{c_H} \tau(c) \cdot \Big[v_E \cdot a_1(c) + \gamma \cdot [\lambda_0 + \lambda_1 \cdot a_1(c)] + z \Big] \, \mathrm{d}F(c).$$

It is convenient now to reformulate the regulator's problem (7) as a linear optimization problem, by introducing the following new variables:

$$\forall c \in [c_L, c_H]: \quad \alpha_1(c) \equiv \tau(c) \cdot a_1(c). \tag{8}$$

The feasibility condition for α_1 in (8) can be written as

$$\forall c \in [c_L, c_H]: \quad 0 \leq \alpha_1(c) \leq \tau(c). \tag{9}$$

The firm's total expected profit function in (1) becomes

$$\hat{\Pi}(c, c') = (1 + \lambda_1 \cdot \gamma) \cdot \alpha_1(c') - (c - \lambda_0 \cdot \gamma) \cdot \tau(c'), \tag{10}$$

and the regulator's expected payoff in (6) becomes

$$U_R(\tau, \alpha_1) \equiv \int_{c_L}^{c_H} [(\lambda_0 \cdot \gamma + z) \cdot \tau(c) + (v_E + \lambda_1 \cdot \gamma) \cdot \alpha_1(c)] \, \mathrm{d}F(c). \tag{11}$$

The regulator's problem can now be expressed as the linear optimization problem

$$\begin{gathered} \max_{\tau, \alpha_1} \; U_R(\tau, \alpha_1), \\ \text{subject to (IC), (IR), (2), and (9)}, \end{gathered} \tag{12}$$

where the firm's payoff function $\hat{\Pi}$ is now defined by equation (10).

We assume that the testing cost never exceeds the maximum total expected profit of the firm, i.e.[5]

Assumption 1. $c_H \leq 1 + \gamma(\lambda_0 + \lambda_1)$.

[5]This is without loss of generality because the firm's total expected profit is bounded by $1 + \gamma(\lambda_0 + \lambda_1)$. Thus any firm with cost $c > 1 + \gamma(\lambda_0 + \lambda_1)$ cannot be incentivized to test in any feasible regulatory policy. For any distribution with support $[0, \infty)$, density $\tilde{f}$, and cdf $\tilde{F}$, we can consider the truncated distribution with support $[0, 1 + \gamma(\lambda_0 + \lambda_1)]$ and cdf

$$F(c) = \min\left\{ \frac{\tilde{F}(c)}{\tilde{F}(1 + \gamma(\lambda_0 + \lambda_1))}, 1 \right\}, \quad \forall c \in \mathbb{R}.$$

2.1 Discussion of the model

In this section, we discuss several features of the model that correspond to stylized facts which are described in the appendix (section B).

In our model, the interaction between the regulator responsible for drug approval and the firm testing a new drug unfolds over two periods. Period 1 begins at time 0, and period 2 begins at time t_1 (Figure 1).

At time 0, the firm has already conducted preliminary tests, which are informative but insufficient to yield statistically significant results. Consequently, the regulator's decision to approve the drug is complicated by residual uncertainty about the efficacy of the drug (stylized fact 1). At time 0, the regulator can grant conditional approval in exchange for confirmatory testing (stylized fact 2), allowing the drug to be sold to a fraction of the market (stylized fact 3).

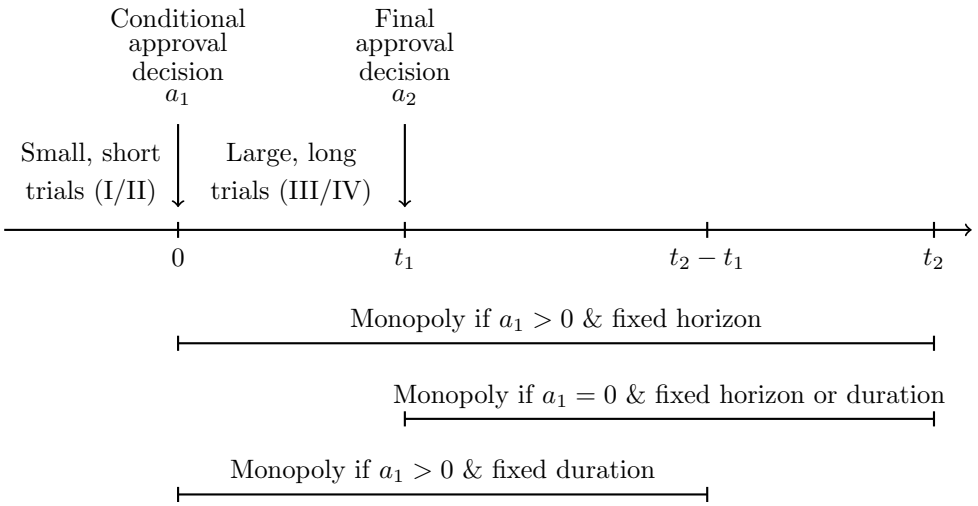

Figure 1: Timeline for testing, regulatory review, and falling prices. After preliminary testing, the regulator may grant conditional approval for a portion of the population (a_1). Later, the regulator can grant full approval (a_2). Source: Authors.

The firm must truthfully disclose the outcomes of all of its clinical trials (stylized fact 4). Therefore, throughout the game, the firm's belief about the drug's efficacy (η) is always shared with the regulator. The loss from approving a bad drug v_L is primarily financial due to wasted money on an ineffective drug (stylized fact 5), but can also include harm to patients.

While there is no private information about efficacy, there may be private information regarding the cost of confirmatory testing. It is difficult to attribute joint testing costs to a

single drug when drug development often involves multiple drugs (stylized fact 6).

If there is no confirmatory testing, either because the regulator does not request it or the firm declines, the game ends with both parties earning zero payoffs. If instead the regulator requests testing and the firm accepts, the firm can start selling the drug in period 1, up to the amount conditionally approved. Testing is the only source of information about efficacy; they do not learn more about the drug's efficacy from sales (stylized fact 7).

The firm has monopoly power until generic entry is allowed or price controls are imposed. The end of monopoly power depends on the drug's characteristics, such as whether it is a small or large molecule, treats a rare or common disease, and has a recent or late patent filing (stylized fact 8). These factors determine whether the monopoly pricing period has a "fixed horizon" or a "fixed duration" (stylized fact 9). For fixed-horizon drugs, the monopoly-pricing period always ends at t_2 (Figure 1). For fixed-duration drugs, the length of the monopoly-pricing period is always $t_2 - t_1$. Hence, if a fixed-duration drug is conditionally approved, its monopoly-pricing period ends earlier.

Figure 2 shows how annual revenue and cumulative profit evolve over time, depending on the conditional approval decision and whether the monopoly-pricing period has a fixed horizon or fixed duration.

Doctors and patients gradually become aware of the drug, hence revenue rises gradually in the years following approval (stylized fact 10), as shown in the top panel of Figure 2. Because of rising revenue, the firm benefits from conditional approval not just because it can begin selling earlier, but also because its later sales are higher.

For drugs with high testing costs, profit can be positive only if the drug is conditionally approved at a large scale and has a fixed horizon (stylized fact 11), as illustrated in Figure 2(b). The cumulative profit in the case of high testing costs without conditional approval remains negative for the duration of the game.

The regulator makes three decisions: (i) whether the firm should conduct confirmatory testing, (ii) at what scale the firm can enter the market before final approval, and (iii) where to set the efficacy threshold for final approval.

In section 4, we consider the case where the regulator's payoff depends continuously on the drug's efficacy and thus may benefit from lowering the efficacy threshold for final approval, in order to motivate more firms to invest in testing.

We derive optimal regulatory policies when the firm knows its confirmatory testing cost privately. The optimal regulatory policy is essentially a take-it-or-leave-it offer: the firm is denied final approval at time 0, unless it conducts confirmatory testing in exchange for early

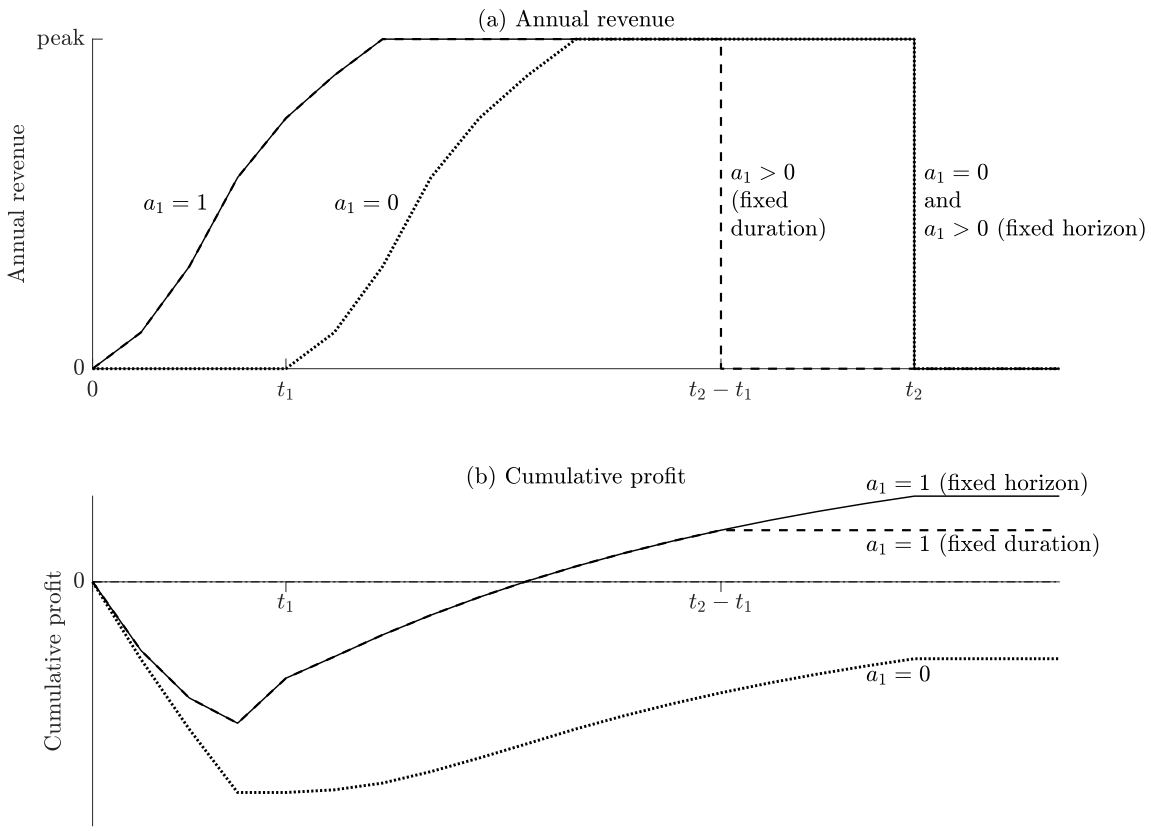

Figure 2: Annual revenue (top) and cumulative profit (bottom) depend on (i) whether the regulator conditionally approves ($a_1 > 0$) or not ($a_1 = 0$) and (ii) whether the monopoly-pricing period has a fixed horizon (solid line) or a fixed duration (dashed line). Source: Authors' analysis based on a representative drug.

market access at an optimally-chosen scale, as shown in the next section.

3 Optimal policies

In this section, we present the regulator's optimal mechanism, which maps the firm's private information on the testing cost to the regulatory decisions. In particular, we study how the optimal mechanism depends on the drug's probability of being efficacious.

Initially, we assume the regulator can grant partial conditional approval. Later, we restrict the regulator to granting only full conditional approval or none, a scenario we refer

to as integer conditional approval (section 3.1).

We also begin by assuming that the duration of the monopoly pricing period is unaffected by conditional approval. Later, we show how the optimal mechanism changes when conditional approval accelerates the end of the monopoly pricing period, thereby reducing the benefit of conditional approval for the firm (section 3.2).

For a drug that is likely to be efficacious, the regulator will grant full-scale conditional approval. Under conditional approval, the firm will always be willing to test, because we assume testing costs are not prohibitively high (Assumption 1). For a drug that is unlikely to be efficacious, the regulator can encourage the firm to bear the testing costs by offering partial conditional approval.

The regulator's problem can be framed as that of a budget-constrained buyer facing a single supplier with a privately known cost. The regulator is interested in learning the drug's efficacy η. The scale of conditional approval, a_1, serves as the payment method used by the regulator to compensate the firm.

Recall that we use three terms regarding a drug's likely goodness: η is the drug's efficacy, γ is the probability that the drug is efficacious, and v_E is the expected value of the drug to the regulator. Additionally, we use two terms for the scale of conditional approval: a_1, which represents the scale, and α_1, which is a_1 multiplied by whether the firm tests, τ. The values of a_1 and α_1 are equal to each other if and only if τ is either 0 or 1.

Without any information about the drug, if the expected value to the regulator is negative ($v_E < 0$), it would be optimal to deny final approval. If the drug's efficacy η becomes known, final approval is granted only if the drug is efficacious ($\eta = 1$). Therefore, the regulator's willingness to pay for this information is the change in expected payoff, $\lambda_0 \gamma + z$. The firm's net cost of testing is $c - \lambda_0 \gamma$, because the regulator grants final approval with probability γ if the firm conducts the test.

The following lemma is commonly used to simplify mechanism design problems involving direct payments between buyers and sellers. For completeness, a proof is provided in the appendix.

Lemma 1. *The incentive constraints* (IC) *and* (IR) *can be replaced by*

$$\tau \text{ non-increasing}, \tag{13}$$

$$\alpha_1(c) = \frac{1}{1 + \lambda_1 \cdot \gamma} \cdot \left[\Pi(c_H) + \int_c^{c_H} \tau(y) \, \mathrm{d}y + (c - \lambda_0 \cdot \gamma) \cdot \tau(c) \right], \tag{14}$$

and

$$\Pi(c_H) \geq 0, \tag{15}$$

where $\Pi(c_H) \equiv \hat{\Pi}(c_H, c_H)$.

