Pricing Mechanisms for Multi-Indication Drugs

Pedro Pita Barros[1], Giovanni Righetti[1], and Luís Sá*[1]

[1] *Nova School of Business and Economics*

July 2025

Abstract

Multi-indication drugs pose pricing challenges, affecting therapeutic benefits and patient access. We model how a manufacturer strategically chooses clinically eligible patient populations for each indication, recognizing that broader populations increase demand but reduce expected therapeutic benefits and thus allowable prices. We identify mechanisms that induce the manufacturer to maximize total benefit while ensuring profits cannot exceed benefit generated: indication-specific prices equal to expected benefits, population-weighted uniform prices, and two-part tariffs. This holds when indications are introduced simultaneously or sequentially, provided prices fully adjust as indications are added. Price caps distort patient selection, explaining empirically observed strategies.

Keywords: Multi-indication drugs; Indication-based pricing; Weighted pricing; Two-part tariffs; Clinical trials.

JEL Classification: I11, I18, L12, L65, O31.

*Corresponding author. Nova School of Business and Economics, Universidade NOVA de Lisboa, Campus de Carcavelos, 2775-405 Carcavelos, Portugal. E-mail: luis.sa@novasbe.pt.

1 Introduction

In July 2025, *The New Yorker* magazine featured a piece asking whether Artificial Intelligence (AI) can find cures for diseases using existing drugs (Khullar, 2025). It discussed the possibility of AI ushering in an era of widespread drug repurposing by streamlining the process of identifying previously unknown applications for known drugs, which have been vetted for safety and whose mechanism of action is understood. Drug repurposing, however, is far from new, and long predates the rise of powerful AI models.

Drugs used to treat multiple indications—different diseases, various stages of the same disease, distinct treatment phases, or in combination with other therapies—are commonly known as multi-indication drugs. Prevalent in oncology, diabetes, respiratory conditions, and immunology, multi-indication drugs may receive approval for a remarkable number of indications. For instance, the cancer drugs nivolumab (Opdivo®) and pembrolizumab (Keytruda®) have been authorized for 13 and 16 different indications, respectively, since their first indication approval in 2014 (Mills et al., 2023). The therapeutic breadth of multi-indication drugs is equally remarkable. Alemtuzumab treats both chronic lymphocytic leukemia (Campath®) and multiple sclerosis (Lemtrada®). Sildenafil, originally marketed for erectile dysfunction (Viagra®), now also treats pulmonary arterial hypertension (Revatio®). These examples illustrate not just therapeutic versatility but also the vast differences in patient populations and treatment values across indications. Against this backdrop, it comes as no surprise that multi-indication drugs are increasingly common (Preckler & Espín, 2022; Mills et al., 2023) and expected to become even more so by industry stakeholders (ISPOR Europe, 2025).

Yet multi-indication drugs pose fundamental challenges for pricing and reimbursement (Cole et al., 2021; Preckler & Espín, 2022; Mills et al., 2023). From the perspective of economic theory, these drugs present a textbook case for price discrimination. A single product serves distinct consumer groups (indications) with heterogeneous valuations—multi-indication drugs often generate vastly different therapeutic benefits across indications (Preckler & Espín, 2022). Moreover, as with all pharmaceutical innovation, R&D incentives require pricing above marginal cost. Given these conditions—identifiable consumer segments, heterogeneous valuations, and market power—the natural question arises: why not price discriminate across indications?

The answer, as international experience reveals, is far from straightforward. Despite the case for indication-based pricing—a form of price discrimination that would align prices with therapeutic

benefits—practical implementation proves elusive, and its theoretical merits disputed.

Consider the remarkable diversity of pricing schemes across countries (Campillo-Artero et al., 2020; Cole et al., 2021; Michaeli et al., 2022; Mills & Kanavos, 2023; Heine et al., 2024; ISPOR Europe, 2025). Anchor pricing sets prices equal to or based on the first indication's price, as seen in Turkey and the Netherlands. Weighted-average uniform pricing (also known as "blended" pricing) establishes a single price given by the average price across indications, weighted by volume in Spain or by both volume and potentially clinical value in Belgium, France, Germany, Australia, and Canada—though the value component often remains unclear or is negotiated confidentially. Indication-based pricing creates indication-specific prices, either by assigning different brand names to different indications and setting prices for each brand or through differential discounting from a single list price, as practiced in England, Scotland, Switzerland, and by some U.S. insurers and pharmacy benefit managers. Notably, no country currently implements pure indication-based pricing with different list prices for different indications of the same drug.[1]

This international heterogeneity reflects deeper theoretical uncertainty. The academic literature, focused on comparing indication-based pricing with variations of uniform pricing, has yielded mixed results. On the one hand, despite simplifying market regulation, uniform pricing can create inefficiencies: high-value indications may subsidize low-value ones, deterring both manufacturers from seeking approval for additional indications and payers from covering low-value uses (Adida, 2024). Moreover, in the presence of competition, uniform pricing might create a softening-of-competition effect, ultimately leading to higher prices (Brekke et al., 2025). On the other hand, indication-based pricing allows manufacturers to extract more surplus through third-degree price discrimination (Adida, 2024; Brekke et al., 2025); the health plan might experience higher prices for patients who benefit the most, higher utilization by patients who benefit least, and higher overall spending (Chandra & Garthwaite, 2017). Recent theoretical and simulation analyses further suggest that whereas indication-based pricing can improve dynamic incentives for innovation, it does not necessarily increase consumer welfare compared to uniform pricing, and can in fact reduce overall patient surplus (Jiang et al., 2024). These mixed findings may reflect differences in institutional assumptions. Goldhaber-Fiebert and Cipriano (2023) formally model the strategic inter-

[1] These categories represent archetypes; most countries employ hybrid approaches. Belgium uses weighted pricing nationally while allowing confidential indication-specific contracts. Spain maintains uniform national pricing with regional variations. Italy transitioned from different prices per indication to a single weighted price (Rossini et al., 2024). Confidential agreements further complicate the precise classification of national approaches.

action between manufacturers and payers under various constraints on pricing and reimbursement across multiple indications, showing that whereas uniform pricing can lead to inefficient coverage or underuse of high-value indications, indication-specific pricing paired with indication-specific reimbursement can achieve first-best outcomes.

This article contributes simultaneously to the ongoing policy debate and to the academic literature by developing a novel model to analyze alternative modes of pricing multi-indication drugs and the therapeutic benefits they generate. We examine schemes implemented in practice—indication-based pricing and several forms of uniform pricing—alongside emerging alternatives like two-part tariffs, to provide a comprehensive account of pricing alternatives and how they interact with Health Technology Assessment (HTA).

To our knowledge, this is the first model to endogenize patient population selection in pharmaceutical pricing. It connects a pharmaceutical manufacturer's choice of clinically eligible patient populations for each indication (alternatively, patient types studied in clinical trials or indication label breadth) to drug prices. In our model, a health plan employs HTA-based coverage criteria, linking prices to estimated expected therapeutic benefits and then offering the drug to all clinically eligible patients. Recognizing this, the manufacturer faces a trade-off: studying broader patient populations in clinical trials expands demand for each indication but includes patients with lower therapeutic benefit, thereby reducing the estimated expected benefit and the price(s) the health plan will pay. By strategically choosing which patients to study—and thus who becomes clinically eligible and ultimately given the drug—the manufacturer simultaneously determines both the demand for each indication and the price(s) it can command.

Pharmaceutical manufacturers retain discretion in determining clinically eligible patient populations through trial study protocols, particularly through patient inclusion and exclusion criteria (Cherubini et al., 2011; Schmidt et al., 2014). These criteria encompass age, gender, disease stage and severity, previous treatments, concomitant medications, and comorbidities. Such eligibility decisions can be implemented directly through explicit enrollment rules or indirectly via recruitment strategies and trial design features—choices that demonstrably affect trial outcomes (Hill et al., 2008; Nordon et al., 2023). Evidence reveals systematic patterns in how these choices are exercised. Clinical trials frequently restrict enrollment to lower-risk patients (Jin et al., 2017), especially when industry-sponsored (Van Spall et al., 2007; Duma et al., 2019). Industry-sponsored trials, in turn, report favorable efficacy outcomes more often than those funded by other sources (Lundh et al., 2017; Siena et al., 2023). These patterns align with manufacturers designing study protocols to

demonstrate higher therapeutic benefit.

We establish that efficient mechanisms—those that maximize total therapeutic benefit across indications by yielding clinically eligible patient populations that include only patients who derive non-negative therapeutic benefit from the drug—align manufacturer profit maximization with total therapeutic benefit maximization. Three mechanisms achieve this alignment when indications are introduced simultaneously: (i) indication-specific prices equal to the indication's expected therapeutic benefit, (ii) a uniform price equal to the patient population-weighted average of expected therapeutic benefits of all indications, and (iii) two-part tariffs composed of unit-prices and a lump-sum payment, regardless of the HTA-based coverage and pricing criteria adopted by the health plan. We also analyze inefficient mechanisms, a uniform price equal to the unweighted average of expected therapeutic benefits of all indications and a uniform price anchored at a single indication's expected benefit, revealing how these institutional designs systematically distort patient selection.

We then assess whether efficiency carries over into a dynamic setting, and show that the three efficient mechanisms remain efficient when indications are introduced sequentially, but only under full price flexibility. In this case, these mechanisms continue to ensure that manufacturers' profits equal—but cannot exceed—the total therapeutic benefit they create, even as new indications are added over time.

When regulatory constraints prevent upward price adjustments, as observed in most countries (Michaeli et al., 2022; Mills & Kanavos, 2023), even forward-looking manufacturers cannot fully restore efficiency. Most strikingly, we show that such constraints can generate a specific pattern of strategic behavior: manufacturers restrict patient populations in first indications to anchor prices at higher levels, then expand populations in subsequent indications to maximize revenues at these elevated prices. This theoretical prediction maps precisely onto the empirical pattern documented across multiple countries, where first indications show high therapeutic benefit for narrow populations, followed by lower-benefit indications for broader populations.

The rest of the article is organized as follows. Section 2 presents the model, introducing the manufacturer's strategic choice of clinically eligible populations. Section 3 analyzes pricing mechanisms when all indications are introduced simultaneously, characterizing both efficient and inefficient mechanisms. Section 4 extends the analysis to sequential indication introduction, examining how price flexibility affects efficiency and strategic behavior. Section 5 discusses the results and offers policy implications.

2 Model Setup

Consider a monopolist pharmaceutical manufacturer that produces a drug with potential therapeutic applications across n distinct indications, indexed by $i \in \{1, 2, ..., n\}$. Production occurs at zero marginal cost. Each indication represents a separate market with a unit mass of patients, all of whom are covered by a single health plan that determines drug coverage based on cost-effectiveness criteria.

2.1 Patient Heterogeneity and Therapeutic Benefits

Within each indication i, patients differ in their therapeutic mismatch with the drug, characterized by a parameter $x_i \in [0, 1]$, independently distributed across indications with continuously differentiable density $f_i(x_i) > 0$ and corresponding cumulative distribution function $F_i(x_i)$.

The therapeutic benefit a patient derives from treatment for indication i is given by $b_i(x_i)$, with $b_i'(x_i) < 0$. Without loss of generality and for clarity of exposition, we assume

$$b_i(x_i) = v_i - \tau_i x_i, \tag{1}$$

where $v_i > 0$ represents the maximum therapeutic benefit (achieved when therapeutic mismatch $x_i = 0$) and $\tau_i > 0$ measures the rate at which benefits decline with therapeutic mismatch. We assume $0 < v_i < \tau_i$, which implies that patients with sufficiently high therapeutic mismatch derive negative benefit from treatment.[2]

2.2 Strategic Choice of Clinically Eligible Patient Populations

The manufacturer strategically chooses the clinically eligible patient population for each indication. For indication i, the manufacturer selects a threshold $\hat{x}_i \in [0, 1]$ that determines which patients to include in clinical trials. Only patients with therapeutic mismatch $x_i \leq \hat{x}_i$ are studied, and this defines the breadth of the drug's label for that indication in terms of patient population. Throughout this article, we use the term "clinically eligible patient population" to refer to patients with $x_i \leq \hat{x}_i$ and the manufacturer's choice of $\hat{x}_i$.

This decision reflects the trade-off inherent in the choice of clinically eligible patient populations: narrower populations (lower $\hat{x}_i$) yield higher expected therapeutic benefits but exclude potential

[2] Our main results follow from the equivalence between manufacturer profits and total therapeutic benefit. See the appendix for a proof that this equivalence is independent of the functional form of $b_i(x_i)$.

patients from treatment, whereas broader populations (higher $\hat{x}_i$) expand market access but dilute measured effectiveness.

2.3 Coverage Determination and Pricing

The health plan's coverage decision follows HTA methodology by relating prices to expected therapeutic benefits. Coverage is granted for indication i if and only if the price p_i satisfies

$$p_i \leq \mathcal{P}\left(\mathbb{E}[b_1(x_1)|x_1 \leq \hat{x}_1], ..., \mathbb{E}[b_i(x_i)|x_i \leq \hat{x}_i], ..., \mathbb{E}[b_n(x_n)|x_n \leq \hat{x}_n]\right), \tag{2}$$

where $\mathcal{P} : \mathbb{R}_+^n \to \mathbb{R}_+$ is non-decreasing in each argument and represents the health plan's willingness to pay as a function of expected therapeutic benefits across indications. Note that although we write $\mathcal{P}$ with all indications' expected benefits as potential arguments, the function may depend on all, a subset, or just a single indication's benefit, capturing different coverage criteria.

This formulation encompasses several economically relevant coverage criteria. When the manufacturer is allowed to charge a different price for each indication (i.e., under indication-based pricing), $\mathcal{P}$ depends only on $\mathbb{E}[b_i(x_i)|x_i \leq \hat{x}_i]$, and each indication is evaluated independently. Alternatively, $\mathcal{P}$ may aggregate information across indications—e.g., as a simple average of expected benefits across all indications, or as a population-weighted average where larger patient populations receive greater weight in the pricing determination. These alternative formulations capture different approaches to value-based pricing in multi-indication settings.

For indications lacking existing treatments, (2) reduces to a standard cost-effectiveness threshold when $\mathcal{P}$ depends only on the expected benefit of the indication being evaluated. Specifically, setting $\mathcal{P}(\mathbb{E}[b_i(x_i)|x_i \leq \hat{x}_i]) = \mathbb{E}[b_i(x_i)|x_i \leq \hat{x}_i]$ and normalizing, without loss of generality, the cost-effectiveness threshold to unity yields the condition that price cannot exceed expected therapeutic benefit.

2.4 Market Structure and the Manufacturer's Problem

Once coverage is granted for indication i, the health plan provides the drug to all patients in the clinically eligible population, generating demand $F_i(\hat{x}_i)$. This institutional feature, whereby coverage decisions determine access for entire clinically eligible patient populations rather than individual patients, transforms the standard monopoly pricing problem.

Instead of facing a traditional downward-sloping demand curve driven by consumer preferences directly, the manufacturer faces a "coverage-mediated demand" where the relationship between

clinically eligible patient population—and hence realized demand, $F_i(\hat{x}_i)$—and price p_i is mediated through the expected therapeutic benefit rather than consumer willingness to pay.

The manufacturer's strategic variables are thus the clinically eligible patient population thresholds $\{\hat{x}_i\}_{i=1}^n$, which simultaneously determine:

(i) The expected therapeutic benefit $\mathbb{E}[b_i(x_i)|x_i \leq \hat{x}_i]$ used in coverage assessment

(ii) The maximum feasible price p_i via the coverage criterion

(iii) The realized demand $F_i(\hat{x}_i)$

This creates the fundamental trade-off that drives the manufacturer's decision: expanding the clinically eligible patient population increases demand $F_i(\hat{x}_i)$ but reduces the maximum feasible price through lower expected therapeutic benefit $\mathbb{E}[b_i(x_i)|x_i \leq \hat{x}_i]$. The manufacturer's optimal strategy must balance these competing forces across multiple indications, accounting for any interdependencies introduced through the coverage criterion $\mathcal{P}$.

3 Pricing Multi-Indication Drugs: Static Analysis

In this section, we assume that all n indications are known ex ante and that the manufacturer determines the clinically eligible patient population for each indication simultaneously.

We begin by analyzing three efficient mechanisms. By mechanisms, we mean a combination of pricing scheme (i.e., indication-based pricing, uniform pricing, or two-part tariffs) and a specific form of the right-hand side of the coverage criterion (2) (e.g., a single indication's expected therapeutic benefit or a weighted or unweighted average of expected therapeutic benefits). By efficient, we mean that the mechanism maximizes total therapeutic benefit across all n indications, implying that only patients with non-negative therapeutic benefit are included in the clinically eligible population of the corresponding indication and hence receive the drug. Although our definition of efficiency might be considered restrictive, it represents the minimum requirement we impose for a mechanism to be considered in the subsequent dynamic analysis of sequentially introduced indications.

We conclude the section with analyses of inefficient mechanisms, which provide a foundation for understanding why analogous inefficiencies might arise even with statically efficient mechanisms when indications are sequentially introduced.

3.1 Indication-Based Pricing

Under indication-based pricing, the manufacturer can charge a different price for each indication, subject to the constraint that prices cannot exceed the expected therapeutic benefit of the corresponding indication. The manufacturer's problem is

$$\max_{\{\hat{x}_i\}_{i=1}^n, \{p_i\}_{i=1}^n} \Pi = \sum_{i=1}^{n} p_i F_i(\hat{x}_i) \tag{3}$$

$$\text{subject to} \quad p_i \leq \mathbb{E}[v_i - \tau_i x_i | x_i \leq \hat{x}_i] \quad \forall i \tag{4}$$

At the optimum, constraint (4) binds for all indications. If the constraint were slack for any indication i, the manufacturer could increase p_i without affecting demand $F_i(\hat{x}_i)$, thereby increasing profit. Thus, optimal prices equal expected therapeutic benefits.[3]

Proposition 1 (Efficiency of Indication-Based Pricing)**.** *Under indication-based pricing with coverage determined by expected therapeutic benefit, the manufacturer:*

(i) Sets the clinically eligible patient population to include exactly those patients with non-negative therapeutic benefit: $\hat{x}_i^* = v_i/\tau_i$ for all i, thereby maximizing the total therapeutic benefit across all indications;

(ii) Extracts the total therapeutic benefit across all indications through indication-specific prices $p_i^* = \mathbb{E}[v_i - \tau_i x_i | x_i \leq v_i/\tau_i]$.

Proof. The expected therapeutic benefit can be expressed as total therapeutic benefit divided by the covered population:

$$\mathbb{E}[v_i - \tau_i x_i | x_i \leq \hat{x}_i] = \frac{\int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i}{F_i(\hat{x}_i)}. \tag{5}$$

Substituting the binding constraint (4) into the objective function (3):

$$\Pi = \sum_{i=1}^{n} \mathbb{E}[v_i - \tau_i x_i | x_i \leq \hat{x}_i] \cdot F_i(\hat{x}_i) \tag{6}$$

$$= \sum_{i=1}^{n} \int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i \tag{7}$$

The demand terms $F_i(\hat{x}_i)$ cancel out, directly transforming the manufacturer's profit function into the total therapeutic benefit across all indications.

[3] Analogous arguments establish that coverage constraints bind at the optimum throughout our analysis.

The first-order condition for indication i is

$$\frac{\partial \Pi}{\partial \hat{x}_i} = (v_i - \tau_i \hat{x}_i) f_i(\hat{x}_i) = 0, \tag{8}$$

yielding $\hat{x}_i^* = v_i/\tau_i$. This solution includes exactly those patients with non-negative therapeutic benefit, as $b_i(x_i) = v_i - \tau_i x_i \geq 0$ if and only if $x_i \leq v_i/\tau_i$.

Because the market for each of the n indications can be analyzed independently and each indication's contribution to profit, $\int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i$, is strictly quasi-concave in $\hat{x}_i$, the solution constitutes the unique global maximum. □

The efficiency result stems directly from the coverage criterion requiring that price equal expected therapeutic benefit. This criterion effectively "undoes" the averaging operation in the manufacturer's objective function: multiplying the expected therapeutic benefit by the population size recovers the total therapeutic benefit. As a result, profit maximization inherently aligns with total therapeutic benefit maximization.

3.2 Population-Weighted Uniform Pricing

Under uniform pricing, the manufacturer must charge a single price p across all indications. Consider a coverage criterion where this uniform price cannot exceed a weighted average of expected therapeutic benefits, with weights proportional to each indication's patient population:

$$\omega_i = \frac{F_i(\hat{x}_i)}{\sum_{j=1}^n F_j(\hat{x}_j)}. \tag{9}$$

The manufacturer's problem becomes:

$$\max_{\{\hat{x}_i\}_{i=1}^n, p} \Pi = p \sum_{i=1}^n F_i(\hat{x}_i) \tag{10}$$

$$\text{subject to} \quad p \leq \sum_{i=1}^n \omega_i \mathbb{E}[v_i - \tau_i x_i | x_i \leq \hat{x}_i] \tag{11}$$

As with indication-based pricing, the coverage constraint binds at the optimum.

Proposition 2 (Efficiency of Population-Weighted Uniform Pricing)**.** *Under uniform pricing with a population-weighted coverage criterion, the manufacturer:*

(i) Sets the clinically eligible patient population to include exactly those patients with non-negative therapeutic benefit: $\hat{x}_i^* = v_i/\tau_i$ for all i, thereby maximizing the total therapeutic benefit across all indications;

(ii) Extracts the total therapeutic benefit across all indications through a uniform price $p^* = \sum_{i=1}^{n} \omega_i^* \mathbb{E}[v_i - \tau_i x_i | x_i \leq v_i/\tau_i]$.