Proof. See Appendix D. ■ □

The "envelope" condition in (14) allows us to eliminate α_1 from the problem and rewrite the regulator's objective function in terms of the "virtual valuation" function w defined by

$$w(c;\gamma) \equiv (\lambda_0 \cdot \gamma + z) + \frac{v_E + \lambda_1 \cdot \gamma}{1 + \lambda_1 \cdot \gamma} \left[c - \lambda_0 \cdot \gamma + \frac{F(c)}{f(c)} \right]. \tag{16}$$

The second best problem (12) can now be rewritten as

$$\max_{\{\tau, \Pi(c_H)\} \in \hat{\Omega}_B} \frac{v_E + \lambda_1 \cdot \gamma}{1 + \lambda_1 \cdot \gamma} \cdot \Pi(c_H) + \int_{c_L}^{c_H} w(c;\gamma) \cdot \tau(c) \, \mathrm{d}F(c), \tag{17}$$

where the feasible set $\hat{\Omega}_B$ is determined by (2), (13), (15), and for any $c \in [c_L, c_H]$,

$$0 \leq \frac{1}{1 + \lambda_1 \cdot \gamma} \left[\Pi(c_H) + \int_c^{c_H} \tau(y) \, \mathrm{d}y + (c - \lambda_0 \, \gamma) \cdot \tau(c) \right] \leq \tau(c). \tag{18}$$

Constraint (18) is the feasibility condition for α_1.

In addition to Assumption 1, we introduce the following regularity assumption on the cost distribution.

Assumption 2. *The density function $f(c)$ is differentiable and non-increasing.*

Assumption 2 holds for various commonly used distributions, including the uniform distribution, exponential distribution, and Pareto distribution. Assumption 2 guarantees that the hazard rate $F(c)/f(c)$ is non-decreasing, and hence the virtual valuation $w(c;\gamma)$ is non-increasing in c when $v_E + \lambda_1 \, \gamma$ is negative, and non-decreasing otherwise. Moreover, $w(\cdot;\gamma)$ changes its sign at most once as c increases in $[c_L, c_H]$ for any given γ. Therefore, when $v_E + \lambda_1 \, \gamma < 0$, w is non-negative if $c < c_0(\gamma)$, and non-positive if $c_0(\gamma) < c$, where the threshold $c_0(\gamma)$ is given by

$$c_0(\gamma) \equiv \sup \{ c \in [c_L, c_H] : w(c;\gamma) \geq 0 \}. \tag{19}$$

The next proposition characterizes the optimal mechanism for the regulator's problem.

Proposition 1. *Under Assumptions 1 and 2, the optimal solution α_1^* and τ^* for the second best problem* (12) *is as follows:*

1. *if $\gamma \in [\check{\gamma}, 1]$, then $\alpha_1^*(c) = 1$ and $\tau^*(c) = 1$ for any $c \in [c_L, c_H]$;*

2. *if $\gamma \in [c_H/\lambda_0, \check{\gamma})$, then $\alpha_1^*(c) = 0$ and $\tau^*(c) = 1$ for any $c \in [c_L, c_H]$;*

3. *if $\gamma \in [0, \min\{c_H/\lambda_0, \check{\gamma}\})$ and $c_0(\gamma) \geq \lambda_0 \gamma$, then*

$$\alpha_1^*(c) = \begin{cases} \dfrac{c_0(\gamma) - \lambda_0 \gamma}{1 + \lambda_1 \gamma}, & \forall c \in [c_L, c_0(\gamma)], \\ 0, & \forall c \in (c_0(\gamma), c_H], \end{cases} \quad \tau^*(c) = \begin{cases} 1, & \forall c \in [c_L, c_0(\gamma)], \\ 0, & \forall c \in (c_0(\gamma), c_H]; \end{cases}$$

4. *if $\gamma \in [0, \min\{c_H/\lambda_0, \check{\gamma}\})$ and $c_0(\gamma) < \lambda_0 \gamma$, then*

$$\alpha_1^*(c) = 0, \ \forall c \in [c_L, c_H], \quad \tau^*(c) = \begin{cases} 1, & \forall c \in [c_L, \lambda_0 \gamma], \\ 0, & \forall c \in (\lambda_0 \gamma, c_H], \end{cases}$$

where $\check{\gamma} = v_L/(1 + v_L + \lambda_1)$ and $c_0(\gamma)$ is defined in (19).

Proof. See Appendix D. ■ □

The optimal mechanism is straightforward to implement. Fixing the known probability of efficacy γ, the regulator offers market share $\alpha_1^*(c)$ as a function of the privately known cost c. A firm has the incentive to take this offer and conduct the test ($\tau^*(c) = 1$) if and only if its cost is below a threshold, which is c_H, c_H, $c_0(\gamma)$, or $\lambda_0 \gamma$, for the four cases, respectively, according to Proposition 1. This is why the optimal mechanism can be interpreted and implemented as a take-it-or-leave-it offer.

The solid line in Figure 3 illustrates the optimal mechanism using model parameters from the calibration section C in the appendix. The dashed line illustrates how the optimal mechanism changes when the drug's potential harm (v_L) is doubled.

Drugs that are likely to be effective are tested. As the probability of efficacy decreases and testing costs increase, the likelihood of testing diminishes (Figure 3(a)).

The scale of conditional approval offered by the regulator can be divided into four regions (Figure 3(b)). First, the regulator grants full conditional approval for drugs likely to be efficacious ($\gamma \geq \check{\gamma}$).

Second, for drugs that are unlikely to be efficacious, but just barely ($\gamma \in (\gamma_2, \check{\gamma})$), the regulator grants partial conditional approval sufficient to motivate testing. Surprisingly, in

this region, there is an inverse relationship between efficacy probability and the scale of conditional approval (Figure 3(b)). One might assume that drugs just missing the threshold for full conditional approval should be granted more conditional approval than those with a lower efficacy probability. However, this is not the case. While the regulator values high-probability drugs more, so does the firm, meaning less conditional approval is needed to motivate testing. While efficient, giving higher-scale approval to drugs with lower efficacy probability might seem unfair. However, the drug's efficacy probability is determined by chance rather than effort. A higher approval share helps offset the bad luck of a lower draw on efficacy probability.

Third, drugs unlikely to be efficacious ($\gamma \in (\gamma_1, \gamma_2)$) receive only as much partial conditional approval as the regulator is willing to provide. The share is only sufficient for firms with relatively low testing costs.

Fourth, for drugs very unlikely to be efficacious ($\gamma \in (0, \gamma_1)$), the regulator grants no conditional approval. Even with low testing costs, the firm will not test these drugs, as they are unlikely to pass the test.

The values in Figures 3 are based on our estimates for the cancer drug avelumab. We assume testing costs were \$98 million to \$631 million, annual sales were \$300 million in the fifth year from approval, the annual social value in the good state was \$108 million, and in the bad state was a loss of \$290 million.

The threshold for conditional approval depends on the potential harm from a drug, v_L, which includes wasted resources and adverse effects from a toxic drug. When the harm from a bad drug increases, fewer drugs are tested and the conditional approval share decreases (comparing the solid and dashed lines in Figure 3).

Based on the solid line in Figure 3, to be a net positive for the regulator and receive 100% conditional approval share, the efficacy probability must be at least 58%. To receive half of the approval share, the efficacy probability must be at least 27%. If offered half of the approval share, only firms with sufficiently low costs will choose to conduct testing. According to the dashed line in Figure 3, on the other hand, the efficacy probability threshold for a drug to be a net positive for the regulator increases to 73%. Additionally, to secure half of the market share in conditional approval, the drug would need an efficacy probability of at least 47%. For more details, see Appendix C.

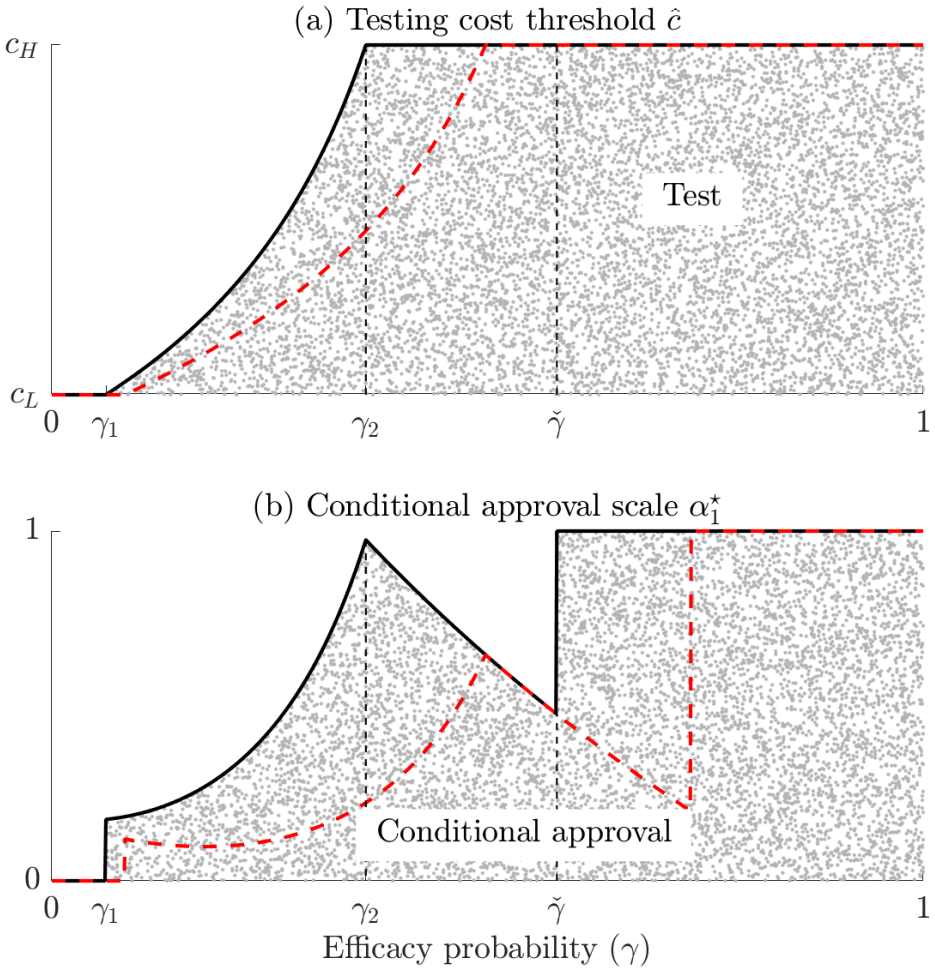

Figure 3: A drug with a high probability of efficacy is more likely to undergo testing (top panel) and typically receives a larger share of conditional approval, though the relationship is not strictly monotonic (bottom panel). When the potential harm from a drug is doubled (dashed line), fewer high-cost drugs are tested, and the regulator may grant reduced share of conditional approval. Source: Authors' analysis based on a representative drug with parameters shown in Table 3 and a uniform distribution of private information c.

3.1 Partial versus integer conditional approval

Up to this point, we assumed that the regulator could grant *partial* conditional approval. Now, we consider the optimal policy when conditional approval is either full or none, which we refer to as integer conditional approval.

We identify the conditions under which: (1) a firm tests because testing is profitable, (2) a firm tests because it is granted conditional approval, and (3) a firm does not test.

We start with the conditions in which the firm will test because the drug is profitable. Here, $a_1(c) = 0$ for any $c \in [c_L, c_H]$.

Proposition 2. *Let $\mathcal{U}_0$ denote the regulator's expected payoff when $a_1(c) = 0$ for any $c \in [c_L, c_H]$, i.e.,*

$$\mathcal{U}_0 \equiv \max_{\tau, a_1:(\text{IC}),(\text{IR}),(2),(3),a_1=0} \tilde{U}_R(\tau, a_1). \tag{20}$$

We have

$$\mathcal{U}_0 = (\lambda_0 \gamma v_H + z) F(\lambda_0 \gamma \pi).$$

Proof of Proposition 2. See Appendix D. ■ □

According to Proposition 2, we can construct the optimal regulatory policy when the scale of conditional approval is restricted to 0 or 1. The formal result is presented in the following proposition.

Proposition 3. *Let $\mathcal{U}_1$ denote the regulator's expected payoff when $a_1(c) = 1$ for any $c \in [c_L, c_H]$, i.e.,*

$$\mathcal{U}_1 \equiv \max_{\tau, a_1:(\text{IC}),(\text{IR}),(2),(3),a_1=1} \tilde{U}_R(\tau, a_1). \tag{21}$$

We have

$$\mathcal{U}_1 = [v_E + \gamma v_H (\lambda_0 + \lambda_1) + z] F(\pi q + \gamma \pi (\lambda_0 + \lambda_1)).$$

When the scale of conditional approval is restricted to 0 or 1, if $\mathcal{U}_0 \geq \mathcal{U}_1$, the optimal mechanism is to set $a_1(c) = 0$ for any $c \in [c_L, c_H]$; otherwise, it is optimal to set $a_1(c) = 1$ for any $c \in [c_L, c_H]$.

Proof of Proposition 3. See Appendix D. ■ □

We illustrate the three testing regions as a function of efficacy probability and testing cost in Figure 4(a). In the first region, where efficacy is high and costs are low, the firm will test because it is profitable, even without conditional approval. A dotted line marks the border of this region.

At the other extreme, in the third region, the firm will not test because the drug has a low probability of efficacy and high testing costs. A solid line indicates the border of this region. In between these extremes are drugs that will only be tested if the regulator grants conditional approval. Drugs in this region are only profitable if the regulator grants conditional approval.