Proof. Substituting the binding constraint (11) into the objective function (10):

$$\Pi = \sum_{i=1}^{n} \omega_i \mathbb{E}[v_i - \tau_i x_i | x_i \leq \hat{x}_i] \cdot \sum_{j=1}^{n} F_j(\hat{x}_j) \tag{12}$$

$$= \sum_{i=1}^{n} \frac{F_i(\hat{x}_i)}{\sum_{j=1}^{n} F_j(\hat{x}_j)} \cdot \frac{\int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i}{F_i(\hat{x}_i)} \cdot \sum_{j=1}^{n} F_j(\hat{x}_j) \tag{13}$$

$$= \sum_{i=1}^{n} \int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i \tag{14}$$

The population weights and demand terms cancel out, transforming the manufacturer's profit into total therapeutic benefit across all indications—identical to expression (7) under indication-based pricing. Thus $\hat{x}_i^* = v_i/\tau_i$ for all i. Because the transformed profit function is separable across indications, with each indication's contribution strictly quasi-concave in $\hat{x}_i$, the solution constitutes the unique global maximum. □

The efficiency result again stems from how the coverage criterion "undoes" the averaging operation, but through a different mechanism than indication-based pricing. Here, the population weights ensure that each indication's contribution to the uniform price is exactly proportional to its contribution to total demand. This proportionality causes the weights and population terms to cancel in the profit function, recovering the total therapeutic benefit.

Despite the constraint of uniform pricing, the population-weighted criterion achieves the same efficient outcome as indication-based pricing: the manufacturer includes exactly those patients with non-negative therapeutic benefit.

3.3 Two-Part Tariffs

Under a two-part tariff, the manufacturer charges unit prices p_i for each indication and receives a lump-sum payment T from the health plan. The manufacturer's problem is

$$\max_{\{\hat{x}_i\}_{i=1}^n, \{p_i\}_{i=1}^n, T} \Pi = \sum_{i=1}^{n} p_i F_i(\hat{x}_i) + T \tag{15}$$

$$\text{subject to} \quad p_i \leq \overline{p}_i \quad \forall i \tag{16}$$

$$T \leq \sum_{i=1}^{n} \int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i - \sum_{i=1}^{n} p_i F_i(\hat{x}_i) \tag{17}$$

where $\{\overline{p}_i\}_{i=1}^n$ denote some unit-price thresholds, which potentially depend on $\{\hat{x}_i\}_{i=1}^n$.

Proposition 3 (Efficiency of Two-Part Tariffs)**.** *Under a two-part tariff, the manufacturer:*

(i) Sets the clinically eligible patient population to include exactly those patients with non-negative therapeutic benefit: $\hat{x}_i^* = v_i/\tau_i$ for all i, thereby maximizing the total therapeutic benefit across all indications;

(ii) Extracts the total therapeutic benefit through the combination of unit prices and lump-sum payment, regardless of the specific unit price values.

Proof. Setting up the Lagrangian:

$$\mathcal{L} = \sum_{i=1}^{n} p_i F_i(\hat{x}_i) + T - \sum_{i=1}^{n} \lambda_i (p_i - \overline{p}_i) - \mu \left[T - \sum_{i=1}^{n} \int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i + \sum_{i=1}^{n} p_i F_i(\hat{x}_i) \right], \quad (18)$$

where λ_i are the multipliers on the unit price constraints and μ is the multiplier on the lump-sum constraint.

The Kuhn-Tucker conditions are:

$$\frac{\partial \mathcal{L}}{\partial \hat{x}_i} = (1-\mu) p_i f_i(\hat{x}_i) + \lambda_i \frac{\partial \overline{p}_i}{\partial \hat{x}_i} + \sum_{j \neq i} \lambda_j \frac{\partial \overline{p}_j}{\partial \hat{x}_i} + \mu (v_i - \tau_i \hat{x}_i) f_i(\hat{x}_i) = 0 \quad \forall i \quad (19)$$

$$\frac{\partial \mathcal{L}}{\partial p_i} = (1-\mu) F_i(\hat{x}_i) - \lambda_i = 0 \quad \forall i \quad (20)$$

$$\frac{\partial \mathcal{L}}{\partial T} = 1 - \mu = 0 \quad (21)$$

with complementary slackness conditions:

$$\lambda_i \geq 0, \quad \lambda_i (p_i - \overline{p}_i) = 0 \quad \forall i \quad (22)$$

$$\mu \geq 0, \quad \mu \left[T - \sum_{i=1}^{n} \int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i + \sum_{i=1}^{n} p_i F_i(\hat{x}_i) \right] = 0 \quad (23)$$

From (21), we have $\mu^* = 1$, which implies the lump-sum constraint binds and the manufacturer extracts the health plan's entire surplus. Substituting into (20) yields $\lambda_i^* = 0$ for all i, indicating that unit price constraints do not affect the optimal clinically eligible populations—the efficiency result holds independently of the specific values of $\overline{p}_i$.

With $\mu^* = 1$ and $\lambda_i^* = 0$ for all i, condition (19) reduces to:

$$(v_i - \tau_i \hat{x}_i) f_i(\hat{x}_i) = 0 \quad \forall i, \quad (24)$$

yielding $\hat{x}_i^* = v_i/\tau_i$ for all i. The manufacturer thus includes exactly those patients with non-negative therapeutic benefit.

The binding lump-sum constraint implies:

$$T^* = \sum_{i=1}^{n} \int_0^{v_i/\tau_i} (v_i - \tau_i x_i) f_i(x_i) dx_i - \sum_{i=1}^{n} p_i^* F_i(v_i/\tau_i), \tag{25}$$

where unit prices p_i^* can take any values satisfying $p_i^* \leq \overline{p}_i$.

Finally, substituting the optimal multipliers $\mu^* = 1$ and $\lambda_i^* = 0$ into the Lagrangian:

$$\mathcal{L}(\{\hat{x}_i\}_{i=1}^n, \{p_i\}_{i=1}^n, T, \{\lambda_i^*\}_{i=1}^n, \mu^*) = \sum_{i=1}^{n} \int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i. \tag{26}$$

Because each indication's contribution $\int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i$ is strictly quasi-concave in $\hat{x}_i$ and $\hat{x}_i^* = v_i/\tau_i$ uniquely maximizes it, the solution constitutes the unique global maximum of the original problem. □

This proof reveals that the two-part tariff structure ensures the manufacturer maximizes total therapeutic benefit regardless of unit price constraints, as the lump-sum payment adjusts to extract the entire health plan's surplus.

The efficiency of two-part tariffs stems from their fundamental decoupling of revenue generation from the determination of clinically eligible patient populations $\{\hat{x}_i\}_{i=1}^n$. Unlike the previous efficient mechanisms, where efficiency requires specific coverage criteria to "undo" averaging operations, two-part tariffs achieve efficiency through the pricing structure itself—independent of whether cost-effectiveness coverage criteria (2) are used to establish unit price caps.

This represents a striking departure from indication-based pricing and population-weighted uniform pricing, where the coverage criteria must be carefully designed to transform the manufacturer's profit maximization into total benefit maximization. With a two-part tariff, the manufacturer's objective inherently becomes maximizing total therapeutic benefit across all indications (determination of clinically eligible patient populations), which it then fully extracts through the combination of unit prices and a correspondingly adjusted lump-sum payment (revenue generation).

Indeed, efficiency holds for any choice of unit prices between zero and the maximum feasible threshold (if any). Although setting unit prices to zero maximizes what can be extracted through the lump-sum payment, any combination of unit prices and lump-sum payment that extracts the total therapeutic benefit leads to the same efficient outcome. This occurs because unit prices and the lump-sum payment are perfect substitutes for revenue generation: any reduction in unit prices is exactly offset by an increase in the feasible lump-sum payment.

This pricing flexibility has gained real-world traction in pharmaceutical "subscription" or "Netflix" models, where payers make fixed payments for potentially unlimited access to medications.

These models, in their purest form, correspond to the case where unit prices equal zero and all revenue is generated through the lump-sum payment. Such two-part tariffs bear resemblance to innovation prizes that reward manufacturers for achieving specific outcomes. A lump-sum prize awarded for maximizing total therapeutic benefit across existing indications maps conceptually onto a two-part tariff with zero unit prices and a lump-sum payment equal to the total therapeutic benefit generated.

3.4 Inefficient Mechanisms

Having characterized three mechanisms that achieve efficiency, we now analyze two mechanisms that fail to maximize total therapeutic benefit: unweighted uniform pricing and anchor pricing. We examine these inefficient mechanisms for two reasons. First, both approaches have real-world traction. Unweighted uniform pricing is attractive for health systems requiring single prices across indications without sophisticated weighting schemes, and anchor pricing occurs when the price for all indications is determined by reference to a single indication's expected benefit—typically the first approved indication or the one with the highest or lowest expected therapeutic benefit. Second, understanding why these mechanisms fail to achieve efficiency, particularly in contrast to the efficient mechanisms analyzed above, provides insights for policy design. The analysis of anchor pricing provides a foundation for understanding how analogous inefficiencies can arise even with statically efficient mechanisms when indications are introduced sequentially rather than simultaneously. The distortions that emerge under anchor pricing in the static setting foreshadow the dynamic inefficiencies we analyze in Section 4, where the sequential nature of indication development allows for similar anchoring effects.

3.4.1 Unweighted Uniform Pricing

Consider a coverage criterion where the uniform price cannot exceed the unweighted average of expected therapeutic benefits across indications. The manufacturer's problem is

$$\max_{\{\hat{x}_i\}_{i=1}^n, p} \Pi = p \sum_{i=1}^{n} F_i(\hat{x}_i) \tag{27}$$

$$\text{subject to} \quad p \leq \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}[v_i - \tau_i x_i | x_i \leq \hat{x}_i] \tag{28}$$

As before, the coverage constraint binds.

Then, the first-order conditions for profit maximization are

$$pf_i(\hat{x}_i) + \frac{\partial p}{\partial \hat{x}_i}\left[\sum_{j=1}^{n} F_j(\hat{x}_j)\right] = 0 \quad \forall i. \tag{29}$$

The the marginal effect on price is

$$\frac{\partial p}{\partial \hat{x}_i} = \frac{1}{n} \cdot \frac{\partial \mathbb{E}[v_i - \tau_i x_i | x_i \leq \hat{x}_i]}{\partial \hat{x}_i} = -\frac{\tau_i f_i(\hat{x}_i)}{n F_i(\hat{x}_i)}\left[\hat{x}_i - \frac{1}{F_i(\hat{x}_i)}\int_0^{\hat{x}_i} x_i f_i(x_i) dx_i\right]. \tag{30}$$

Substituting into (29) and rearranging yields the n first-order conditions that implicitly define $\{\hat{x}_i^*\}_{i=1}^n$

$$\frac{\sum_{j=1}^n \mathbb{E}[v_j - \tau_j x_j | x_j \leq \hat{x}_j^*]}{\sum_{j=1}^n F_j(\hat{x}_j^*)} = \frac{\tau_i}{F_i(\hat{x}_i^*)}\left[\hat{x}_i^* - \frac{1}{F_i(\hat{x}_i^*)}\int_0^{\hat{x}_i^*} x_i f_i(x_i) dx_i\right] \quad \forall i. \tag{31}$$

These conditions reveal a crucial insight about uniform pricing: because the revenue from an additional patient is the same across all indications—captured by the left-hand side of (31)—the manufacturer optimally selects the clinically eligible patient populations such that the negative, indication-specific marginal price effects—captured by the right-hand side of (31)—are equalized across indications.

To see why this is generally inefficient, consider the case of indication-based pricing. Under indication-based pricing, revenue generation from an additional patient is also indication-specific because prices are indication-specific. For each indication, the manufacturer efficiently selects clinically eligible patient populations such that revenue generation and the indication-specific marginal effect on expected therapeutic benefit (i.e., price) balance. Under unweighted uniform pricing, and with heterogeneous indications, equalizing the marginal effects on expected therapeutic benefits across indications will generally lead to clinically eligible patient populations that differ from those under indication-based pricing.

Then, to understand the direction of this inefficiency, we further rewrite (31) as

$$\left[\frac{F_i(\hat{x}_i^*)}{\sum_{j=1}^n F_j(\hat{x}_j^*)}\right]\sum_{j=1}^{n} \mathbb{E}[v_j - \tau_j x_j | x_j \leq \hat{x}_j^*] = \tau_i\left[\hat{x}_i^* - \mathbb{E}[x_i | x_i \leq \hat{x}_i^*]\right] \quad \forall i, \tag{32}$$

compare it with the first-order condition under indication-based pricing[4]

$$\mathbb{E}[v_i - \tau_i x_i | x_i \leq \hat{x}_i^*] = \tau_i\left[\hat{x}_i^* - \mathbb{E}[x_i | x_i \leq \hat{x}_i^*]\right] \quad \forall i. \tag{33}$$

and impose the following regularity condition:

[4] This first-order conditions follows immediately from solving the manufacturer's profit-maximization problem without transforming the objective function into the total therapeutic benefit and using $\frac{\partial p}{\partial \hat{x}_i}$.

Assumption 1 (Monotonicity)**.** *For each indication i, the distribution F_i satisfies*

$$\frac{f_i(\hat{x}_i)}{F_i(\hat{x}_i)}\left[\hat{x}_i - \mathbb{E}[x_i|x_i \leq \hat{x}_i]\right] < 1 \quad \forall \hat{x}_i \in (0,1). \tag{34}$$

This condition ensures that $\tau_i[\hat{x}_i - \mathbb{E}[x_i|x_i \leq \hat{x}_i]]$ is increasing in $\hat{x}_i$ over the relevant range.

Under Assumption 1, the clinically eligible patient population for indication i will be inefficiently restricted (expanded) when its population share $F_i(\hat{x}_i^*)/\sum_{j=1}^n F_j(\hat{x}_j^*)$ is less than (exceeds) its share of therapeutic benefit $\mathbb{E}[v_i - \tau_i x_i|x_i \leq \hat{x}_i^*]/\sum_{j=1}^n \mathbb{E}[v_j - \tau_j x_j|x_j \leq \hat{x}_j^*]$. Indications with high expected therapeutic benefit but relatively small patient populations contribute disproportionately to raising the uniform price through the unweighted average, yet their restricted patient populations (at the margin) limit the manufacturer's ability to capture this value through volume. Conversely, indications with low expected therapeutic benefit but large patient populations (at the margin) are made available to inefficiently broad patient populations to increase revenue volume, even though this dilutes the average price. The manufacturer's optimal response amplifies these initial disparities—effectively "doubling down" on each indication's comparative advantage. High-value indications are further restricted to maximize their contribution to price, whereas high-volume indications are further expanded to maximize their contribution to quantity.

3.4.2 Anchor Pricing

Consider a coverage criterion where the uniform price cannot exceed the expected therapeutic benefit of a specific indication, which we call the anchor indication:

$$p \leq \mathbb{E}[v_a - \tau_a x_a|x_a \leq \hat{x}_a], \tag{35}$$

where $a \in \{1, 2, ..., n\}$ denotes the anchor indication.

Before proceeding, we must address which indications are covered under anchor pricing. One might consider selective coverage where only indications with expected benefits exceeding the uniform price are included. However, this approach leads to a fundamental problem. If the anchor is the indication with the highest expected benefit, only that indication would be covered, as all others have lower expected benefits than the price. If the anchor is the indication with the lowest expected benefit, all indications would be covered as they all exceed this minimum. For any other anchor, only indications with expected benefits at least equal to the anchor's would be covered, but among these, the binding constraint becomes the lowest expected benefit—making it equivalent to

anchoring on the minimum among covered indications. This reveals that any anchor pricing mechanism with selective coverage effectively collapses to having the minimum expected benefit (among covered indications) as the anchor. Meaningful anchor pricing therefore requires all indications to be covered at the uniform price, regardless of their individual expected benefits. We thus assume the health plan covers all approved indications at this uniform price.

The manufacturer's problem is

$$\max_{\{\hat{x}_i\}_{i=1}^n, p} \Pi = p \sum_{i=1}^{n} F_i(\hat{x}_i) \tag{36}$$

$$\text{subject to} \quad p \leq \mathbb{E}[v_a - \tau_a x_a | x_a \leq \hat{x}_a] \tag{37}$$

As before, constraint (37) binds at the optimum.

The first-order conditions are:

$$\frac{\partial p}{\partial \hat{x}_a} \sum_{j=1}^{n} F_j(\hat{x}_j) + p f_a(\hat{x}_a) = 0 \tag{38}$$

$$\frac{\partial p}{\partial \hat{x}_i} \sum_{j=1}^{n} F_j(\hat{x}_j) + p f_i(\hat{x}_i) = 0 \quad \forall i \neq a \tag{39}$$

Because only changes in $\hat{x}_a$ affect the price through constraint (37), we have $\frac{\partial p}{\partial \hat{x}_i} = 0$ for all $i \neq a$. Thus, condition (39) becomes:

$$p f_i(\hat{x}_i) = 0 \quad \forall i \neq a. \tag{40}$$

Given $p > 0$ and $f_i(\hat{x}_i) > 0$ for $\hat{x}_i \in [0, 1]$, this condition cannot hold for interior solutions. As the manufacturer's profit from indication $i \neq a$ is $pF_i(\hat{x}_i)$, which is increasing in $\hat{x}_i$ for fixed p, profit maximization requires:

$$\hat{x}_i^* = 1 \quad \forall i \neq a. \tag{41}$$

For the anchor indication, substituting $\frac{\partial p}{\partial \hat{x}_a} = -\frac{\tau_a f_a(\hat{x}_a)}{F_a(\hat{x}_a)} [\hat{x}_a - \mathbb{E}[x_a | x_a \leq \hat{x}_a]]$ into (38) and using $F_i(\hat{x}_i^*) = 1$ for all $i \neq a$:

$$\frac{(n-1)}{F_a(\hat{x}_a^*)} [\hat{x}_a^* - \mathbb{E}[x_a | x_a \leq \hat{x}_a^*]] + \hat{x}_a^* = \frac{v_a}{\tau_a}. \tag{42}$$

Because the left-hand side includes the positive term $\frac{(n-1)}{F_a(\hat{x}_a^*)} [\hat{x}_a^* - \mathbb{E}[x_a | x_a \leq \hat{x}_a^*]]$, we have $\hat{x}_a^* < v_a/\tau_a$, confirming that the anchor indication's clinically eligible population is inefficiently restricted.

The fundamental inefficiency of anchor pricing stems from the asymmetric role of indications in price determination. The anchor indication faces the usual price-quantity trade-off, as expanding its clinically eligible population reduces the expected therapeutic benefit and hence the uniform

price applied to all indications. Non-anchor indications, conversely, face no such trade-off; they are effectively irrelevant for price setting and contribute to profits only through quantity. This asymmetry leads the manufacturer to inefficiently restrict the anchor indication's clinically eligible patient population while maximizing coverage in all others.

The resulting allocation involves two types of inefficiency: patients with positive therapeutic benefits are excluded from the anchor indication, whereas all patients in non-anchor indications receive treatment regardless of their therapeutic match—including those with negative therapeutic benefit (i.e., $x_i > v_i/\tau_i$).

These results hold regardless of how the anchor indication is chosen. Whether the anchor has the highest expected benefit (leading to high prices and hence increased R&D funding, but restricted access in the high-value indication), the lowest expected benefit (leading to low prices, less R&D funding, and broad access), or is chosen for other reasons (e.g., previously introduced indication, as analyzed below), the same pattern emerges: inefficient restriction in the price-setting indication and full coverage elsewhere.

4 Pricing Multi-Indication Drugs: Dynamic Analysis

4.1 Pricing Flexibility and Dynamic Efficiency

We now examine whether the static efficiency results extend to a dynamic setting where indications are introduced sequentially over time. We consider a dynamic model where the manufacturer decides whether to introduce indication i at the beginning of period $t \in \{1, 2, ..., \mathcal{T}\}$. $\mathcal{T}$ denotes the number of periods that the monopoly profits associated with a mechanism are guaranteed (e.g., the manufacturer is shielded from competition or the agreement with the health plan holds). Once introduced, an indication remains available in all subsequent periods. Let I_t denote the set of indications available at the end of period t, with $I_0 = \emptyset$, and let $\delta \in (0, 1)$ denote the discount factor.

For any pricing mechanism, let Π_i denote the per-period profit from indication i given the current set of available indications and resulting prices. The present value of introducing indication i in period t is

$$\text{PV}_i^t = \sum_{s=t}^{\mathcal{T}} \delta^{s-t} \left[\sum_{j \in I_t} \Pi_j - \sum_{j \in I_{t-1}} \Pi_j \right], \tag{43}$$

where $I_t = I_{t-1} \cup \{i\}$. This captures the incremental profit stream from introducing indication i,

accounting for any effects on the profitability of previously introduced indications.

When prices can freely adjust to accommodate new indications, the three statically efficient mechanisms preserve their efficiency properties dynamically. To see this, recall from our static analysis that under each efficient mechanism, total profit equals total therapeutic benefit across all indications:

$$\sum_{j \in I_t} \Pi_j = \sum_{j \in I_t} \int_0^{\hat{x}_j} (v_j - \tau_j x_j) f_j(x_j) dx_j. \tag{44}$$

This relationship holds whether we have indication-based pricing with expected therapeutic benefit coverage criteria, population-weighted uniform pricing, or two-part tariffs. In each case, profit maximization aligns with total therapeutic benefit maximization.

In the dynamic setting, when the manufacturer introduces indication i in period t, taking $\{\hat{x}_j\}_{j \in I_{t-1}}$ as given, the present value becomes

$$\mathrm{PV}_i^t = \sum_{s=t}^{\mathcal{T}} \delta^{s-t} \left[\sum_{j \in I_t} \int_0^{\hat{x}_j} (v_j - \tau_j x_j) f_j(x_j) dx_j - \sum_{j \in I_{t-1}} \int_0^{\hat{x}_j} (v_j - \tau_j x_j) f_j(x_j) dx_j \right] \tag{45}$$

$$= \sum_{s=t}^{\mathcal{T}} \delta^{s-t} \int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i, \tag{46}$$

implying that

$$\arg\max_{\hat{x}_i} \mathrm{PV}_i^t = \arg\max_{\hat{x}_i} \sum_{s=t}^{\mathcal{T}} \delta^{s-t} \int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i = \frac{v_i}{\tau_i}. \tag{47}$$

Thus, in each period t, the manufacturer maximizes the incremental profit stream from introducing indication i by setting the clinically eligible patient population to include exactly those patients with non-negative therapeutic benefit: $\hat{x}_i^* = v_i/\tau_i$ for each newly introduced indication i. By doing so, with $\{\hat{x}_j^*\}_{j \in I_{t-1}}$ having been previously determined, the manufacturer maximizes the total profit stream from period t onwards.