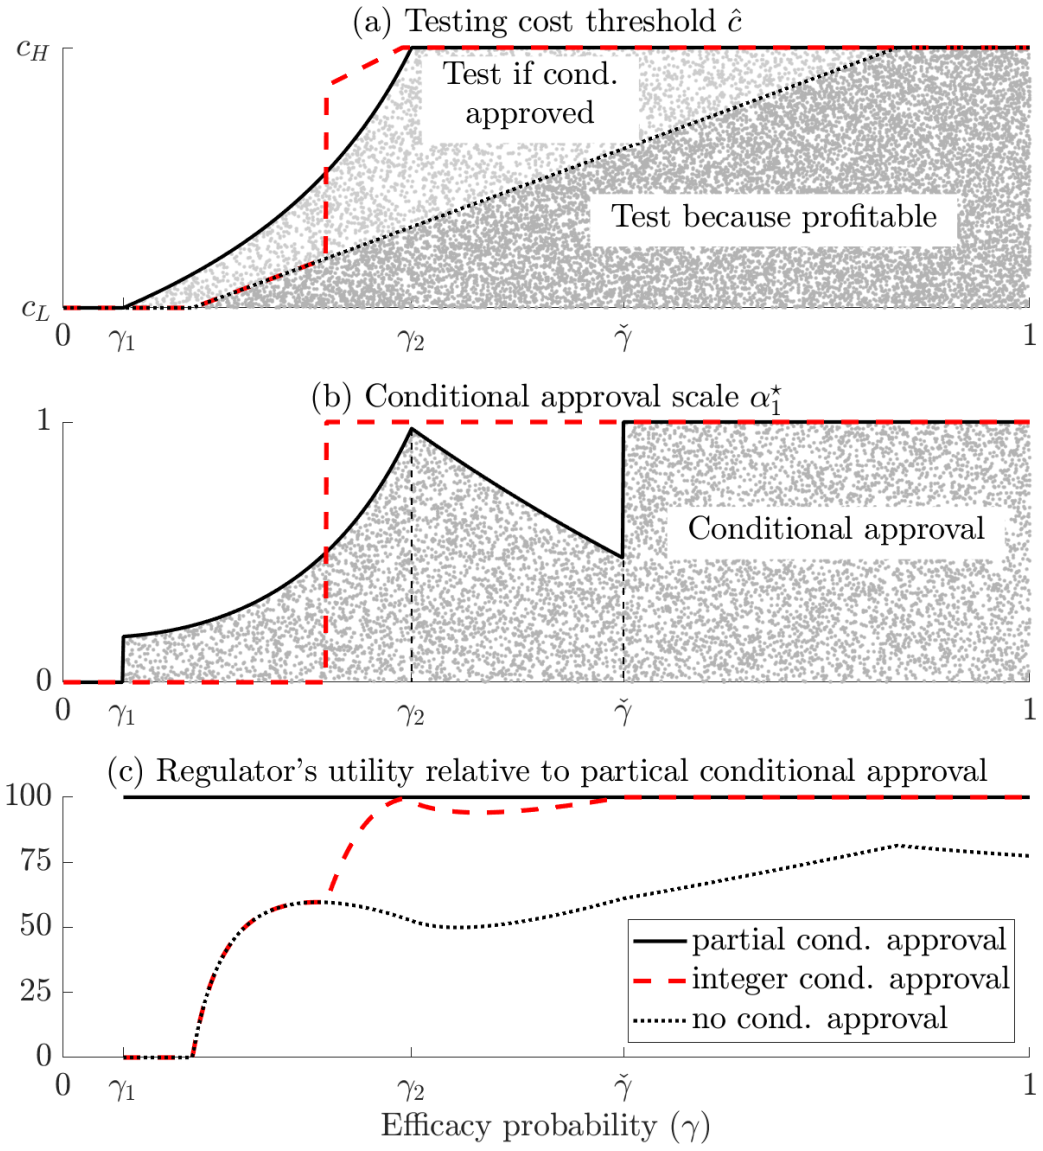

Figure 4: Thresholds for testing and the scale of conditional approval under three scenarios: (i) partial conditional approval is allowed; (ii) conditional approval is restricted to scales 0 or 1; and (iii) conditional approval is unavailable. Source: Authors' analysis based on a representative drug with parameters shown in Table 3 and a uniform distribution of private information c.

We also illustrate integer conditional approval with a dashed line in Figure 4. According to panel (b), the regulator should offer conditional approval if and only if the probability of

being efficacious is higher than a threshold, which is intuitive. On the left of this threshold, restricting to integer rather than partial conditional approval results in less testing (Figure 4(a)) because no conditional approval (Figure 4(b)) implies less profit to compensate the testing cost. Conversely, on the right of the threshold, full conditional approval (Figure 4(b)) implies more funding to cover testing costs, which allows more drugs to be tested (Figure 4(a)).

The regulator's utility under partial conditional approval is greater than or equal to its utility under integer conditional approval for drugs unlikely to be efficacious (Figure 4(c)). For drugs likely to be efficacious, the regulator's utility is higher under both partial and integer conditional approval.

3.2 Monopoly duration

The typical innovative (not generic) drug will have some monopoly-pricing power after approval but will lose that monopoly-pricing power when generic drugs arrive or when price controls begin. In the United States, price controls can begin as early as nine years after conditional approval.

A monopoly pricing period can have fixed duration or fixed horizon. Fixed duration means that earlier approval leads to an earlier end to monopoly pricing (stylized fact 9). Recall that Figure 1 illustrates the monopoly-pricing periods as functions of conditional approval and whether there is a fixed horizon or a fixed duration.

Under fixed duration, some of the firm's benefits of conditional approval are lost because earlier approval causes an earlier end to monopoly pricing. Hence, we will show that the regulator must offer a larger share of conditional approval if the firm's monopoly-pricing period has a fixed duration.

When changing from fixed horizon to fixed duration, some of the sales fall out of the monopoly-pricing period. The amount of doses lost from the monopoly-pricing period is $\tilde{\Delta} \in [0, \lambda_0]$. Therefore, the total discounted doses in the monopoly-pricing period after final approval is $\lambda_0 + \lambda_1 a_1 - \tilde{\Delta}\, \mathbb{1}_{a_1>0}$.[6]

In this scenario, the firm's expected payoff when its testing cost is c and it reports c', defined in (1), becomes

$$\tilde{\Pi}(c, c') \equiv \pi \cdot q \cdot \alpha_1(c') - c \cdot \tau(c') + \gamma \cdot \pi \cdot \left[\lambda_0 \cdot \tau(c') + \lambda_1 \cdot \alpha_1(c') - \tilde{\Delta} \cdot \mathbb{1}_{\alpha_1(c')>0} \cdot \tau(c')\right], \quad (22)$$

[6]We assume that lower prices generate savings for the regulator, but do not increase doses, because an insurer is paying most of a drug's price (Lakdawalla and Philipson, 2012).

and the regulator's expected payoff (6) becomes

$$\tilde{U}_R(\tau,\alpha_1) \equiv \int_{c_L}^{c_H} \Big\{ v_E \cdot q \cdot \alpha_1(c) + \gamma \cdot \Big[\lambda_0 + \lambda_1 \cdot a_1(c) - \tilde{\Delta} \cdot \mathbb{1}_{\alpha_1(c')>0}\Big] \cdot v_H \cdot \tau(c) \\ + \gamma \cdot \Big[\Delta + \tilde{\Delta} \cdot \mathbb{1}_{\alpha_1(c')>0}\Big] \cdot (v_H + \sigma) \cdot \tau(c) \Big\} \, \mathrm{d}F(c), \tag{23}$$

where the indicator function captures the existence of earlier price control if conditional approval is granted.

Under fixed duration, even a small scale of conditional approval can accelerate the end to monopoly pricing, so the firm might prefer to exit the market. Hence, the scale of partial conditional approval must increase to motivate testing.

The incentive compatibility constraint can be formulated as follows:

$$\forall (c, c') \in [c_L, c_H]^2 : \quad \tilde{\Pi}(c,c) \geq \tilde{\Pi}(c,c'), \tag{24}$$

guaranteeing that it is weakly dominant to report truthfully. The individual rationality constraints capture the idea that the firm must receive non-negative expected profit and must not be worse off than if it only applied for normal approval. The two constraints can be formulated as

$$\forall c \in [c_L, c_H] : \quad \tilde{\Pi}(c,c) \geq 0 \tag{25}$$

and

$$\forall c \in [c_L, c_H] : \quad \tilde{\Pi}(c,c) \geq -c + \lambda_0 \cdot \gamma \cdot \pi, \tag{26}$$

respectively. The right-side of the second inequality constraint represents the firm's expected profit when it conducts the testing and applies for normal approval.

The regulator's problem boils down to

$$\tilde{\mathcal{U}}_R \equiv \max_{\tau,\alpha_1:(2),(9),(24),(25),(26)} \tilde{U}_R(\tau,\alpha_1). \tag{27}$$

Because of the indicator function in both parties' payoff functions $\tilde{\Pi}$ and $\tilde{U}_R$, the regulator's problem is nonlinear. We propose an alternative way to solve the problem in two steps.

1. We solve the following two linear optimization problems separately:

   (a) we set $\alpha_1(c) = 0$ for any $c \in [c_L, c_H]$, and the problem is reduced to the one without conditional approval. In this case, firms with costs less than $\lambda_0 \cdot \gamma \cdot \pi$ will

conduct the testing; and then,

(b) we keep α_1 as a decision variable, replace all the indicator functions with value 1, and solve the optimization problem with constraint (26), i.e., capturing the idea that the firm should be weakly better off by applying for conditional approval.

2. We compare the two objective values in the previous step, and if the first one is greater, the optimal policy is to not grant conditional approval in period one. Conversely, if the second one is greater, the optimal mechanism is to grant conditional approval.

The critical scale of conditional approval is the level at which the firm is indifferent between applying or not, i.e.,

$$\pi \cdot q \cdot \hat{\alpha}_1 + \gamma \cdot \pi \cdot (\lambda_0 + \lambda_1 \cdot \hat{\alpha}_1 - \tilde{\Delta}) - c = \lambda_0 \cdot \gamma \cdot \pi - c \quad \Rightarrow \quad \hat{\alpha}_1 = \frac{\gamma \cdot \tilde{\Delta}}{q + \lambda_1 \cdot \gamma}.$$

With fixed duration, to induce the firm's application for conditional approval, the scale must be at least $\hat{\alpha}_1$. Therefore, when the optimal scale of conditional approval in the fixed horizon setting is smaller than $\hat{\alpha}_1$, the regulator faces a trade-off between granting $\hat{\alpha}_1$ or not granting at all. As γ increases, under the calibrated model parameters, the regulator consistently grants a larger scale of conditional approval (Figure 5(b)) to incentivize testing. Due to the higher net present value after the monopolistic period, the regulator's expected payoff is higher in most cases.

Figure 5 displays the optimal testing threshold and scale of conditional approval with fixed horizon and fixed duration under the calibrated model parameters (appendix section C). Drugs very unlikely to be efficacious ($\gamma < \gamma_1$) will not be tested and will not receive conditional approval. Drugs very likely to be efficacious ($\gamma > \check{\gamma}$) will receive full conditional approval. In between, the amount of conditional approval depends on whether the drug has a fixed duration or fixed horizon. The share of conditional approval will be higher for a drug with a fixed duration than a drug with a fixed horizon, in order to induce testing.

In summary, we have demonstrated that (in most cases) the regulator should offer a larger share of conditional approval if the firm's monopoly-pricing period has a fixed duration, in order to compensate the firm for an earlier pricing reduction.

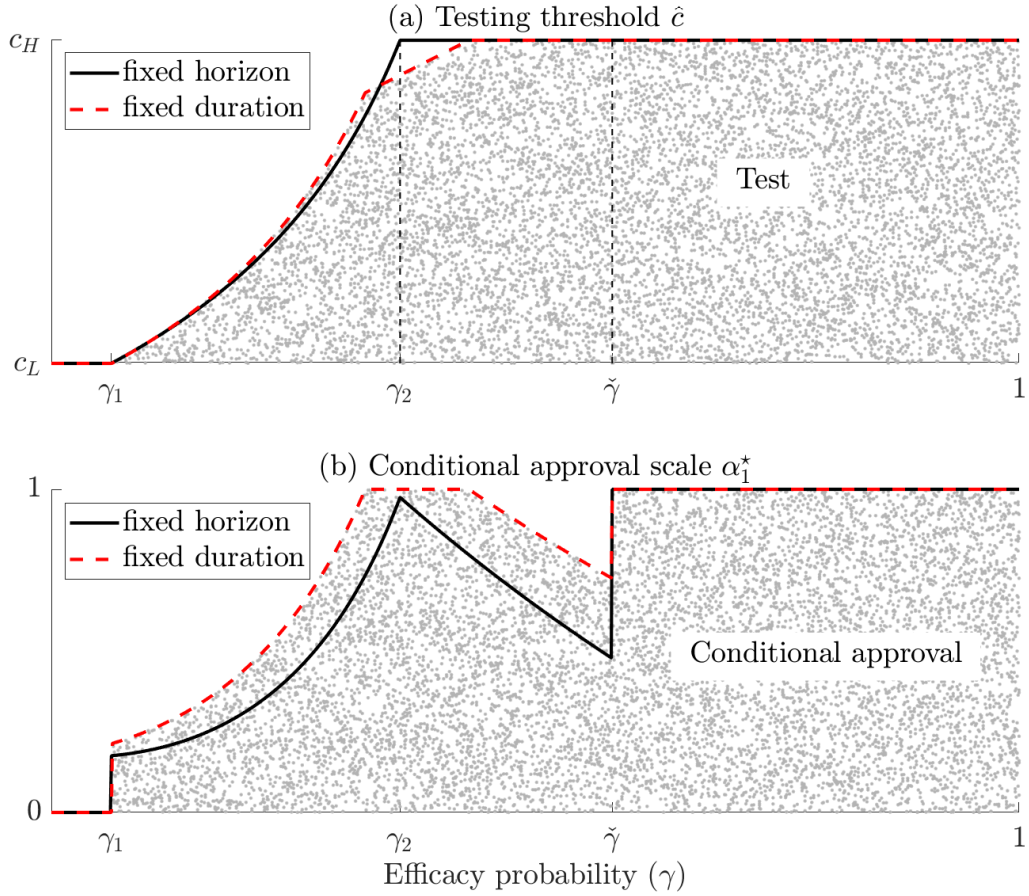

Figure 5: Optimal testing threshold and conditional approval scale depend on whether the monopoly-pricing period has a fixed duration or fixed horizon. Source: Authors' analysis based on a representative drug with parameters shown in Table 3 and a uniform distribution of private information c.

4 Continuous drug efficacy levels

Up to this point, we modeled drug efficacy as binary (good or bad) rather than as a continuum. Here, we allow efficacy to vary continuously, meaning a drug can narrowly miss the efficacy threshold.[7]

[7]In the continuous framework, we only report partial conditional approval for drugs with a fixed horizon, instead of a fixed duration. We do not report on fixed duration because the range of partial conditional approval options and strategic final approvals would be too extensive. However, the same intuition from the binary case can apply to fixed duration. Specifically, under fixed duration, the firm is less inclined to accept partial conditional approval, so the regulator must offer a larger share of partial conditional approval on the intensive margin.

By allowing continuous drug efficacy, we reveal an additional tool the regulator can use to incentivize the firm to test the drug. We will show that it is sometimes optimal for the regulator to commit to granting final approval of a drug that misses the threshold for efficacy.

In practice, missing the threshold might mean that the trials were not conclusive in demonstrating a benefit for the average patient, but the regulator asks for no more. The regulator can commit to a lower threshold during discussions with the firm before confirmatory testing. Indeed, the FDA grants what it calls "breakthrough status" to certain drugs, providing the developer with timely advice on clinical trial design (stylized fact 14).

We generalize the model from section 3 by allowing the drug's efficacy, η, to take values in $[0, 1]$ and the regulator's payoff, v, to be continuous in η, with $v(0) < 0 < v(1)$.

Both the regulator and the firm believe that η is drawn from a distribution with cumulative distribution function $G(\eta)$ and density function $g(\eta)$. The cdf G is strictly increasing, hence the density g is strictly positive on $[0, 1]$.