Furthermore, such dynamic efficiency arises regardless of whether the manufacturer is myopic or forward-looking. Because price flexibility ensures that, in each t, the choice of $\hat{x}_i^*$ is independent of $\{\hat{x}_j^*\}_{j \in I_{t-1}}$, the manufacturer need not consider the impact of $\hat{x}_i^*$ on future prices nor on the clinically eligible populations of subsequent indications. Therefore, myopic and forward-looking behavior coincide: in each t, the manufacturer is free to maximize indication i's stream of total therapeutic benefit.

These results hold for all three efficient mechanisms.

Under indication-based pricing with expected therapeutic benefit coverage criteria, each indication's price depends only on its own expected therapeutic benefit, so introducing indication i

does not affect the prices or profits from previously introduced indications. The incremental profit comes solely from indication i's total therapeutic benefit.

Under population-weighted uniform pricing, introducing indication i changes the uniform price by incorporating its expected therapeutic benefit with appropriate population weight. The population-weighted structure ensures that despite affecting all indications' revenues, the incremental profit equals indication i's total therapeutic benefit.

Under two-part tariffs, the combination of unit prices and lump-sum payment adjusts to maintain the equality between total profit and total therapeutic benefit across all indications. Regardless of how unit prices are set (subject to any coverage criteria), the lump-sum payment adjusts so that incremental profit equals indication i's total therapeutic benefit.

The analysis extends naturally to the case where $\mathcal{T} \to \infty$, with δ incorporating both time discounting and the probability that the monopoly arrangement continues.

The key insight is that all three mechanisms preserve dynamic efficiency by ensuring that the incremental profit stream from introducing each indication equals its total therapeutic benefit stream. However, this dynamic efficiency depends on the ability of prices to adjust when new indications are introduced. In practice, regulatory constraints, contractual commitments, or political economy considerations often prevent upward price adjustments. For example, in six out of the nine countries analyzed by Mills and Kanavos (2023), upward price adjustments upon the launch of new indications where either prohibited or, despite theoretically permitted, not observed in practice.

4.2 Constraints on Price Setting and Dynamic Efficiency

The dynamic efficiency of indication-based pricing with expected therapeutic benefit coverage criteria and population-weighted uniform pricing depends on unconstrained price setting. More precisely, dynamic efficiency depends on the price of a new indication being independent of those of previously introduced indications in the former case and on allowing the population-weighted uniform price to adjust both upward and downward as new indications are introduced in the latter.

We focus on these two mechanisms because they are susceptible to constraints on unit price upward adjustments, commonly observed in practice. In contrast, two-part tariffs are robust to such constraints. Under the more innovative two-part tariffs, efficiency depends solely on the ability to adjust the lump-sum payment when new indications are introduced, not unit prices. Regardless of how these are set, provided that the lump-sum payment adjusts so that incremental profit equals incremental therapeutic benefit for each newly introduced indication, efficiency is preserved. In

the limiting case of a pure subscription model, the manufacturer eschews constraints on unit prices entirely by setting them equal to zero.

4.2.1 Constraints on Price Setting With a Myopic Manufacturer

Consider the case of a myopic manufacturer that evaluates each indication's introduction sequentially, disregarding implications for subsequently introduced indications. For clarity, we conduct the analysis focusing on the more computationally demanding case of population-weighted uniform pricing, then show how the results carry over to indication-based pricing (with expected therapeutic benefit coverage criteria).

Define p_{t-1} as the population-weighted uniform price in period $t-1$ and p_t as the fully adjusted population-weighted uniform price that would result from the introduction of indication i in period t if there were no constraints. Suppose further that p_{t-1} acts as price cap on p_t; i.e., $p_t \leq p_{t-1}$. For simplicity but without loss of generality, we assume that period t is the first period when the price cap is potentially binding.

When the constraint is not binding, the manufacturer's profit function in period t is

$$\Pi_t = p_t \left[F_i(\hat{x}_i) + \sum_{j \in I_{t-1}} F_j(\hat{x}_j) \right], \tag{48}$$

where p_t is the population-weighted average of expected therapeutic benefits across all indications in $I_{t-1} \cup \{i\}$. Again, this profit equals the sum of total therapeutic benefits across all indications. When maximizing with respect to $\hat{x}_i$, only the term corresponding to indication i varies, yielding

$$\arg\max_{\hat{x}_i} \left\{ \int_0^{\hat{x}_i} (v_i - \tau_i x_i) f_i(x_i) dx_i + \sum_{j \in I_{t-1}} \int_0^{\hat{x}_j} (v_j - \tau_j x_j) f_j(x_j) dx_j \right\} = \frac{v_i}{\tau_i}. \tag{49}$$

When the constraint binds, the manufacturer's profit function in period t is

$$\Pi_t = p_{t-1} \left[F_i(\hat{x}_i) + \sum_{j \in I_{t-1}} F_j(\hat{x}_j) \right]. \tag{50}$$

Because p_{t-1} equals the population-weighted average of expected therapeutic benefits for indications in I_{t-1}, the profit from previously introduced indications remains $\sum_{j \in I_{t-1}} \int_0^{\hat{x}_j} (v_j - \tau_j x_j) f_j(x_j) dx_j$. The manufacturer thus solves

$$\max_{\hat{x}_i} \left\{ p_{t-1} F_i(\hat{x}_i) + \sum_{j \in I_{t-1}} \int_0^{\hat{x}_j} (v_j - \tau_j x_j) f_j(x_j) dx_j \right\}. \tag{51}$$

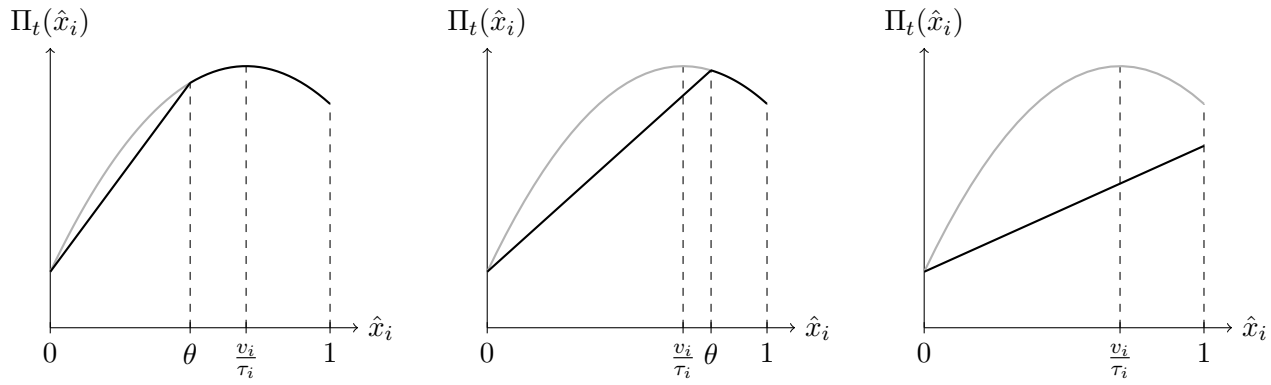

Figure 1: Illustrative examples of the myopic manufacturer's profits in period t as a function of indication i's clinically eligible patient population. The solid grey line represents $\sum_{j \in I_t} \int_0^{\hat{x}_j} (v_j - \tau_j x_j) f_j(x_j) dx_j$: profits and total therapeutic benefit achieved when prices adjust freely. The solid black line shows actual profits. When $\hat{x}_i \geq \theta$, profits and total therapeutic benefit coincide.

As only $F_i(\hat{x}_i)$ varies with $\hat{x}_i$ and is strictly increasing, this profit function is strictly increasing in $\hat{x}_i$, leading the manufacturer to expand the clinically eligible patient population as much as possible while the constraint binds.

Define threshold θ such that

$$\mathbb{E}(v_i - \tau_i x_i | x_i \leq \theta) = p_{t-1}. \tag{52}$$

Because $\mathbb{E}[v_i - \tau_i x_i | x_i \leq \hat{x}_i]$ is strictly decreasing in $\hat{x}_i$ and p_{t-1} is given, θ is unique when it exists in $[0, 1]$.. The constraint binds for $\hat{x}_i \leq \theta$; for $\hat{x}_i > \theta$, indication i's expected benefit falls below p_{t-1}, pulling the uniform price below the cap.

The manufacturer's profit function can therefore be rewritten as

$$\Pi_t = \begin{cases} \sum_{j \in I_{t-1}} \int_0^{\hat{x}_j} (v_j - \tau_j x_j) f_j(x_j) dx_j + F_i(\hat{x}_i) p_{t-1} & \text{if } 0 \leq \hat{x}_i \leq \theta \\ \sum_{j \in I_{t-1}} \int_0^{\hat{x}_j} (v_j - \tau_j x_j) f_j(x_j) dx_j + F_i(\hat{x}_i)\, \mathbb{E}(v_i - \tau_i x_i | x_i \leq \hat{x}_i) & \text{if } \theta < \hat{x}_i \leq 1 \end{cases}. \tag{53}$$

From (52), the profit function is continuous at $\hat{x}_i = \theta$, but exhibits a kink, as the derivative changes from strictly positive below the threshold to potentially negative above it. Note, additionally, that the "unconstrained" profit function branch lies above the "constrained" branch when $\hat{x}_i < \theta$ because $\mathbb{E}(v_i - \tau_i x_i | x_i \leq \theta) > p_{t-1}$.

Figure 1 presents illustrative examples for the graph of (53).

Having characterized the shape of the profit function when the population-weighted uniform price can only be revised downwards, we are now ready to derive the manufacturer's optimal choice of $\hat{x}_i$.

When $\theta \in [0, v_i/\tau_i]$, and due to the continuity of the profit function at $\hat{x}_i = \theta$, the profit function is strictly increasing on $[0, v_i/\tau_i)$ and strictly decreasing on $(v_i/\tau_i, 1]$. The profit function is strictly quasi-concave, and the manufacturer optimally sets $\hat{x}_i^* = v_i/\tau_i$. Efficiency is preserved, with the incremental profit from indication i equaling its total therapeutic benefit. This case is shown in the leftmost panel of Figure 1.

When $\theta \in (v_i/\tau_i, 1]$, and due to the continuity of the profit function at $\hat{x}_i = \theta$, the profit function is strictly increasing on $[0, \theta)$ and strictly decreasing on $(\theta, 1]$. The profit function is strictly quasi-concave, and the manufacturer optimally sets $\hat{x}_i^* = \theta$, inefficiently including patients with negative therapeutic benefit. This case is shown in the middle panel of Figure 1.

This solution naturally accounts for two special cases. When $\mathbb{E}(v_i - \tau_i x_i | x_i \leq \hat{x}_i) < p_{t-1} \; \forall \; \hat{x}_i \in [0, 1]$ (i.e., $\theta < 0$), the price cap never binds, and the manufacturer chooses the efficient coverage $\hat{x}_i^* = v_i/\tau_i$. When $\mathbb{E}(v_i - \tau_i x_i | x_i \leq \hat{x}_i) > p_{t-1} \; \forall \; \hat{x}_i \in [0, 1]$ (i.e., $\theta > 1$), the price constraint keeps the price at p_{t-1}. The manufacturer maximizes profit by setting $\hat{x}_i = 1$, including all patients regardless of therapeutic benefit. The latter case is shown in the rightmost panel of Figure 1.