In this setting, the regulatory policy specifies one more decision in period two: a probability of final approval $a_2(c, \eta)$, for each cost value $c \in [c_L, c_H]$ and efficacy level $\eta \in [0, 1]$. The expected value of v is given by

$$v_E \equiv \int_0^1 v(\eta)\,\mathrm{d}G(\eta).$$

With a slight abuse of notation, we still use $\hat{\Pi}$ and $\tilde{U}_R$ to denote the firm and the regulator's payoff functions. If the firm's true cost is c, and it reports c' to the regulator, its expected total profit is

$$\hat{\Pi}(c, c') \equiv \tau(c') \cdot \left[a_1(c') - c + [\lambda_0 + \lambda_1 \cdot a_1(c')] \cdot \int_0^1 a_2(c', \eta)\,\mathrm{d}G(\eta)\right]. \tag{28}$$

The probability a_2 must be between 0 and 1:

$$\forall (c, \eta) \in [c_L, c_H] \times [0, 1]: \qquad 0 \leq a_2(c, \eta) \leq 1. \tag{29}$$

The regulator's expected payoff consists of two parts, corresponding to the monopoly-pricing and competitive periods. During the monopoly-pricing period, similar to equation (6), the regulator's expected payoff is

$$\int_{c_L}^{c_H} \tau(c) \cdot \left[v_E \cdot a_1(c) + [\lambda_0 + \lambda_1 \cdot a_1(c)] \cdot \int_0^1 v(\eta) \cdot a_2(c, \eta)\,\mathrm{d}G(\eta)\right] \mathrm{d}F(c).$$

Recall that σ represents the price drop during the competitive period. In this scenario, the regulator's marginal benefit from a drug with efficacy η increases from $v(\eta)$ to $v(\eta)+\sigma$. The regulator's expected payoff in the competitive period boils down to

$$\Delta \cdot \int_{c_L}^{c_H} \int_0^1 \tau(c) \cdot [v(\eta)+\sigma] \cdot a_2(c,\eta)\ \mathrm{d}G(\eta)\ \mathrm{d}F(c).$$

Therefore, the regulator's payoff function is given as follows:

$$\begin{aligned}\tilde{U}_R(\tau, a_1, a_2) \equiv \int_{c_L}^{c_H} \tau(c) \Bigg[& v_E\, a_1(c) + [\lambda_0 + \lambda_1\, a_1(c)] \int_0^1 v(\eta)\, a_2(c,\eta)\ \mathrm{d}G(\eta) \\ & + \Delta \int_0^1 [v(\eta)+\sigma]\ a_2(c,\eta)\ \mathrm{d}G(\eta) \Bigg] \mathrm{d}F(c). \end{aligned} \tag{30}$$

The regulator's problem is

$$\begin{aligned} U_R^* \equiv & \max_{\tau, a_1, a_2}\ \tilde{U}_R(\tau, a_1, a_2), \\ & \text{subject to}\ \ (\mathrm{IC}), (\mathrm{IR}), (2), (3),\ \text{and}\ (29). \end{aligned} \tag{31}$$

The regulator's problem is nonlinear, so we focus on a relaxed linear optimization problem, whose objective function value serves as an upper bound for the original problem. Besides α_1 defined in (8), we introduce another two decision variables, i.e., for any $(c,\eta) \in [c_L, c_H] \times [0,1]$,

$$\alpha_2(c,\eta) \equiv \tau(c)\ [\lambda_0 + \lambda_1\, a_1(c)]\ a_2(c,\eta) = [\lambda_0\, \tau(c) + \lambda_1\, \alpha_1(c)]\ a_2(c,\eta),$$

and

$$\alpha_3(c,\eta) \equiv \tau(c)\, a_2(c,\eta).$$

Both parties' payoff functions defined in (28) and (30) can be rewritten as

$$\hat{\Pi}(c,c') = \alpha_1(c') - c\,\tau(c') + \int_0^1 \alpha_2(c',\eta)\, \mathrm{d}G(\eta), \tag{32}$$

and

$$\hat{U}_R(\tau, \alpha_1, \alpha_2, \alpha_3) = \int_{c_L}^{c_H} \Bigg[v_E\, \alpha_1(c) + \int_0^1 v(\eta)\, \alpha_2(c,\eta)\, \mathrm{d}G(\eta)$$

$$+\Delta \int_0^1 [v(\eta)+\sigma]\ \alpha_3(c,\eta)\,\mathrm{d}G(\eta)\Big]\,\mathrm{d}F(c). \tag{33}$$

We impose the following feasibility conditions on the newly defined decision variables α_2 and α_3, i.e., for any $(c,\eta)\in[c_L,c_H]\times[0,1]$,

$$0\le\alpha_2(c,\eta)\le\lambda_0\,\tau(c)+\lambda_1\,\alpha_1(c), \tag{34}$$

$$\lambda_0\,\alpha_3(c,\eta)\le\alpha_2(c,\eta)\le(\lambda_0+\lambda_1)\,\alpha_3(c,\eta), \tag{35}$$

$$\alpha_3(c,\eta)\le\tau(c). \tag{36}$$

The relaxed regulator's problem becomes

$$\begin{aligned}\bar{U}_R^*\equiv\ &\max_{\tau,\alpha_1,\alpha_2,\alpha_3}\ \hat{U}_R(\tau,\alpha_1,\alpha_2,\alpha_3), \\ &\text{subject to}\ \ (\mathrm{IC}),(\mathrm{IR}),(2),(9),(34),(35),\ \text{and}\ (36),\end{aligned} \tag{37}$$

where $\hat{\Pi}$ and $\hat{U}_R$ follow (32) and (33), respectively.

The next proposition formalizes that the objective function value of the optimization problem (37) is an upper bound for (31).

Proposition 4. $\bar{U}_R^*\ge U_R^*$.

Proof of Proposition 4. See Appendix D. ■ □

The following proposition states that, given the optimal solution to the relaxed problem (37), we can always construct a feasible solution to the original problem (31), with its objective function value serving as a lower bound.

Proposition 5. *Let $\hat{\mathcal{M}}\equiv\{\hat{\tau},\hat{\alpha}_1,\hat{\alpha}_2,\hat{\alpha}_3\}$ denote the optimal solution to the relaxed problem (37). The following mechanism, $\mathcal{M}\equiv\{\hat{\tau},\hat{\alpha}_1,\tilde{\alpha}_2\}$, is feasible to the original problem (31). The decision variable $\tilde{\alpha}_2$ is constructed as follows: for any $(c,\eta)\in[c_L,c_H]\times[0,1]$,*

$$\tilde{\alpha}_2(c,\eta)=\begin{cases}\lambda_0\,\hat{\tau}(c)+\lambda_1\,\hat{\alpha}_1(c), & \forall\eta\in[\check{\eta}(c),1],\\ 0, & \forall\eta\in[0,\check{\eta}(c)),\end{cases}$$

where the threshold $\check{\eta}(c)$ is uniquely determined by

$$\int_0^1\hat{a}_2(c,\eta)\,\mathrm{d}G(\eta)=1-G(\check{\eta}(c)),$$

with

$$\hat{a}_2(c,\eta) = \begin{cases} \frac{\hat{\alpha}_2(c,\eta)}{\lambda_0\,\hat{\tau}(c)+\lambda_1\,\hat{\alpha}_1(c)}, & \textit{if } \max\{\hat{\tau}(c),\hat{\alpha}_1(c)\} > 0, \\ 0, & \textit{otherwise.} \end{cases}$$

Therefore, we have

$$\tilde{U}_R^* \equiv \tilde{U}_R(\hat{\tau},\hat{\alpha}_1,\tilde{\alpha}_2) \leq U_R^*.$$

Proof of Proposition 5. See Appendix D. ■ □

We use the calibrated model parameters and solve the relaxed problem numerically for various G, based on which we construct feasible solutions to the original problem. The candidate solution constructed from Proposition 5 performs well, consistently achieving 97.5% of the upper bound. See Appendix E for more details.

Our mechanism's testing and final approval decisions are straightforward—request testing if the cost c is low (below the threshold $\hat{c}$), and grant final approval if efficacy η is high (above the threshold $\hat{\eta}$). Against the corresponding expected values v_E, we plot the testing threshold $\hat{c}$ (Figure 6(a)), the scale of conditional approval $\alpha_1^\star$ (Figure 6(b)), and the threshold of final approval $\hat{\eta}$ (Figure 6(c)).

We partition the v_E space into five intervals under the calibrated model parameters. First, in the most pessimistic scenario ($v_E < v_1$), the regulator grants no conditional approval, and the firm does not test. Second, if beliefs are only slightly pessimistic ($v_E \in [v_1, v_2)$), the regulator still does not grant conditional approval, but the firm will test if its testing costs are very low. Third, as the expected value increases ($v_E \in [v_2, v_3)$), the regulator grants conditional approval, and the firm tests drugs with a higher threshold for testing costs. Fourth, as the expected value rises further ($v_E \in [v_3, v_4)$), the regulator can reduce the conditional approval granted, knowing the firm will still test. Here, the regulator grants just enough conditional approval to incentivize the firm with the maximum testing cost (c_H) to test. Fifth, if beliefs are optimistic ($v_E \in [v_4, v_H]$), the regulator grants full conditional approval, and the firm tests.

The regulator's final approval threshold ($\hat{\eta}$) is lower for drugs with lower expected value (v_E) under the parameters in our calibration exercise (Figure 6(c)). At v_4, the final approval threshold jumps up because the scale of conditional approval jumps up. With higher conditional approval, the regulator can set a higher final approval threshold without discouraging the firm from testing.

We compare policy performance with and without the tools of conditional approval and strategic final approval (Figure 7). First, consider the scenario where the regulator uses conditional approval but not strategic final approval. Here, it fixes a naive final approval

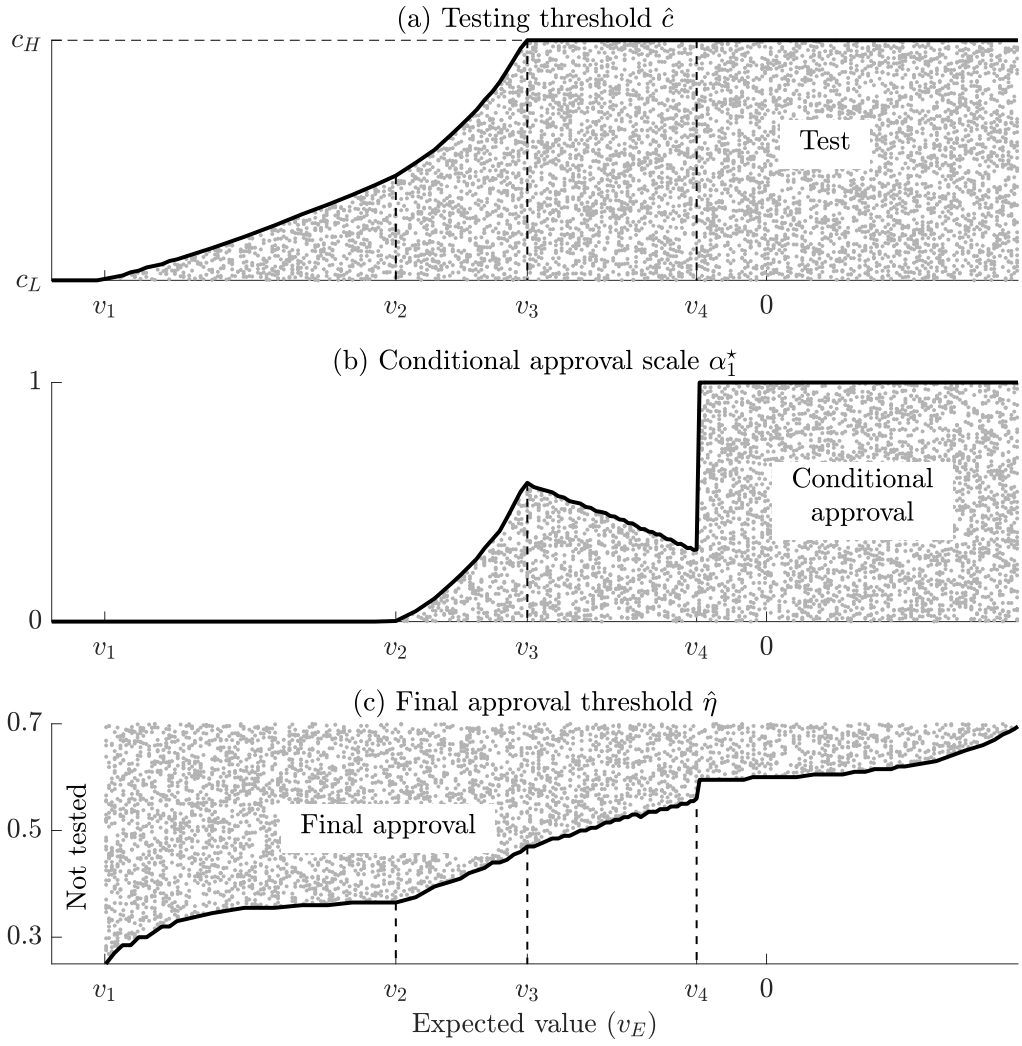

Figure 6: Drugs with low expected value and high testing cost are not tested (top panel). Drugs with expected values that are not too negative receive partial conditional approval (middle panel). The regulator's final approval threshold is lower for drugs with lower expected value (bottom panel). Source: Authors' analysis based on a representative drug with parameters shown in Table 3, a uniform distribution of private information c, efficacy distributed according to $G(\eta) = \eta^k,\ \forall \eta \in [0,1]$, and linear regulator's utility function v.

threshold, $\bar{\eta}$, defined as the point where the weighted average of the regulator's expected payoff is zero:

$$\bar{\eta} \equiv \inf \left\{ \eta \in [0,1] : \lambda_0 \, v(\eta) + \Delta \, [v(\eta) + \sigma] \geq 0 \right\}.$$

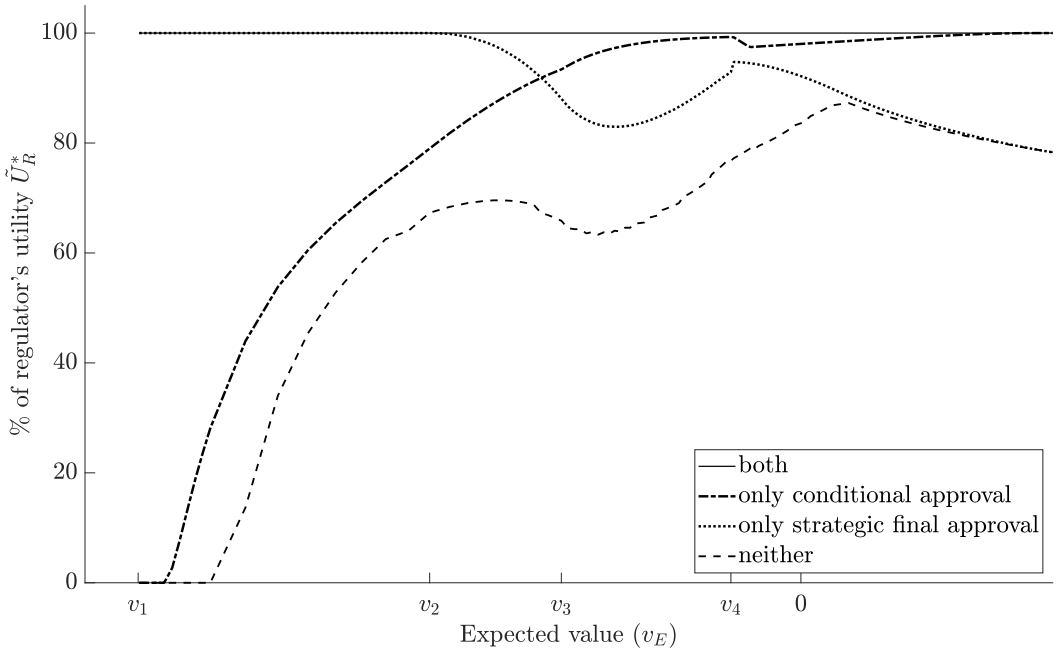

Figure 7: The regulator's utility is highest when it uses both conditional approval and a strategic final approval threshold, and lowest when it uses neither tool. Source: Authors' analysis based on a representative drug with parameters shown in Table 3 and uniform distribution of private information c.