Finally, consider how these results carry over to the analysis of indication-based pricing with expected therapeutic benefit coverage criteria. Suppose that in period t, the price of indication i cannot exceed the lowest price among the previously introduced indications. In this case, $p_{t-1} = \mathbb{E}(v_a - \tau_a x_a | x_a \leq v_a/\tau_a)$, with $\mathbb{E}(v_a - \tau_a x_a | x_a \leq v_a/\tau_a) < \mathbb{E}(v_j - \tau_j x_j | x_j \leq \hat{x}_j) \; \forall \; a, j \in I_{t-1}$, and is taken as given in period t. The manufacturer's profit function can be written once more as (53), and the choice of $\hat{x}_i^*$ follows accordingly.

We summarize the results for the two mechanisms in the following lemma.

Lemma 1. *Under indication-based pricing, if the price of existing indications acts as a cap on the price of a new indication, or under population-weighted uniform pricing, if the price cannot adjust upward, the myopic manufacturer may inefficiently expand the clinically eligible patient population of a new indication. When there exists a threshold $\theta > v_i/\tau_i$, below which an indication's expected therapeutic benefit exceeds the relevant price cap, the manufacturer chooses $\hat{x}_i^* = \min\{\theta, 1\}$, inefficiently including patients with negative therapeutic benefit. The manufacturer expands the clinically eligible population precisely to the point where the price cap is preserved, maximizing quantity while maintaining the highest feasible price.*

This analysis reveals how prohibiting price increases may fundamentally break the equivalence property that makes indication-based with expected therapeutic benefit coverage criteria and

population-weighted uniform pricing efficient. When prices cannot adjust upward, the equality between profit maximization and total therapeutic benefit maximization may no longer hold for new indications. In such cases, the manufacturer responds to the price cap by expanding coverage beyond the efficient level, including patients with negative therapeutic benefit to maximize revenue at the constrained price.

This distortion partly mirrors the inefficiency of anchor pricing analyzed in Section 3.4.2. While the price cap binds, indication i becomes effectively irrelevant for price setting and contributes to profits only through quantity, making it profitable for the manufacturer to expand the clinically eligible population to exactly the point where price (and profit) erosion would begin.

The potential inefficiency identified under myopic behavior raise a natural question: can forward-looking behavior mitigate these distortions? To address this, we now turn to the case of a forward-looking manufacturer with complete information about all future indications.

4.2.2 Constraints on Price Setting With a Forward-Looking Manufacturer

For clarity and tractability, we focus on a two-period case where the forward-looking manufacturer with complete information introduces one indication in each period. We label these as indication 1 in period 1 and indication 2 in period 2. Under both indication-based pricing (with expected therapeutic benefit coverage criteria) and population-weighted uniform pricing, the pricing constraint is that indication 2's price cannot exceed that of indication 1.

Whereas the myopic manufacturer treats the previous period's price as an exogenous cap, the forward-looking manufacturer recognizes it can strategically set this price *anchor* through its choice of $\hat{x}_1$. To see how, recall that the threshold θ is implicitly defined by $\mathbb{E}[v_2 - \tau_2 x_2 | x_2 \leq \theta] = p_1$ and that when $\theta \geq v_2/\tau_2$ the manufacturer optimally sets $\hat{x}_2^* = \theta$—expanding this indication's clinically eligible patient population to exactly where p_1 is preserved, as any further expansion would trigger price and profit erosion. The key insight is that the forward-looking manufacturer can manipulate where this price erosion occurs. It anticipates how its choice of $\hat{x}_1$ affects the choice of $\hat{x}_2$ through the pricing constraint; more precisely, how its choice of $\hat{x}_1$ affects θ. Thus, the threshold θ becomes a function of $\hat{x}_1$ and, because $p_1 = \mathbb{E}[v_1 - \tau_1 x_1 | x_1 \leq \hat{x}_1]$ under both mechanisms, it is implicitly defined by

$$\mathbb{E}[v_2 - \tau_2 x_2 | x_2 \leq \theta(\hat{x}_1)] = \mathbb{E}[v_1 - \tau_1 x_1 | x_1 \leq \hat{x}_1]. \tag{54}$$

For $\hat{x}_2 \leq \theta(\hat{x}_1)$, p_1 is preserved; for $\hat{x}_2 > \theta(\hat{x}_1)$, indication 2's expected therapeutic benefit falls below p_1, pulling the price below the anchor p_1.

Applying the implicit function theorem yields

$$\frac{\partial\theta(\hat{x}_1)}{\partial\hat{x}_1} = \frac{\tau_1 f_1(\hat{x}_1)[\hat{x}_1 - \mathbb{E}[x_1|x_1 \leq \hat{x}_1]]/F_1(\hat{x}_1)}{\tau_2 f_2(\theta)[\theta - \mathbb{E}[x_2|x_2 \leq \theta]]/F_2(\theta)} > 0, \tag{55}$$

which confirms that restricting $\hat{x}_1$ reduces $\theta(\hat{x}_1)$. As a lower θ corresponds to a higher binding price for indication 2, the forward-looking manufacturer can effectively anchor indication 2's price at a higher value through strategic restriction of the clinically eligible population for indication 1.

Define $\hat{x}_1^\theta$ as the threshold satisfying $\theta(\hat{x}_1^\theta) = v_2/\tau_2$. Given the monotonicity of $\theta(\hat{x}_1)$, this threshold is unique when it exists in $[0, 1]$. It divides the forward-looking manufacturer's problem into two regions. For $\hat{x}_1 \leq \hat{x}_1^\theta$, we have $\theta(\hat{x}_1) \leq v_2/\tau_2$, making efficient patient selection profit-maximizing for indication 2. For $\hat{x}_1 > \hat{x}_1^\theta$, we have $\theta(\hat{x}_1) > v_2/\tau_2$, leading to over-inclusion in indication 2.

Under both indication-based and population-weighted uniform pricing, the forward-looking manufacturer's problem may be written as

$$\begin{aligned}
\max_{\hat{x}_1} \Pi &= \mathbb{E}[v_1 - \tau_1 x_1 | x_1 \leq \hat{x}_1] \cdot F_1(\hat{x}_1) \\
&\quad + \delta\Big[\mathbb{E}[v_1 - \tau_1 x_1 | x_1 \leq \hat{x}_1] \cdot F_1(\hat{x}_1) + \mathbb{E}[v_2 - \tau_2 x_2 | x_2 \leq \hat{x}_2^*(\hat{x}_1)] \cdot F_2(\hat{x}_2^*(\hat{x}_1))\Big] \\
&= \int_0^{\hat{x}_1} (v_1 - \tau_1 x_1) f_1(x_1) dx_1 + \delta\left[\int_0^{\hat{x}_1} (v_1 - \tau_1 x_1) f_1(x_1) dx_1 + \int_0^{\hat{x}_2^*(\hat{x}_1)} (v_2 - \tau_2 x_2) f_2(x_2) dx_2\right],
\end{aligned} \tag{56}$$

where $\hat{x}_2^*(\hat{x}_1) = \max\{v_2/\tau_2, \theta(\hat{x}_1)\}$.

Efficient Clinically Eligible Patient Populations

When $\hat{x}_1^\theta \geq v_1/\tau_1$, the constraint never forces inefficient over-inclusion in indication 2 for any feasible $\hat{x}_1 \in [0, v_1/\tau_1]$. The manufacturer optimally chooses $\hat{x}_2^* = v_2/\tau_2$, and it immediately follows that

$$\arg\max_{\hat{x}_1} \left\{(1+\delta)\int_0^{\hat{x}_1} (v_1 - \tau_1 x_1) f_1(x_1) dx_1 + \delta \int_0^{v_2/\tau_2} (v_2 - \tau_2 x_2) f_2(x_2) dx_2\right\} = \frac{v_1}{\tau_1}. \tag{57}$$

The efficient outcome is achieved in both indications. In this case, the manufacturer has no incentive to inefficiently under-include patients in indication 1 to anchor the price at a higher value. The profit-maximizing patient population for indication 2 will be efficiently set at v_2/τ_2, as the manufacturer finds it profitable to expand indication 2's population from $\theta(v_1/\tau_1)$ to v_2/τ_2, accepting the associated price reduction, because the gain from efficient demand expansion dominates the revenue loss from the lower price.

Figure 2 illustrates this outcome.

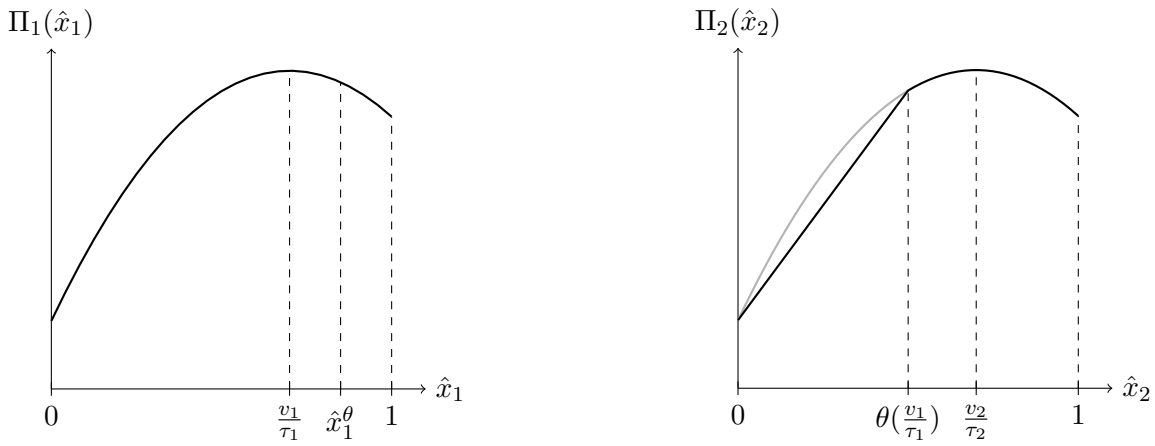

Figure 2: Illustrative examples of the forward-looking manufacturer's instantaneous profits from each indication as functions of their clinically eligible patient populations when efficient outcomes are intertemporally profit-maximizing. For indication 1, profits equal total therapeutic benefit at all $\hat{x}_1$. For indication 2, the solid grey line represents profits and total therapeutic benefit achieved when prices adjust freely, whereas the solid black line shows actual profits. When $\hat{x}_2 \geq \theta(\hat{x}_1)$, profits and total therapeutic benefit coincide.