After observing the efficacy level, the regulator makes no further decisions, effectively reducing the problem to the binary framework.

When the regulator uses only a strategic final approval threshold and not conditional approval, we have $a_1(c) = 0$ for all $c \in [c_L, c_H]$. With neither tool, the final approval threshold is $\bar{\eta}$, and only firms with low testing costs proceed with trials.

The regulator's utility is highest when it uses both conditional approval and a strategic final approval threshold. Both tools are needed for inducing a firm to test when beliefs are very pessimistic, particularly when testing costs are high.

5 Conclusion

Our analysis leads to five recommendations for drug regulators. First, regulators should not be discouraged by critics who oppose conditional approval of indications that have a low efficacy probability, as long as the drugs are likely to be safe. Second, regulators should

lower the bar for conditional approval, but the conditional approval share should be limited. Practically speaking, the share can be limited through a narrower list of approved indications or insurers only covering certain patients. Third, regulators should not, however, limit conditional approval share as much for indications with fixed-length monopolies, because those will have earlier price drops. For example, a drug given conditional approval might be subject to earlier Medicare price controls, so it should be given a higher conditional approval share. Fourth, regulators should commit to lowering the bar for *final* approval (in addition to conditional approval) to encourage more testing. Fifth, regulators should not automatically grant more conditional approvals for indications that narrowly miss the efficacy probability threshold, as these indications are likely to be tested regardless.

Our fifth recommendation might seem counterintuitive. A regulator should sometimes grant a higher conditional approval share for an indication with a lower efficacy probability. An indication that narrowly misses the threshold requires only a relatively small conditional approval share to incentivize testing, while indications below it need more. This approach is not only efficient but also arguably fair. A higher approval share helps offset the bad luck of a low draw on efficacy probability.

Conditional approval carries risks, such as wasting money and harming health, but the benefits could be substantial. Conditional approval accelerates access to treatments, leads to earlier price reductions, and incentivizes testing of more indications. Another potential benefit, which we do not model, is the opportunity to learn from real-world use. Conditional approval could generate valuable data, reducing the number of patients needed in clinical trials and lowering testing costs. However, due to selection bias, regulators typically do not rely on such evidence. For instance, if only high-income people use the drug, their outcomes could reflect advantages unrelated to the drug, creating a misleading correlation between the drug and positive outcomes. Typically, confirmatory evidence comes from randomized-controlled tests.

non-AER

References

Aitken, Murray, Michael Kleinrock, and Jamie Pritchett. 2024. "Global Use of Medicines: Outlook to 2028." Tech. rep., IQVIA Institute for Human Data Science. URL `www.iqvia institute.org`.

Armstrong, Annalee. 2024. "Preventing "3 billion visits" to FDA a year with CBER's Peter

Marks." *Fierce Biotech* URL https://www.fiercebiotech.com/biotech/preventing-3-billion-visits-fda-year-cbers-peter-marks. Accessed 2024-07-02.

Belluck, Pam and Rebecca Robbins. 2021. "Three FDA advisers resign over agency's approval of Alzheimer's drug." *The New York Times* URL https://www.nytimes.com/2021/06/10/health/aduhelm-fda-resign-alzheimers.html.

Berndt, Ernst R, Rachel Glennerster, Michael R Kremer, Jean Lee, Ruth Levine, Georg Weizsäcker, and Heidi Williams. 2007. "Advance market commitments for vaccines against neglected diseases: Estimating costs and effectiveness." *Health Economics* 16 (5):491–511.

Bognar, Katalin, John A. Romley, Jay P. Bae, James Murray, Jacquelyn W. Chou, and Darius N. Lakdawalla. 2017. "The role of imperfect surrogate endpoint information in drug approval and reimbursement decisions." *Journal of Health Economics* 51:1–12.

Budish, Eric, Benjamin N. Roin, and Heidi Williams. 2015. "Do firms underinvest in long-term research? Evidence from cancer clinical trials." *American Economic Review* 105 (7):2044–2085.

Carpenter, Daniel and Michael M. Ting. 2007. "Regulatory errors with endogenous agendas." *American Journal of Political Science* 51 (4):835–852.

Centers for Medicare & Medicaid Services. 2022. "CMS finalizes Medicare coverage policy for monoclonal antibodies directed against amyloid for the treatment of Alzheimer's disease." URL https://www.cms.gov/newsroom/press-releases/cms-finalizes-medicare-coverage-policy-monoclonal-antibodies-directed-against-amyloid-treatment.

———. 2024. "Medicare & Medicaid Spending by Drug." URL https://data.cms.gov/summary-statistics-on-use-and-payments/medicare-medicaid-spending-by-drug. Baltimore, MD. Accessed on December 2, 2024. URL: https://data.cms.gov/summary-statistics-on-use-and-payments/medicare-medicaid-spending-by-drug.

Chandra, Amitabh, Jennifer Kao, Kathleen L. Miller, and Ariel D. Stern. 2024. "Regulatory incentives for innovation: The FDA's breakthrough therapy designation." *The Review of Economics and Statistics* :1–46.

Collard-Wexler, Allan, Matthew Grennan, and Andrew Steck. 2024. "Regulating new product testing: The FDA vs. the invisible hand." *Working Paper* :1–32.

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

Dubois, Pierre, Paul-Henri Moisson, and Jean Tirole. 2022. "The economics of transferable patent extensions." *Toulouse School of Economics Working Paper* 1377:1–55.

Feuerstein, Adam. 2024. "FDA 's Peter Marks seems inclined to grant full approval to Sarepta's Duchenne gene therapy." URL `https://www.statnews.com/2024/01/09/fda-peter-marks-full-approval-sarepta-duchenne-gene-therapy/`.

Florko, Nicholas. 2021. "Ellis Unger, an outspoken FDA veteran, retires." URL `https://www.statnews.com/2021/08/12/ellis-unger-an`.

Frank, Richard G, Mahnum Shahzad, and Ezekiel J Emanuel. 2022. "Accelerated approval of cancer drugs: No economic reward for drug makers that conduct confirmatory trials." *Health Affairs* 41 (9):1273–1280.

Gans, Joshua S. and David B. Ridley. 2013. "Innovation incentives under transferable fast-track regulatory review." *Journal of Industrial Economics* 61 (3):789–816.

Grennan, Matthew and Robert J. Town. 2020. "Regulating innovation with uncertain quality: Information, risk, and access in medical devices." *American Economic Review* 110 (1):120–161.

Henry, Emeric, Marco Loseto, and Marco Ottaviani. 2022. "Regulation with experimentation: Ex ante approval, ex post withdrawal, and liability." *Management Science* 68 (7):5330–5347.

Henry, Emeric and Marco Ottaviani. 2019. "Research and the approval process: The organization of persuasion." *American Economic Review* 109 (3):911–955.

Kremer, Michael, Jonathan Levin, and Christopher M Snyder. 2022. "Designing advance market commitments for new vaccines." *Management Science* 68 (7):4786–4814.

Lakdawalla, Darius and Tomas Philipson. 2012. "Does intellectual property restrict output? An analysis of pharmaceutical markets." *The Journal of Law and Economics* 55 (1):151–187.

Light, Donald W. and Rebecca N. Warburton. 2005. "Extraordinary claims require extraordinary evidence." *Journal of Health Economics* 24 (5):1030–1033.

Liu, Ian T. T., Aaron S. Kesselheim, and Edward R Scheffer Cliff. 2024. "Clinical benefit and regulatory outcomes of cancer drugs receiving accelerated approval." *Journal of the American Medical Association* 331 (17):1471–1479.

McClellan, Andrew. 2022. "Experimentation and approval mechanisms." *Econometrica* 90 (5):2215–2247.

Myerson, Roger B. 1981. "Optimal auction design." *Mathematics of operations research* 6 (1):58–73.

Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov. 2020. "Persuading the principal to wait." *Journal of Political Economy* 128 (7):2542–2578.

Peltzman, Sam. 1973. "An evaluation of consumer protection legislation: The 1962 drug amendments." *Journal of Political Economy* 81 (5):1049–1091.

Philipson, Tomas J. and Eric Sun. 2008. "Is the Food and Drug Administration safe and effective?" *Journal of Economic Perspectives* 22 (1):85–102.

Ridley, David B. and Chung-Ying Lee. 2020. "Does Medicare reimbursement drive up drug launch prices?" *The Review of Economics and Statistics* 102 (5):980–993.

Ridley, David B. and Stephane A. Regnier. 2016. "The commercial market for priority review vouchers." *Health Affairs* 35 (5):776–783.

Robey, Seth and Frank S. David. 2016. "Drug launch curves in the modern era." *Nature Reviews Drug Discovery* 16 (1):13–14.

Snyder, Christopher M., Kendall Hoyt, and Dimitrios Gouglas. 2023. "An optimal mechanism to fund the development of vaccines against emerging epidemics." *Journal of Health Economics* 91:102795.

US Congress. 2023. "Consolidated Appropriations Act, 2023." URL `https://www.congress.gov/117/bills/hr2617/BILLS-117hr2617enr.pdf`. Pub. L. No. 117-328, div. FF, title III, §3210(d)(D), 136 Stat. 4459.

Wouters, Olivier J, Martin McKee, and Jeroen Luyten. 2020. "Estimated research and development investment needed to bring a new medicine to market, 2009-2018." *Jama* 323 (9):844–853.

Xu, Liang, Hui Zhao, and Nicholas C Petruzzi. 2021. "Inducing compliance with postmarket studies for drugs under FDA's accelerated approval pathway." *Manufacturing & Service Operations Management* 23 (1):170–190.

A Online appendix

This online appendix includes stylized facts, model calibration, and proofs of lemmas and propositions.

B Stylized facts

Stylized facts about drug regulation and competition provide a foundation for the model. The first stylized facts describe what is in the model, while the remaining stylized facts explain what is not included and the reasons for the exclusion.

Stylized fact 1: Phase II clinical testing provides information about a drug's effectiveness, but the number of patients enrolled is typically too small to yield statistically significant results. Phase I testing often includes tens of patients, phase II hundreds, and phase III thousands for diseases common enough to enroll that many patients. Additionally, phase III trials tend to be longer than phase II. Regulators and firms prefer shorter, smaller phase I and II trials to minimize financial waste and patient harm, given that most drugs entering phase II do not progress to phase III (DiMasi, Grabowski, and Hansen, 2016). We treat phase I and II testing costs as sunk, but the model could be reframed to account for earlier decision-making by the firm.

Stylized fact 2: The regulator can require confirmatory testing as a condition for granting conditional approval. Beginning in fiscal year 2023, the U.S. Congress granted the FDA the authority to mandate that trials begin before a drug receives conditional approval (US Congress, 2023). Previously, some firms selling conditionally-approved drugs cited difficulty enrolling enough patients in trials for rare diseases. This rationale was often a convenient excuse, as confirmatory testing could reveal that the drug should be withdrawn (Xu, Zhao, and Petruzzi, 2021; Frank, Shahzad, and Emanuel, 2022). Among 46 cancer drugs granted conditional approval between 2013 and 2017, 15% had not completed testing after a median of six years (Liu, Kesselheim, and Cliff, 2024) (Figure 8). However, we expect that under the new policy, the share of drugs failing to do timely testing will drop to near zero.

Stylized fact 3: The regulator can grant conditional approval for a portion of the population. The regulator can approve the drug for only certain patients, such as those with the most severe conditions. Alternatively, the regulator can approve the drug for everyone but only pay for it for some patients, leaving the remainder to pay out-of-pocket. For example, after the conditional approval of an Alzheimer's disease drug, Medicare announced it would only cover the drug for patients who agreed to have their data collected for a study

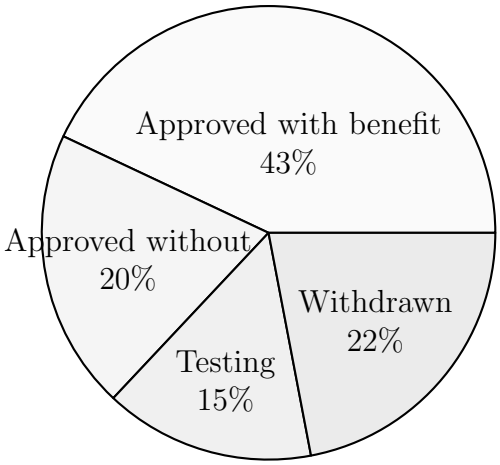

Figure 8: For 46 cancer indications conditionally approved in the U.S. between 2013 and 2017, 43% received final approval after demonstrating statistically significant benefits in confirmatory trials, 20% received final approval despite Liu, Kesselheim, and Cliff (2024) finding no evidence of statistically significant benefits, 22% were withdrawn, and 15% remained in ongoing testing. Source: Authors' figure based on data from Liu, Kesselheim, and Cliff (2024).

.

of the drug's efficacy (Centers for Medicare & Medicaid Services, 2022). For simplicity, we will assume that the regulator conditionally approves for a share of uniform patients.

Stylized fact 4: The firm must truthfully disclose the outcomes of its clinical trials. These results inform both the expected efficacy after phase II testing and the confirmed efficacy after phase III testing. Firms are required to submit their results to the regulator.

Stylized fact 5: Conditionally-approved drugs are typically safe but might be ineffective. Some drugs have safety risks which increase the harm from conditional approval. However, the health risks of conditional approval are generally small. Drugs given conditional approval have passed safety testing in phases I and II. Some drugs that are conditionally approved for one disease indication have already passed phase III testing for a different indication. Furthermore, conditionally-approved drugs treat serious diseases with no or few existing treatments. Hence, the potential harm from conditional approval is primarily financial, involving the waste of money on an ineffective drug. The loss from a bad drug, whether wasted money or side effects, is v_L.

Stylized fact 6: Testing costs are privately known by the firm. According to Light and Warburton (2005), "this science-based industry refuses to allow independent parties to check the validity of their cost data and analyze it so that policy can be based on solid, objective, reproducible evidence." Firms do not disclose testing costs, because it is hard to allocate

joint costs. Also, firm executives fear being held up. Disclosing testing costs could lead the government to reimburse only the testing costs, without compensating for the failures of other drugs or the opportunity cost of capital. For external parties, it is especially hard to benchmark costs for conditionally-approved drugs because they tend to be novel with few comparable examples.

Stylized fact 7: The regulator does not learn more about the drug's efficacy from its sales. Information from use in patients outside of controlled testing is less reliable due to selection bias. For example, if high-income people are generally healthy and have the money to pay for a drug, then the drug will look efficacious, due to selection. Hence, so-called "real-world evidence" is rarely used. If it were used, then it would reduce testing cost for the firm. Hence, our analysis might underestimate the benefits of conditional approval.

Stylized fact 8: Drugs given conditional approval often have several years of monopoly pricing. Market exclusivity lasts seven years in the U.S. and ten years in Europe for drugs for rare diseases. From 2014 to 2023, more than 80% of FDA conditional approvals were for rare diseases (see Figure 9). Additionally, markets for rare diseases are often too small to attract multiple firms.

Stylized fact 9: Patent expiration has a fixed date, but exclusivity and price control initiation dates have fixed durations from the approval date. We refer to drugs with long patent lives remaining as "fixed horizons" and drugs with monopoly-pricing periods tied to approval as "fixed duration."

The end of monopoly is the minimum of the drug's price control initiation date and the end of the longest-running exclusivity or patent.

A drug will have a fixed-duration monopoly if it is subject to price controls or if its remaining patent is short relative to its exclusivity.

A drug will be subject to Medicare price controls – and thus have a fixed-duration monopoly – if it is a top seller with multiple indications. For top-selling small-molecule drugs, Medicare price controls start after nine years. For top-selling large-molecule drugs, Medicare price controls start after twelve years. Drugs for rare diseases with only one indication are not subject to Medicare price controls.

A drug with a short patent life will have protection mainly provided by exclusivities which have a fixed duration. Exclusivities include orphan drug exclusivity of seven years and large-molecule exclusivity of twelve years from approval in the United States. A drug will have a shorter effective patent life if it has a short testing period, due to conditional approval, because of provisions in the Hatch-Waxman Act. The act restores half of a drug's

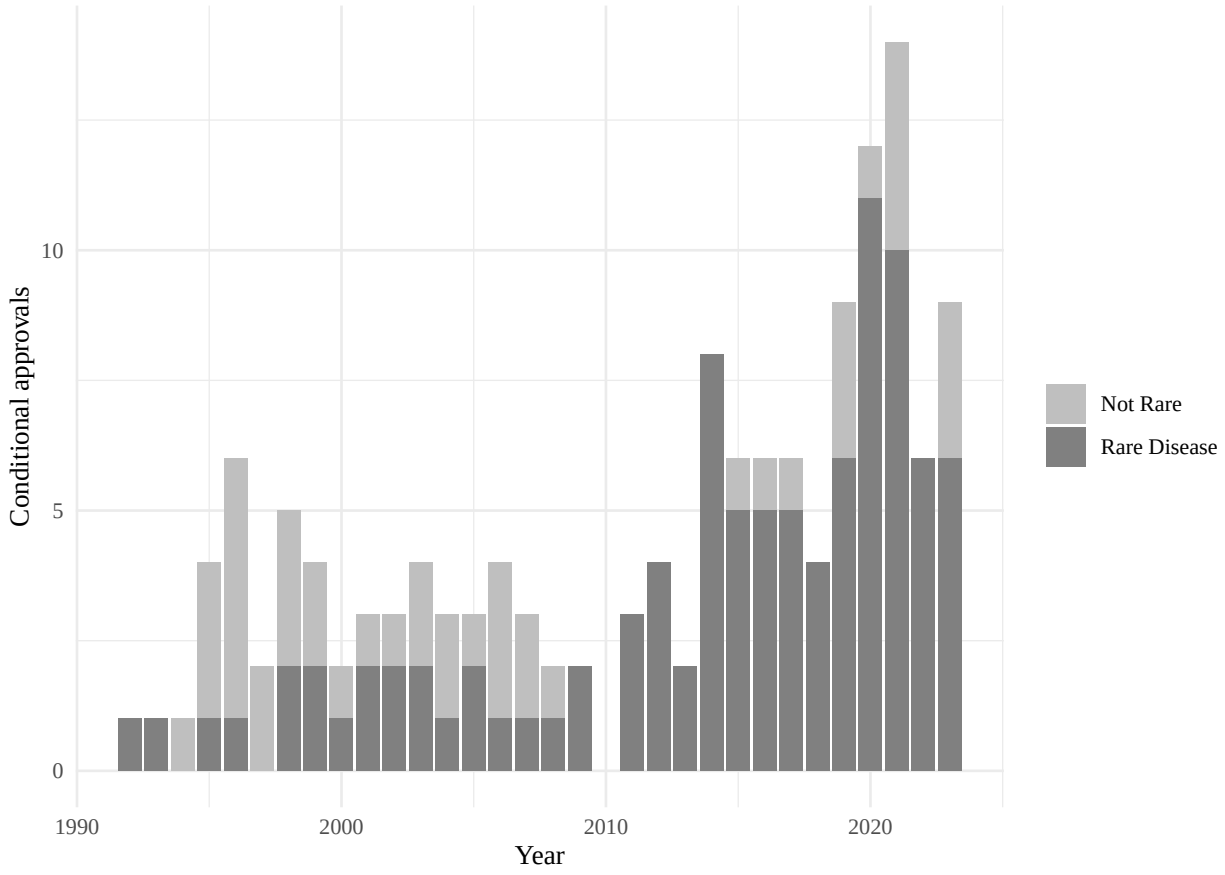

Figure 9: The FDA granted conditional approval to 80 drugs between 2014 and 2023. Most of the conditionally-approved drugs treat rare diseases. The official FDA term for conditional approval is "accelerated approval" and for rare disease is "orphan disease." Source: Authors' analysis of data on approvals of new molecules by the Center for Drug Evaluation and Research.

effective patent life lost during human testing. Shorter testing means shorter effective patent life.

We treat the monopoly-pricing period as exogenous rather than a policy option, because changing it would introduce complexity to the model and require legislative action from Congress and the President. Moreover, altering patent durations is also challenging due to international agreements.

Stylized fact 10: Conditional approval increases the firm's revenue during conditional approval and several years afterward. Drug revenue typically rises during the first five years after approval (Robey and David, 2016), as physicians and patients become more familiar with the drug. Demand is higher for drugs with which physicians and patients have more

experience (Ridley and Lee, 2020).

Stylized fact 11: Conditional approval provides earlier revenue for some firms that would otherwise not continue testing. According to the head of the FDA center responsible for gene therapies and vaccines, "The wherewithal to do a three-year study or a four-year study without having a revenue stream, it's just beyond many companies that are startups. So having the accelerated approval process is a way to get there" (Armstrong, 2024).

Stylized fact 12: A discrete-time model is appropriate in our setting because human clinical trials typically involve pre-specified endpoints and results that are announced all at once. Typically, phase III trials have pre-specified endpoints, followed by a comprehensive analysis by the firm before regulatory submission. However, there are exceptions. Occasionally, pre-planned points are set along the testing process to check for early safety signals that might necessitate halting the test. A regulator may also conduct a rolling review, as the FDA did with the COVID-19 vaccine. Finally, animal testing, which occurs before human testing, has a more flexible and exploratory set of protocols.

Stylized fact 13: The regulator cannot provide direct financing for late-stage testing. Phase III testing for a single drug often costs more than $100 million (DiMasi, Grabowski, and Hansen, 2016), which would be a massive burden on a government budget. Rather than financing late-stage testing for firms, the National Institutes of Health (NIH) generally funds low-cost, early-stage testing at universities. A notable exception was the financing of late-stage testing for COVID-19 vaccines through Operation Warp Speed (Snyder, Hoyt, and Gouglas, 2023).

Stylized fact 14: The regulator can signal to the firm that it will grant final approval to a drug that narrowly misses the efficacy threshold. The FDA has granted final approval to several drugs that did not prove efficacy. Furthermore, the FDA will advise some firms in advance of confirmatory testing about how to design its trial and what evidence will merit approval. Under a 2012 law, FDA staff support the developer of a potential breakthrough drug by "providing timely advice to, and interactive communication with, the sponsor regarding the development of the drug to ensure that the development program to gather the nonclinical and clinical data necessary for approval is as efficient as practicable".

C Calibration

We calibrate the model to illustrate its application and to inform policy recommendations. Our approach proceeds in five steps: First, we select a representative drug. Second, we

use the literature on drug development to determine values for the probability of technical success and the discount rate. Third, we estimate the testing cost for a representative drug. Fourth, we estimate drug doses and revenue for a representative drug. Fifth, we apply these parameter values to calculate the cumulative net present value for a representative drug.

We select a representative drug that received conditional approval: avelumab (brand name Bavencio). In 2017, the FDA granted conditional approval for avelumab to treat skin cancer based on tumor shrinkage. In 2023, full approval was granted based on clinical benefits. Following its initial approval for skin cancer, avelumab was also approved for bladder and kidney cancers.

The probability of receiving final approval (γ) is set at 0.45 (Aitken, Kleinrock, and Pritchett, 2024). The discount rate (δ) is set at 10.5% (DiMasi, Grabowski, and Hansen, 2016), and for simplicity, we assume that both the regulator and the firm use the same discount rate.

C.1 Estimated testing costs

We assume that Phase III clinical testing costs are evenly distributed over the first three years, with the regulator making the final approval decision at the start of the fifth year. The firm sells at monopoly prices for thirteen years, after which we assume the net revenue following generic entry is negligible.

To establish the lower and upper bounds for avelumab's testing costs, we use estimates from similar drugs. A prior study analyzed financial filings from companies with few drugs and limited joint costs, which made it easier to attribute expenses to a specific drug. Among 355 drugs approved by the FDA between 2009 and 2018, Wouters, McKee, and Luyten (2020) identified sufficient testing cost data for 63 drugs. Of these, 13 drugs had Phase III testing cost estimates in the same therapeutic class (antineoplastic and immunomodulating agents) as avelumab. For the low estimate, we used brigatinib's Phase III testing cost of \$98.1 million (in 2018 dollars). For the high estimate, we used \$630.7 million for sarilumab. However, these estimates are imprecise. For example, the authors expressed low confidence in the Phase III cost estimate for sarilumab.

C.2 Estimated doses and sales

Next, we estimate the total doses and revenue for the representative drug. We have only publicly available data on Medicare and Medicaid patients so we extrapolate to the total

market. Also, our analysis uses data on total sales of avelumab across all indications, meaning the reported sales figures are larger than those specific to the skin cancer indication. Table 1 summarizes the annual sales data for avelumab in 2022.

Table 1: Annual doses and prices for avelumab in 2022.

	Avelumab
Medicaid+Medicare ($MM)[a]	121
Total sales ($MM)[b]	290
Net revenue ($MM)[c]	217
Price per dose ($)[d]	86
Production cost per dose ($)[e]	21.5
Total doses (MM)[f]	3.37
Willingness to pay per dose ($)[g]	118

[a] Medicaid and Medicare Part B spending (Source: (Centers for Medicare & Medicaid Services, 2024))

[b] "Medicaid+Medicare" multiplied by 2.4 because Medicaid and Medicare accounted for 42% of total U.S. prescription drug spending at the time (Source: National Health Expenditures)

[c] "Total sales" multiplied by 0.75 to reflect cost of goods sold is 25% of total sales (Source: Harmand, `https://hardmanandco.com/2021-pharma-statistics-long-term-cost-underlying-ebit-analysis/`)

[d] We use Medicare prices because they are negotiated by commercial insurers and are similar to commercial prices. We use Part B spending for avelumab, because it is mainly administered by providers. (Source: Centers for Medicare and Medicaid Services)

[e] "Price per dose" multiplied by 0.25 to reflect production cost of goods sold is 25% of price (Source: Harmand, `https://hardmanandco.com/2021-pharma-statistics-long-term-cost-underlying-ebit-analysis/`)

[f] "Total sales"/"Price per dose"

[g] We assume that the buyer's surplus is half of the seller's. Hence, the willingness to pay is 1.375 times the unit price.

The price per dose is \$86, the production cost per dose is \$21.5, and the patient's willingness to pay \$118 (Table 1), and we have

$$\pi = \$86 - \$21.5 = \$64.5, \quad v_L = \$86, \quad v_H = \$118 - \$86 = \$32. \tag{38}$$

Table 2 summarizes the (discounted) annual doses of avelumab. When full conditional

approval is granted, the firm's total sales in the first 4 years are 4.36 million doses,[8] and from years 5 to 17, they are 15.43 million doses.[9] When no conditional approval is granted, if the drug receives final approval, the firm's total sales are 11.34 million.[10] By their definitions in Table 3, the values of q, λ_0, and λ_1 are determined by

$$\begin{cases} q = 4.36 \\ \lambda_0 = 11.34 \\ \lambda_0 + \lambda_1 = 15.43 \end{cases} \Rightarrow \begin{cases} q = 4.36 \\ \lambda_0 = 11.34 \\ \lambda_1 = 4.09 \end{cases} . \tag{39}$$

The discounted quantities sold from year 18 onward are given by[11]

$$\Delta = \sum_{i=18}^{\infty} \frac{3.37}{(1 + 10.5\%)^i} = 5.88. \tag{40}$$

Under fixed duration, if conditional approval is granted, the firm is only able to sell at the monopolistic price ($86) until year 13. In this case, the total discounted doses in the final approval period decrease by 2.89 million.[12] Conversely, the total discounted doses during the competitive period increase by the same amount, as Years 14 to 17 are now included in the competitive period.