Inefficient Clinically Eligible Patient Populations

When $\hat{x}_1^\theta < v_1/\tau_1$, choosing $\hat{x}_1 \geq \hat{x}_1^\theta$ implies $\hat{x}_2 = \theta(\hat{x}_1) \geq v_2/\tau_2$. From (56), the first-order condition for profit maximization is

$$(1+\delta)(v_1 - \tau_1\hat{x}_1)f_1(\hat{x}_1) + \delta[v_2 - \tau_2\theta(\hat{x}_1)]f_2[\theta(\hat{x}_1)]\frac{\partial\theta(\hat{x}_1)}{\partial\hat{x}_1} = 0. \tag{58}$$

The optimal $\hat{x}_1^*$ must lie in the interval $(\hat{x}_1^\theta, v_1/\tau_1)$. In this case, the manufacturer anticipates that the profit-maximizing patient population for indication 2 will be inefficiently expanded beyond v_2/τ_2. Within the interval $[0, \theta(\hat{x}_1)]$, indication 2 contributes to profits only through quantity—exactly as in the static anchor pricing case analyzed in Section 3.4.2—making this expanded population more profitable when served at a higher price. The manufacturer thus finds it profitable to sacrifice demand in indication 1 to anchor this higher price, as the revenue gain from serving indication 2's expanded population at the elevated price dominates the loss from under-inclusion in indication 1. It then expands $\hat{x}_2$ to exactly $\theta(\hat{x}_1^*)$, maximizing quantity while preserving the anchor price.

To see why $\hat{x}_1^* \in (\hat{x}_1^\theta, v_1/\tau_1)$, consider the following steps.

For $\hat{x}_1 < \hat{x}_1^\theta$, the manufacturer sacrifices profits in indication 1 (moving away from its maximum at v_1/τ_1) whereas indication 2 already achieves its unconstrained maximum. This is strictly dominated by $\hat{x}_1 = \hat{x}_1^\theta$. However, at $\hat{x}_1 = \hat{x}_1^\theta$, which implies $\theta(\hat{x}_1) = v_2/\tau_2$, the first-order condition

becomes

$$(1+\delta)(v_1 - \tau_1 \hat{x}_1^\theta) f_1(\hat{x}_1^\theta) > 0. \tag{59}$$

Therefore, $\hat{x}_1^* > \hat{x}_1^\theta$.

For $\hat{x}_1 > v_1/\tau_1$, the manufacturer forgoes profits in both indications 1 (moving past its maximum) and 2 (as $\theta(\hat{x}_1)$ increases further beyond v_2/τ_2). This is strictly dominated by $\hat{x}_1 = v_1/\tau_1$. At $\hat{x}_1 = v_1/\tau_1$, which implies $\theta(\hat{x}_1) > v_2/\tau_2$, the first-order condition becomes

$$\delta[v_2 - \tau_2 \theta(v_1/\tau_1)] f_2[\theta(v_1/\tau_1)] \frac{\partial \theta(v_1/\tau_1)}{\partial \hat{x}_1} < 0. \tag{60}$$

Therefore, $\hat{x}_1^* < v_1/\tau_1$.

Let $H(\hat{x}_1)$ denote the left-hand side of (58). We have established that $H(\hat{x}_1^\theta) > 0$ and $H(v_1/\tau_1) < 0$. Given the continuity of $H(\hat{x}_1)$ and by the intermediate value theorem, the manufacturer's optimal choice exists and satisfies $\hat{x}_1^* \in (\hat{x}_1^\theta, v_1/\tau_1)$. If $H'(\hat{x}_1) < 0$, which is economically plausible, $H(\hat{x}_1)$ is monotonically decreasing, and the solution is unique.

This analysis extends to the special case where $\hat{x}_1^\theta \leq 0$. When $\hat{x}_1^\theta = 0$, we have $\theta(0) = v_2/\tau_2$, and (58) at $\hat{x}_1 = 0$ yields $(1+\delta)v_1 f_1(0) > 0$, confirming that $0 < \hat{x}_1^* < v_1/\tau_1$. When $\hat{x}_1^\theta < 0$, even setting $\hat{x}_1 = 0$ forces over-inclusion in indication 2. If the left-hand side of (58) at $\hat{x}_1 = 0$ is positive, an interior solution exists; otherwise, $\hat{x}_1^* = 0$, though this corner solution is less economically relevant as it implies the manufacturer would forgo indication 1 (i.e., zero demand).

Figure 3 illustrates the solution $\hat{x}_1^* \in (\hat{x}_1^\theta, v_1/\tau_1)$ and $\hat{x}_2^*(\hat{x}_1^*) = \theta(\hat{x}_1^*) > v_2/\tau_2$.

Within the interval $(\hat{x}_1^\theta, v_1/\tau_1)$, the manufacturer faces a price-quantity trade-off. Reducing $\hat{x}_1$ below v_1/τ_1 lowers profits from indication 1 due to lower demand for this indication, but raises profits from indication 2 by anchoring the price at higher value, enabling the manufacturer to serve the inefficiently expanded patient population in indication 2 at this elevated price. The optimum balances these opposing effects.

From direct inspection of (58) and the equivalence between manufacturer profits and total therapeutic benefit at $(\hat{x}_1^*, \hat{x}_2^*)$, this trade-off can be alternatively interpreted in terms of therapeutic benefit. Reducing $\hat{x}_1$ below v_1/τ_1 lowers profits from indication 1 by including fewer patients with positive therapeutic benefit. This effect is captured by the first term on the left-hand side of (58). However, it also raises profits from indication 2 by allowing the manufacturer to profitably exclude more patients with negative therapeutic benefit. This effect is captured by the second term on the left-hand side of (58).

We summarize the results in the following lemma.

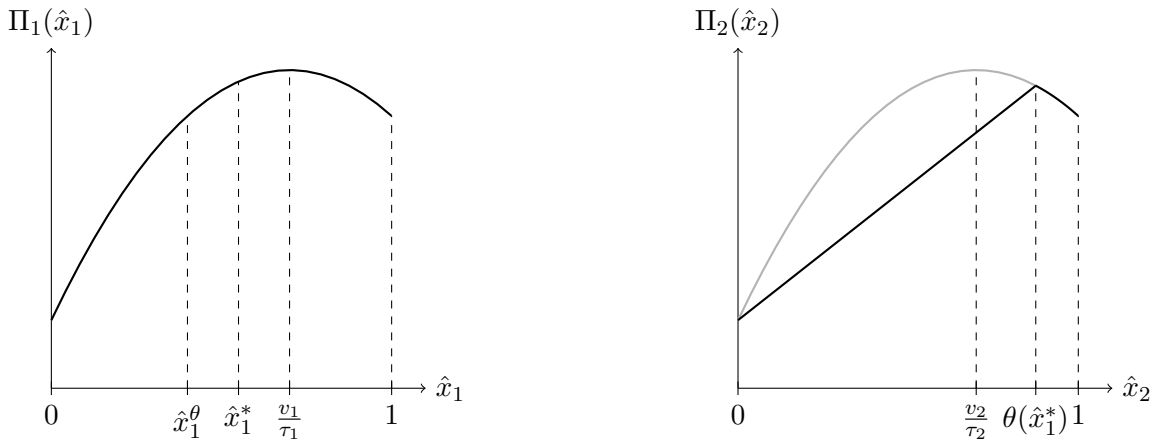

Figure 3: Illustrative examples of the forward-looking manufacturer's instantaneous profits from each indication as functions of their clinically eligible patient populations when inefficient outcomes are intertemporally profit-maximizing. For indication 1, profits equal total therapeutic benefit at all $\hat{x}_1$. For indication 2, the solid grey line represents profits and total therapeutic benefit achieved when prices adjust freely, whereas the solid black line shows actual profits. When $\hat{x}_2 \geq \theta(\hat{x}_1)$, profits and total therapeutic benefit coincide.

Lemma 2. *Under indication-based pricing, if the price of a new indication cannot exceed the price of existing indications, or under population-weighted uniform pricing, if the price cannot adjust upward, the forward-looking manufacturer with complete information may inefficiently choose the clinically eligible patient populations of two indications. When there exists a threshold $\hat{x}_1^\theta < v_1/\tau_1$ such that choosing $\hat{x}_1 > \hat{x}_1^\theta$ leads to profitable over-inclusion in indication 2:*

1. *Both indications have inefficient patient selection: under-inclusion in indication 1 ($\hat{x}_1^* < v_1/\tau_1$) and over-inclusion in indication 2 ($\hat{x}_2^* > v_2/\tau_2$).*

2. *The forward-looking manufacturer strategically under-includes patients in indication 1 to anchor the price at a higher value, then expands indication 2's clinically eligible patient population to exactly where this anchor price is preserved.*

3. *Despite creating inefficiency in both indications, forward-looking behavior leads to higher total therapeutic benefit across indications than myopic behavior.*

As with myopic behavior, the dual inefficiency mirrors that of anchor pricing analyzed in Section 3.4.2: under-inclusion in the price-setting indication to create a high price anchor, and over-inclusion elsewhere to maximize revenues at this anchor price.

Importantly, the improvement in total therapeutic benefit under forward-looking behavior fol-

lows from revealed preference and the equivalence between manufacturer profits and total therapeutic benefit at $(\hat{x}_1^*, \hat{x}_2^*)$ under both types of behavior. The forward-looking manufacturer could choose the myopic solution ($\hat{x}_1^* = v_1/\tau_1$ and $\hat{x}_2^* = \theta(v_1/\tau_1)$) but opts for different clinically eligible patient populations that yield higher discounted profits—and, consequently, total therapeutic benefit across indications.

We summarize the results on dynamic efficiency, manufacturer behavior, and constraints on price setting in the following proposition.

Proposition 4. *Under indication-based pricing with expected therapeutic benefit coverage criteria, population-weighted uniform pricing, and two-part tariffs—provided that upward payment adjustments are allowed as indications are sequentially introduced—the manufacturer efficiently selects the clinically eligible patient population for each indication, thereby including exactly the patients with non-negative therapeutic benefit. This efficient outcome arises regardless of whether the manufacturer is myopic or forward-looking.*

Under indication-based pricing with expected therapeutic benefit coverage criteria, if the price of a new indication cannot exceed the price of existing indications, or under population-weighted uniform pricing, if the price cannot adjust upward, inefficiencies might arise that cannot be fully eliminated by forward-looking behavior. Whereas the myopic manufacturer may inefficiently expand the clinically eligible patient population for a new indication to the point where the price cap binds, the forward-looking manufacturer may also strategically restrict the clinically eligible patient population for an earlier indication to set a higher price anchor based on its increased expected therapeutic benefit, then expand the subsequent indication's population to exactly where this anchor price is preserved. However, despite creating dual inefficiencies, forward-looking behavior improves the total therapeutic benefit across indications.

Proof. Follows directly from Section 4.1, Lemma 1, and Lemma 2. □

5 Discussion and Concluding Remarks

This article explores the pricing of multi-indication drugs by developing a model that connects a pharmaceutical manufacturer's choice of clinically eligible patient populations for each indication to drug prices. Faced with a health plan that offers the drug to all clinically eligible patients, the manufacturer chooses patient populations strategically, simultaneously determining realized demands and the expected therapeutic benefits used in coverage and pricing decisions by the health

plan. Therefore, our model provides a theoretical foundation for manufacturer decisions observed empirically, revealing they constitute strategic responses to regulation within the context of drug development and pricing.

Defining efficiency as setting the clinically eligible patient population for each indication to include exactly those patients with non-negative therapeutic benefit, we identify three efficient mechanisms when indications are simultaneously introduced: (i) indication-based pricing with expected therapeutic benefit coverage criteria, (ii) population-weighted uniform pricing, and (iii) two-part tariffs, regardless of the coverage and pricing criteria adopted by the health plan. Although all three mechanisms achieve efficiency through perfect incentive alignment—i.e., by equating profit maximization with total therapeutic benefit maximization—they do so in conceptually different manners.