The estimated model parameters are summarized in Table 3.

C.3 Estimated cumulative net present value

Next, we apply the parameter values to estimate the cumulative net present value for avelumab when the testing cost is at its maximum value, c_H = \$630.7 million, is evenly distributed over the first three years, and full conditional approval is granted. Table 4 displays the firm's cumulative net present value.

Figure 10 shows the firm's cumulative net present value under various scales of conditional approval. The upper and lower curves in each panel represent the firm's cumulative net profit when testing costs are at their lower bound (c_L) and upper (c_H) bound. For testing costs that fall within $[c_L, c_H]$, the corresponding curve lies in the shaded area. In the absence

[8] Summation of Years 1 to 4 of "Discounted annual sales" for "Full conditional approval" in Table 2.

[9] Summation of Years 5 to 17 of "Discounted annual sales" for "Full conditional approval" in Table 2.

[10] Summation of Years 5 to 17 of "Discounted annual sales" for "No conditional approval" in Table 2.

[11] The values 3.37 and 10.5% come from the "Total doses" line in Table 1 and the discount rate (δ), respectively.

[12] Summation of Years 14 to 17 of "Discounted annual sales" for "Full conditional approval" in Table 2.

of conditional approval ($\alpha_1 = 0$), firms with high testing costs will forgo testing, causing a potential societal loss.

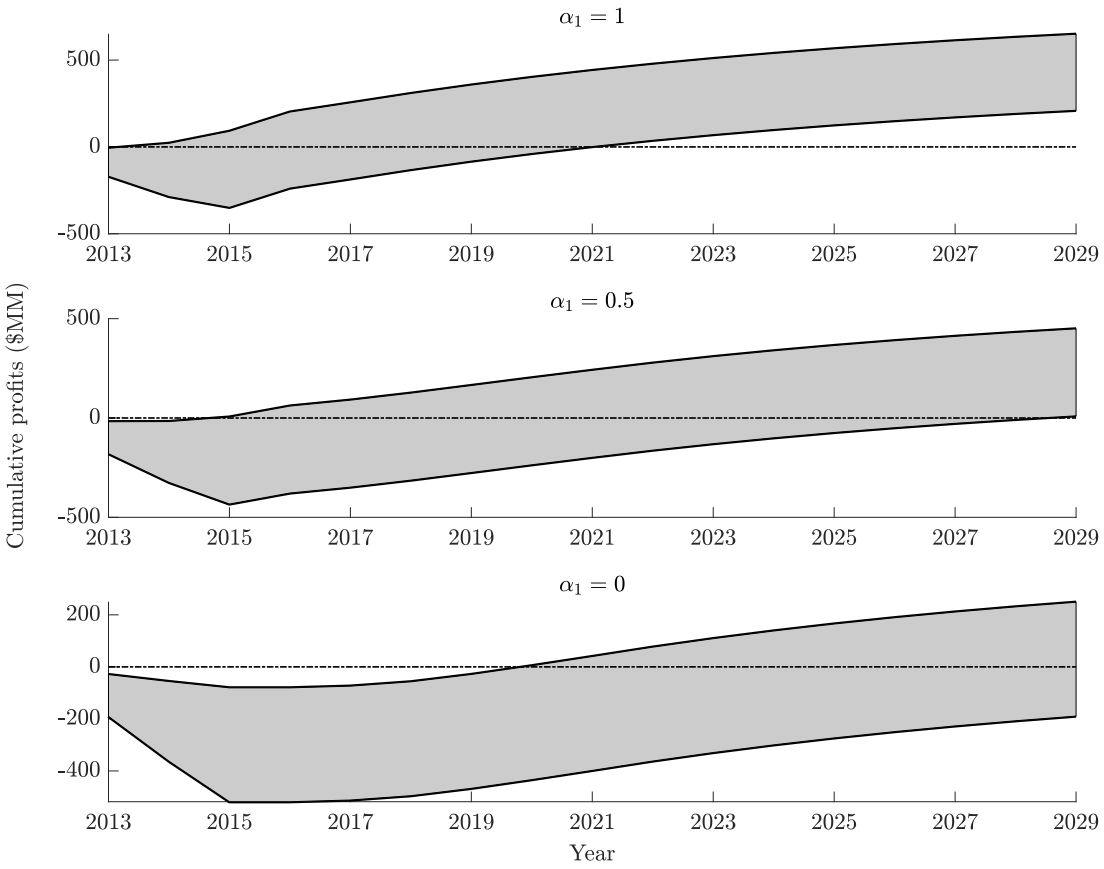

Figure 10: Firm's cumulative net profit under different scales of conditional approval and testing costs. Source: Authors' analysis using a representative drug avelumab.

Table 2: The firm's annual (discounted) doses when full/no conditional approval is granted to avelumab.

Year	Full conditional approval			No conditional approval		
	% of peak sales[a]	Annual doses (MM)[b]	Discounted annual doses (MM)[c]	% of peak sales[a]	Annual doses (MM)[b]	Discounted annual doses (MM)[c]
1	11%	0.37	0.34			
2	31%	1.05	0.86			
3	58%	1.96	1.45			
4	76%	2.56	1.72			
5	89%	3.00	1.82	11%	0.37	0.23
6	100%	3.37	1.85	31%	1.05	0.57
7	100%	3.37	1.68	58%	1.96	0.97
8	100%	3.37	1.52	76%	2.56	1.15
9	100%	3.37	1.37	89%	3.00	1.22
10	100%	3.37	1.24	100%	3.37	1.24
11	100%	3.37	1.12	100%	3.37	1.12
12	100%	3.37	1.02	100%	3.37	1.02
13	100%	3.37	0.92	100%	3.37	0.92
14	100%	3.37	0.83	100%	3.37	0.83
15	100%	3.37	0.75	100%	3.37	0.75
16	100%	3.37	0.68	100%	3.37	0.68
17	100%	3.37	0.62	100%	3.37	0.62

[a] Source: Robey and David (2016)

[b] Annual doses=Peak annual doses $\times$ % of peak sales

[c] Discounted annual doses in year i=Annual doses in year $i/(1+\delta)^i$, where $\delta = 10.5\%$ is the annual discount rate.

Table 3: Estimated model parameters for avelumab.

Notation	Description	Value	Source
δ	discount rate	10.5%	DiMasi, Grabowski, and Hansen (2016)
γ	probability of passing Phase III clinical trial	0.45	Aitken, Kleinrock, and Pritchett (2024)
c_L	lower bound of the testing cost	98.1($MM)	The estimated Phase III testing cost of a similar drug (brigatinib) as estimated by Wouters, McKee, and Luyten (2020).
c_H	upper bound of the testing cost	630.7($MM)	The estimated Phase III testing cost of a similar drug (sarilumab) as estimated by Wouters, McKee, and Luyten (2020).
v_L	regulator's loss from administering an ineffective dose	$86	See equation (38).
v_H	regulator's payoff from administering an effective dose	$32	See equation (38).
σ	price drop per dose after the monopolistic period	$64.5	See equation (38).
q	total discounted doses with full conditional approval before final approval	4.36(MM)	See equation (39).
λ_0	total discounted doses from final approval when $\alpha_1 = 0$	11.34(MM)	See equation (39).
λ_1	total boosted discounted doses in the monopolistic period after final approval when $\alpha_1 = 1$	4.09(MM)	See equation (39).
Δ	total discounted doses after the monopolistic period	5.88(MM)	See equation (40).
$\tilde{\Delta}$	total discounted doses lost in the monopolistic period after final approval due to earlier price control	2.89(MM)	The total discounted doses from years 14 to 17 amount to 2.89 million.

Table 4: The firm's cumulative net present value when its testing cost is the maximum value.

Year	Discounted net revenue (if in market) ($MM)[a]	Discounted expected net revenue ($MM)[b]	Annual testing cost ($MM)[c]	Discounted annual testing cost ($MM)	Discounted annual net profit ($MM)[d]	Cumulative net present value ($MM)[e]
1	22	22	210	190	-169	-169
2	55	55	210	172	-117	-286
3	93	93	210	156	-63	-348
4	111	111	0	0	111	-238
5	117	53	0	0	53	-185
6	119	54	0	0	54	-131
7	108	49	0	0	49	-83
8	98	44	0	0	44	-39
9	88	40	0	0	40	1
10	80	36	0	0	36	37
11	72	33	0	0	33	70
12	65	29	0	0	29	99
13	59	27	0	0	27	126
14	54	24	0	0	24	150
15	49	22	0	0	22	172
16	44	20	0	0	20	191
17	40	18	0	0	18	209

[a] "Discounted annual doses" column from Table 2 times "Price per dose" from Table 1

[b] Discounted expected net revenue = Discounted net revenue if the drug is in mraket × Pr(the drug is in market), and the probabilistic term is 1 and 0.45 in years 1 to 4 and 5 to 17, respectively.

[c] The testing cost $ 630.7 million is evenly distributed over the first three years.

[d] Discounted annual net profit = Discounted expected net revenue (column 3) - Discounted annual testing cost (column 5).

[e] Cumulative NPV in year i = sum of the discounted annual net profit from years 1 to i.

D Proofs

D.1 Proofs of section 3

Proof of Lemma 1. The incentive-compatibility constraint can be rewritten as

$$\Pi(c) \geq \Pi(c') + \tau(c') \cdot (c' - c),$$

where $\Pi(c) \equiv \hat{\Pi}(c, c)$. Flipping the role of c and c' yields

$$\Pi(c') \geq \Pi(c) + \tau(c) \cdot (c - c').$$

Combining the two inequalities above yields the monotonicity of τ.

By envelope theorem, we have

$$\Pi'(c) = -\tau(c) \leq 0,$$

which implies that

$$\Pi(c) = \Pi(c_H) + \int_c^{c_H} \tau(y)\,\mathrm{d}y.$$

Therefore, it is sufficient to impose the individual rationality constraint on type c_H and

$$\alpha_1(c) = \frac{1}{1 + \lambda_1\,\gamma} \cdot \left[\Pi(c_H) + \int_c^{c_H} \tau(y)\,\mathrm{d}y + (c - \lambda_0\,\gamma) \cdot \tau(c)\right].$$

To prove the other direction, for any $c < c'$, we have

$$\begin{aligned}
\Pi(c) &= \Pi(c_H) + \int_c^{c_H} \tau(y)\,\mathrm{d}y \\
&= \Pi(c_H) + \int_{c'}^{c'} \tau(y)\,\mathrm{d}y + \int_{c'}^{c_H} \tau(y)\,\mathrm{d}y \\
&= \Pi(c') + \int_c^{c'} \tau(y)\,\mathrm{d}y \\
&\geq \Pi(c') + \tau(c') \cdot (c' - c),
\end{aligned}$$

where the inequality follows from the monotonicity of τ. Constraint (IR) is implied by

$$\Pi(c) \geq \Pi(c_H) \geq 0.$$

This completes the proof. ■ □

Proof of Proposition 1. In the optimization problem (17), the value of γ is fixed. Throughout this proof, we omit the dependence of w on γ for notation simplicity.

In the first case, we have

$$\gamma \geq \check{\gamma} \quad \Leftrightarrow \quad v_E + \lambda_1 \gamma \geq 0.$$

Ignoring all constraints, it is optimal to set $\tau^*(c) = \alpha_1^*(c) = 1$ for any $c \in [c_L, c_H]$. Since the solution to the relaxed problem is feasible, it remains optimal for the original problem.

In the second case, we have

$$\gamma < \check{\gamma} \quad \Leftrightarrow \quad v_E + \lambda_1 \gamma < 0.$$

Ignoring all constraints, it is optimal to set $\tau^*(c) = 1$ and $\alpha_1^*(c) = 0$ for any $c \in [c_L, c_H]$. Since the solution to the relaxed problem is feasible, it remains optimal for the original problem.

In the third case, we first ignore constraints (13) and (18), and solving the relaxed problem point-wise yields

$$\tau(c) = \begin{cases} 1, & \forall c \in [c_L, c_0], \\ 0, & \forall c \in (c_0, c_H], \end{cases}$$

satisfying (13) automatically. Plugging the expression of τ into (14) yields

$$\alpha_1(c) = \begin{cases} \dfrac{c_0 - \lambda_0 \gamma}{1 + \lambda_1 \gamma}, & \forall c \in [c_L, c_0], \\ 0, & \forall c \in (c_0, c_H]. \end{cases}$$

Under Assumption 1, when $c_0 \geq \lambda_0 \gamma$, the feasibility constraint (18) holds automatically, and thus the point-wise solution above is feasible.

In the last case, it is straightforward to verify that our candidate solution is feasible. When $c_L \leq c_0 < \lambda_0 \gamma$, define

$$\begin{aligned} \phi &= \frac{1}{c_1 - c_2} \int_{c_L}^{c_1} w(y)\, f(y)\, \mathrm{d}y - \frac{v_E + \lambda_1 \gamma}{1 + \lambda_1 \gamma}, \\ \nu(c) &= \frac{1}{(c_2 - c)^2} \left[\int_{c_L}^{c} w(y)\, f(y)\, \mathrm{d}y + (c_2 - c)\, w(c)\, f(c) \right], && \forall c \in [c_L, c_1], \\ \sigma(c) &= \frac{1}{c_2 - c_1} \int_{c_L}^{c_1} w(y)\, f(y)\, \mathrm{d}y + w(c)\, f(c), && \forall c \in [c_1, c_2], \text{ and} \end{aligned}$$

$$\mu(c) = -\frac{1}{c_2 - c_1} \int_{c_L}^{c_1} w(y)\, f(y)\, \mathrm{d}y - w(c)\, f(c), \qquad \forall c \in [c_2, c_H],$$

where $c_2 = \lambda_0\, \gamma$ and c_1 is determined by

$$\int_{c_L}^{c_1} w(y)\, f(y)\, \mathrm{d}y + (c_2 - c_1)\, w(c_2)\, f(c_2) = 0.$$

We first verify the monotonicity of $w(c)\, f(c)$ on $[c_L, c_2]$. By definition, we have

$$w(c)\, f(c) = \frac{1}{1 + \lambda_1\, \gamma} \left[\left[(\lambda_0\, \gamma + z)\, (1 + \lambda_1\, \gamma) + \underbrace{(v_E + \lambda_1\, \gamma)\, (c - \lambda_0\, \gamma)}_{\text{non-negative and decreasing in } c} \right] f(c) + \underbrace{(v_E + \lambda_1\, \gamma)}_{<0} F(c) \right].$$

Under Assumption 2, we know that $w(c)\, f(c)$ is decreasing in c.

Then we prove the unique existence of c_1 on $[c_L, c_2]$. Let

$$\psi(c) \equiv \int_{c_L}^{c} w(y)\, f(y)\, \mathrm{d}y + (c_2 - c)\, w(c_2)\, f(c_2), \quad \forall c \in [c_L, c_2].$$

Because

$$\begin{aligned} &\psi'(c) = w(c)\, f(c) - w(c_2)\, f(c_2) \geq 0, \\ &\psi(c_L) = (c_2 - c_L)\, w(c_2)\, f(c_2) < 0, \text{ and} \\ &\psi(c_2) = \int_{c_L}^{c_2} w(y)\, f(y)\, \mathrm{d}y > 0, \end{aligned}$$

the unique existence of c_1 is guaranteed. The first inequality follows from the monotonicity of $w(c)\, f(c)$, and the last inequality follows from the fact that under the mechanism granting no conditional approval, the regulator's expected payoff is strictly positive.

Next, we will prove that ϕ, ν, σ, and μ are all non-negative. By definition of c_1, we have

$$\phi = w(c_2)\, f(c_2) - \frac{v_E + \lambda_1\, \gamma}{1 + \lambda_1\, \gamma} = (\lambda_0\, \gamma + z)\, f(c_2) - \frac{\overbrace{v_E + \lambda_1\, \gamma}^{<0}}{1 + \lambda_1\, \gamma} \left[1 - F(c_2)\right] \geq 0.$$

The nonnegativity of ν boils down to

$$h(c) \equiv \int_{c_L}^{c} w(y)\, f(y)\, \mathrm{d}y + (c_2 - c)\, w(c)\, f(c) \geq 0, \quad \forall c \in [c_L, c_1].$$

Taking derivative yields

$$h'(c) = (c_2 - c)\,\frac{d}{dc}\left(w(c)\,f(c)\right) \leq 0, \quad \forall c \in [c_L, c_1],$$

where the inequality follows from the monotonicity of $w(c)\,f(c)$. Therefore, the desired property is reduced to

$$h(c_1) = \int_{c_L}^{c_1} w(y)\,f(y)\,\mathrm{d}y + (c_2 - c_1)\,w(c_1)\,f(c_1) \geq 0,$$

which follows from

$$\int_{c_L}^{c_1} w(y)\,f(y)\,\mathrm{d}y + (c_2 - c_1)\,w(c_1)\,f(c_1) \geq \int_{c_L}^{c_1} w(y)\,f(y)\,\mathrm{d}y + (c_2 - c_1)\,w(c_2)\,f(c_2) = 0.$$

For any $c \in [c_1, c_2]$ and $c \in [c_2, c_H]$, we have

$$\sigma'(c) = \frac{d}{dc}\left(w(c)\,f(c)\right) \leq 0$$

and

$$\mu(c) \geq \mu(c_2),$$

respectively. The first and second inequalities follow from the monotonicity of $w(c)\,f(c)$ on $c \in [c_L, c_2]$ and

$$\begin{aligned}
\mu(c) \geq \mu(c_2) \;&\Leftrightarrow\; w(c)\,f(c) \leq w(c_2)\,f(c_2) \\
&\Leftrightarrow\; \left[(\lambda_0\,\gamma + z)\,(1 + \lambda_1\,\gamma) + (v_E + \lambda_1\,\gamma)\,(c - \lambda_0\,\gamma)\right]\,f(c) + (v_E + \lambda_1\,\gamma)\,F(c) \\
&\leq \left[(\lambda_0\,\gamma + z)\,(1 + \lambda_1\,\gamma) + (v_E + \lambda_1\,\gamma)\,(c_2 - \lambda_0\,\gamma)\right]\,f(c_2) + (v_E + \lambda_1\,\gamma)\,F(c_2)
\end{aligned}$$

for any $c \in [c_2, c_H]$, where the last step follows from

$$v_E + \lambda_1\,\gamma < 0, \quad f(c) \leq f(c_2), \quad \text{and} \quad c - \lambda_0\,\gamma > 0 = c_2 - \lambda_0\,\gamma.$$

The desired non-negativity conditions are automatically implied by $\sigma(c_2) = \mu(c_2) = 0$.

Because

$$\int_{c_L}^{c} \nu(y) = \frac{1}{c_2 - c}\int_{c_L}^{c} w(y)\,f(y)\,\mathrm{d}y, \quad \forall c \in [c_L, c_1],$$

we have

$$
\begin{aligned}
&-\phi - \int_{c_L}^{c_1} \nu(y)\,\mathrm{d}y = \frac{v_E + \lambda_1\,\gamma}{1+\lambda_1\,\gamma}, \\
&(c_2 - c)\,\nu(c) - \int_{c_L}^{c} \nu(y)\,\mathrm{d}y = w(c)\,f(c), && \forall c \in [c_L, c_1], \\
&\sigma(c) - \int_{c_L}^{c_1} \nu(y)\,\mathrm{d}y = w(c)\,f(c), && \forall c \in [c_1, c_2], \text{ and} \\
&-\mu(c) - \int_{c_L}^{c_1} \nu(y)\,\mathrm{d}y = w(c)\,f(c), && \forall c \in [c_2, c_H].
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
&\phi \cdot [-\Pi(c_H)] - \int_{c_L}^{c_1} \left[\Pi(c_H) + \int_{c}^{c_H} \tau(y)\,\mathrm{d}y + (c - \lambda_0\,\gamma)\cdot\tau(c)\right]\cdot\nu(c)dc \\
&\quad + \int_{c_1}^{c_2} \tau(c)\cdot\sigma(c)dc - \int_{c_2}^{c_H} \tau(c)\cdot\mu(c)dc \\
=&\left[-\phi - \int_{c_L}^{c_1}\nu(c)dc\right]\cdot\Pi(c_H) + \int_{c_L}^{c_1}\left[(c_2-c)\,\nu(c) - \int_{c_L}^{c}\nu(y)\,\mathrm{d}y\right]\cdot\tau(c)dc \\
&\quad + \int_{c_1}^{c_2}\left[\sigma(c) - \int_{c_L}^{c_1}\nu(y)\,\mathrm{d}y\right]\cdot\tau(c)dc + \int_{c_2}^{c_H}\left[-\mu(c) - \int_{c_L}^{c_1}\nu(y)\,\mathrm{d}y\right]\cdot\tau(c)dc \\
=&\frac{v_E + \lambda_1\,\gamma}{1+\lambda_1\,\gamma}\cdot\Pi(c_H) + \int_{c_L}^{c_H} w(c)\cdot\tau(c)\;\mathrm{d}F(c).
\end{aligned}
$$

Therefore, the objective function is bounded by $\int_{c_1}^{c_2}\sigma(c)dc$. Substituting our candidate primal solution into the objective function yields $\int_{c_L}^{c_2} w(y)\,f(y)\,\mathrm{d}y$, which is the same value as the upper bound, i.e.,

$$
\begin{aligned}
\int_{c_1}^{c_2}\sigma(y)\,\mathrm{d}y &= \int_{c_1}^{c_2}\left[\frac{1}{c_2-c_1}\int_{c_L}^{c_1} w(z)\,f(z)dz + w(y)\,f(y)\right]\mathrm{d}y \\
&= \int_{c_L}^{c_1} w(y)\,f(y)\,\mathrm{d}y + \int_{c_1}^{c_2} w(y)\,f(y)\,\mathrm{d}y \\
&= \int_{c_L}^{c_2} w(y)\,f(y)\,\mathrm{d}y.
\end{aligned}
$$

This establishes the optimality of our candidate solution.

When $c_0 = -\infty$, the virtual valuation function w is always non-positive on $[c_L, c_H]$. Moreover, the coefficient of $\Pi(c_H)$ is also negative; hence, it is evident that the optimal mechanism is to not request testing and to grant no conditional approval.

This completes the proof. ■ □

Proof of Proposition 2. When $a_1(c) = 0$ for any $c, c' \in [c_L, c_H]$, both parties' payoff functions can be written as

$$\hat{\Pi}(c, c') = \tau(c') \left(-c + \lambda_0 \gamma \pi\right)$$

and

$$\tilde{U}_R(\tau, 0) = \int_{c_L}^{c_H} \tau(c) \left(\lambda_0 \gamma v_H + z\right) \mathrm{d}F(c).$$

From the regulator's perspective, it is always optimal to induce participation. The individual rationality constraint implies that only firms with costs below $\lambda_0 \gamma \pi$ have the incentive to conduct the testing. Therefore, the optimal mechanism is given by

$$\tau(c) = \begin{cases} 1, & \forall c \in [c_L, \lambda_0 \gamma \pi], \\ 0, & \forall c \in (\lambda_0 \gamma \pi, c_H], \end{cases}$$

and the corresponding regulator's expected payoff is $\left(\lambda_0 \gamma v_H + z\right) F(\lambda_0 \gamma \pi)$.

This completes the proof. ■ □

Proof of Proposition 3. The proof closely follows the structure of the proof of Proposition 2, and thus the details are omitted here. When $a_1 = 1$, firms capable of self-funding will participate.

When the scale of conditional approval is restricted to 0 or 1, the regulator compares its expected payoff under both options and commits to the one that generates a higher expected payoff.

This completes the proof. ■ □

D.2 Proofs of section 4

Proof of Proposition 4. Let Ω and $\bar{\Omega}$ denote the feasible sets of the optimization problems (31) and (37), respectively. To prove the desired result, it suffices to show that $\Omega \subseteq \bar{\Omega}$.

Let $\mathcal{M} \equiv \{\tau, a_1, a_2\}$ denote a feasible solution to the optimization problem (31). It is straightforward to verify that $\bar{\mathcal{M}} \equiv \{\tau, \alpha_1, \alpha_2, \alpha_3\}$, where

$$\begin{aligned} &\alpha_1(c) = \tau(c)\, a_1(c), && \forall c \in [c_L, c_H], \\ &\alpha_2(c, \eta) = \tau(c)\, \left[\lambda_0 + \lambda_1\, a_1(c)\right]\, a_2(c, \eta), && \forall (c, \eta) \in [c_L, c_H] \times [0, 1], \\ &\alpha_3(c, \eta) = \tau(c)\, a_2(c, \eta), && \forall (c, \eta) \in [c_L, c_H] \times [0, 1], \end{aligned}$$

is feasible to the relaxed problem (37) and yields the same objective function value. The feasibility follows from

$$\frac{\alpha_2(c,\eta)}{\lambda_0+\lambda_1} \leq \alpha_3(c,\eta) = \frac{\alpha_2(c,\eta)}{\lambda_0+\lambda_1\, a_1(c)} \leq \frac{\alpha_2(c,\eta)}{\lambda_0},$$
$$\alpha_3(c,\eta) = \tau(c)\, a_2(c,\eta) \leq \tau(c),$$

where the first and second inequality conditions are implied by $0 \leq a_1(c) \leq 1$ and $0 \leq a_2(c,\eta) \leq 1$, respectively. The same objective function value is guaranteed by definition.

This completes the proof. ■ □

Proof of Proposition 5. The proof is straightforward, as the feasible set of the original problem (31) is determined by the constraints independent of α_3 in the relaxed problem (37). By construction, we have

$$\int_0^1 \tilde{\alpha}_2(c,\eta)\,\mathrm{d}G(\eta) = \int_0^1 \hat{\alpha}_2(c,\eta)\,\mathrm{d}G(\eta), \quad \forall c \in [c_L, c_H].$$

Hence, constraints (IC), (IR), (2), (9), and (34) hold automatically. Because U_R^* is the optimal value of the objective function, the feasibility property implies that

$$\tilde{U}_R(\hat{\tau}, \hat{\alpha}_1, \tilde{\alpha}_2) \leq U_R^*.$$

This completes the proof. ■ □

E Performance of the approximation

Figure 11 illustrates the performance ratio (PR) of the candidate solution compared with the upper bound, i.e.,

$$\mathrm{PR} = \frac{\tilde{U}_R^*}{\bar{U}_R^*}.$$

Under the given model parameters, the candidate solution consistently achieves a performance ratio exceeding 0.975, implying that it is essentially optimal.

[13] The unsmooth part arises from the discretization when solving the linear optimization problem numerically.

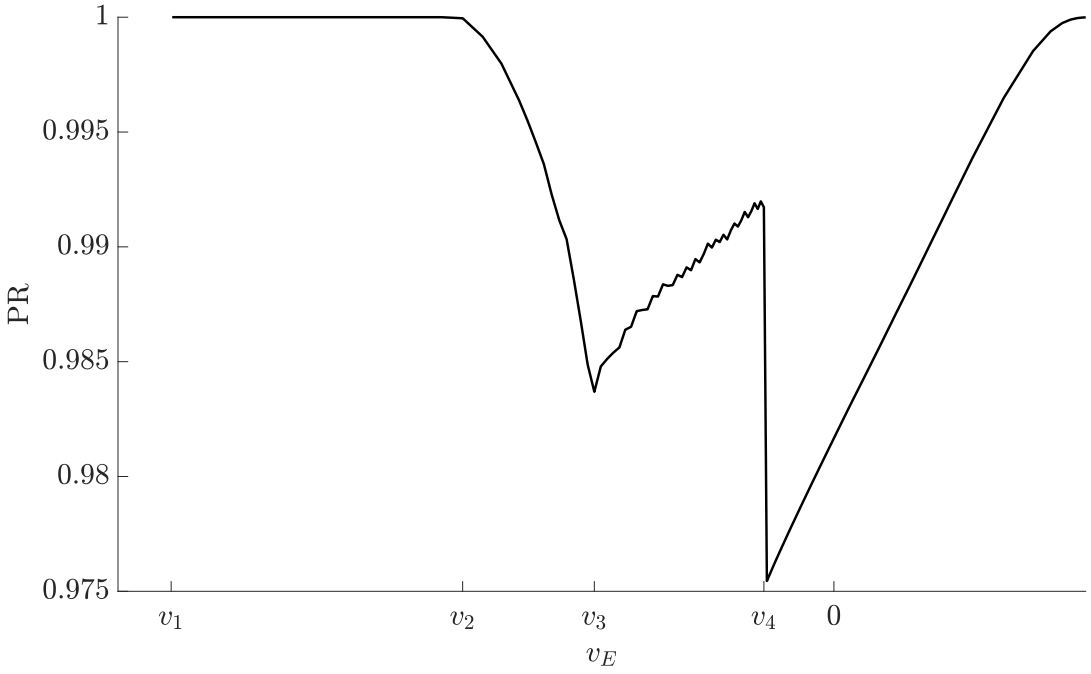

Figure 11: Performance ratio of our candidate mechanism compared to the upper bound.[13]The model parameters are summarized in Table 3. The firm's private information c is uniformly distributed, and the regulator's utility function v is linear.