The first two mechanisms effectively *undo* the averaging operation in the manufacturer's profit function. Under indication-based prices equal to the indication's expected therapeutic benefit, multiplying the expected (average) therapeutic benefit by the patient population size recovers the total therapeutic benefit for each indication. Although the single price requires population weights to cancel out patient populations, the same recovering of the total therapeutic benefit across all indications occurs under population-weighted uniform pricing. Two-part tariffs achieve efficiency by *decoupling* the determination of clinically eligible patient populations from revenue generation. Under the other two efficient mechanisms, choosing clinically eligible patient populations affects unit prices and hence revenue generation. With two-part tariffs, conversely, the determination of clinically eligible patient populations aims uniquely at maximizing total therapeutic benefit, whereas revenue generation is left for the lump-sum payment, independently of how unit prices are set.

Due to their real-world traction and despite their being inefficient, we also look at unweighted uniform pricing and anchor pricing. The former *doubles down* on each indication's initial attributes, restricting high-value indications and expanding high-volume ones. The latter makes all but one indication *irrelevant for price setting* and contribute to profits only through quantity.

We also show that the three efficient mechanisms continue to yield the efficient outcome in the arguably more realistic case where indications are sequentially introduced, provided there is price flexibility; i.e., prices (or lump-sum payments) are allowed to adjust both downward and upward.

This article makes a conceptual contribution, and our results offer several policy implications.

First, our model extends beyond multi-indication drugs to inform the broader understanding

of regulated monopoly markets in healthcare. In our model, the health plan purchases and offers the drug to all patients in the clinically eligible population, and although this could be seen as modeling price-insensitive consumers, we are effectively recasting the standard monopoly problem within pharmaceutical markets. The HTA-based coverage criteria create coverage-mediated demand curves whereby the relationship between patient population sizes and prices is mediated through expected therapeutic benefit rather than consumer willingness-to-pay. By setting clinically eligible populations through late-stage R&D (i.e., phase III clinical trials), the manufacturer chooses its position (a price-quantity pair) along these demand curves. Moreover, this framework shows that regulation specifics matter: different coverage criteria (e.g., weighted vs unweighted average) yield different demand curves, explaining why they lead to different outcomes and revealing why two-part pricing allows the manufacturer to bypass such criteria entirely. In sum, price regulation does not eliminate monopoly behavior within pharmaceutical markets; it channels it. This leads to our key policy insight that proper design of mechanisms can align profit maximization with therapeutic benefit maximization.

Second, our results provide policymakers with a menu of efficient mechanisms, each suited to different institutional constraints and political economies of drug pricing. Health systems unable to implement indication-based pricing (e.g., because price transparency rules preclude confidential discounts) might achieve identical efficiency through population-weighted uniform pricing. Systems where drug prices are politically sensitive might achieve efficiency through subscription-based models that decouple access from per-unit costs. Indeed, upward adjustments to fixed payments when new patient populations gain access might face fewer institutional barriers and political pushback than per-unit price increases, as they could potentially be justified by the expanded treatment population rather than appearing as price inflation. Systems that lack formal HTA processes can still achieve efficiency through two-part tariffs, whereas the other two efficient mechanisms require specific coverage and pricing criteria that depend on therapeutic benefit assessment.

Third, a critical property of these three mechanisms is that manufacturer profits exactly equal the total therapeutic benefit generated, even when price constraints force inefficient patient selection. This ensures that payers receive value commensurate with their payments, as manufacturers cannot extract profits beyond the therapeutic benefit they create. Although the political economy of allowing manufacturers to capture the full value of their innovations remains complex, this equivalence provides a theoretical foundation for value-based pharmaceutical pricing.

Fourth, our results demonstrate that pricing flexibility ensures that the three efficient mech-

anisms remain efficient when indications are sequentially introduced. However challenging the political economy of price increases might be, such flexibility yields a policy-relevant property: it renders the sophistication of manufacturer behavior irrelevant. In the absence of price constraints, the manufacturer maximizes the profit stream from each indication simply by maximizing that indication's stream of therapeutic benefit—and it is free to pursue this objective without concerns about the impact on future prices. Conversely, constraints on price setting create potential inefficiencies in patient selection that even forward-looking behavior fails to fully eliminate. Policymakers should therefore select one of the efficient mechanisms and let it function unimpeded.

Fifth, our analysis provides a theoretical foundation for understanding the empirically documented pattern of high expected therapeutic benefit, narrower patient population first indications, followed by lower benefit, wider patient population indications. Michaeli et al. (2022) report that average incremental quality-adjusted life years (QALY) and life years gained were significantly lower, and patient populations larger, for the second and third launched indications in the United States, Germany, France, England, Canada, Australia, and Scotland. Importantly, this pattern held in countries that resorted to weighted-average pricing or differential discounts (a form of indication-based pricing), where no upward price revisions were observed. This pattern maps exactly onto our result that price caps lead to under-inclusion in the first indication and over-inclusion in a subsequent indication under both indication-based pricing (with expected therapeutic benefit coverage criteria) and population-weighted uniform pricing, when the manufacturer is forward-looking. In a note to policymakers, Michaeli et al. (2022) write that manufacturers maximize revenues "by sequencing indication launches to set the highest possible drug price: cancer drugs are first launched for rare diseases that offer significant QALY gains and then extended to indications that deliver lower QALY gains to more eligible patients."[5] Our analysis reveals that the above-described pattern may reflect intertemporal optimization whereby manufacturers strategically balance inefficiencies across indications rather than simple maximization of the first indication's therapeutic benefit. Forward-looking manufacturers respond to the ability to endogenously set future price anchors through strategic under-inclusion in the first indication, thereby increasing its expected therapeutic benefit and setting a high price anchor. They then pursue over-inclusion in subsequently launched indications to maximize revenues at this anchor price, expanding patient populations to exactly where the anchor price is preserved.

Finally, our analysis also has implications for upstream R&D decisions and pharmaceutical

[5] This is explicitly advocated by industry consultants. See the "narrow first strategy" in Gores and Scott (2023).

innovation more broadly. We have shown that when price caps prevent upward adjustments, they may force manufacturers into inefficient patient selection, reducing total profits from multi-indication drugs below their unconstrained potential. This profit reduction alters the relative attractiveness of two innovation strategies. Consider a manufacturer choosing between developing an entirely new drug or repurposing an existing one for additional indications. The new drug enters unconstrained and can optimize both price and patient population. The new indication, however, inherits pricing constraints that force the inefficient patient selections we have analyzed. This may lead to a reversal in R&D priorities: cases where repurposing would be more profitable under flexible pricing may see new drug development under price caps. Whether this shift enhances welfare—by spurring novel innovation—or reduces it—by discouraging efficient repurposing that leverages existing safety data—remains an open question. What is clear is that pricing mechanisms for multi-indication drugs have consequences beyond their immediate market effects, potentially reshaping the direction of pharmaceutical innovation.

References

Adida, E. (2024). Indication-based pricing for multi-indication drugs. *Management Science*, *70*(11), 7506–7523. https://doi.org/10.1287/mnsc.2022.01721

Brekke, K. R., Dalen, D. M., & Straume, O. R. (2025). *Competition matters: Uniform vs. indication-based pricing of pharmaceuticals* (Discussion Paper No. 01/2025). NHH Dept. of Economics. https://doi.org/10.2139/ssrn.5087646

Campillo-Artero, C., Puig-Junoy, J., Segú-Tolsa, J. L., & Trapero-Bertran, M. (2020). Price models for multi-indication drugs: A systematic review. *Applied Health Economics and Health Policy*, *18*(1), 47–56. https://doi.org/10.1007/s40258-019-00517-z

Chandra, A., & Garthwaite, C. (2017). The economics of indication-based drug pricing. *New England Journal of Medicine*, *377*(2), 103–106. https://doi.org/10.1056/NEJMp1705035

Cherubini, A., Oristrell, J., Pla, X., Ruggiero, C., Ferretti, R., Diestre, G., ... Mills, G. H. (2011). The persistent exclusion of older patients from ongoing clinical trials regarding heart failure. *Archives of Internal Medicine*, *171*(6), 550–556. https://doi.org/10.1001/archinternmed.2011.31

Cole, A., Neri, M., & Cookson, G. (2021). *Payment models for multi-indication therapies* (OHE Consulting Report). Office of Health Economics. https://www.ohe.org/

wp-content/uploads/2021/11/OHE-Report_Cole-et-al.-2021_Payment-Models-for-Multi-indication_Therapies-.pdf

Duma, N., Kothadia, S. M., Azam, T. U., Yadav, S., Paludo, J., Vera Aguilera, J., ... Adjei, A. A. (2019). Characterization of comorbidities limiting the recruitment of patients in early phase clinical trials. *The Oncologist*, *24*(1), 96–102. https://doi.org/10.1634/theoncologist.2017-0687

Goldhaber-Fiebert, J. D., & Cipriano, L. E. (2023). Pricing treatments cost-effectively when they have multiple indications: Not just a simple threshold analysis. *Medical Decision Making*, *43*(7–8), 914–929. https://doi.org/10.1177/0272989X231197772

Gores, M., & Scott, K. (2023). *Success multiplied: Launch excellence for multi-indication assets* (White Paper). IQVIA. https://www.iqvia.com/library/white-papers/success-multiplied-launch-excellence-for-multi-indication-assets

Heine, R. J., Mathijssen, R. H., Verbeek, F. A., Van Gils, C., & Uyl-de Groot, C. A. (2024). Market entry agreements for innovative pharmaceuticals subject to indication broadening: A case study for Pembrolizumab in the Netherlands. *Value in Health*, *27*(10), 1367–1372. https://doi.org/10.1016/j.jval.2024.06.003

Hill, N. S., Preston, I. R., & Roberts, K. E. (2008). Patients with pulmonary arterial hypertension in clinical trials: Who are they? *Proceedings of the American Thoracic Society*, *5*(5), 603–609. https://doi.org/10.1513/pats.200803-032SK

ISPOR Europe. (2025). *Pricing and reimbursement of multiple indication medicines: Can a balance be found between different stakeholder perspectives to optimise value and access for patients, while ensuring sustainable and affordable innovation?* (Panel Report). Dolon. https://dolon.com/rare-knowledge/ispor-europe-panel-report-pricing-and-reimbursement-of-multiple-indication-medicines-can-a-balance-be-found-between-different-stakeholder-perspectives-to-optimise-value-and-access-for-patients-whil

Jiang, Y., Li, M., Jiang, S., Si, L., & Gu, Y. (2024). Patient welfare implications of indication-specific value-based pricing of multi-indication drugs. *Value in Health*, *27*(3), 273–277. https://doi.org/10.1016/j.jval.2023.11.008

Jin, S., Pazdur, R., & Sridhara, R. (2017). Re-evaluating eligibility criteria for oncology clinical trials: Analysis of investigational new drug applications in 2015. *Journal of Clinical Oncology*, *35*(33), 3745–3752. https://doi.org/10.1200/JCO.2017.73.4186

Khullar, D. (2025, July 15). Can A.I. find cures for untreatable diseases—using drugs we already

have? *The New Yorker*. https://www.newyorker.com/culture/open-questions/can-ai-find-cures-for-untreatable-diseases-using-drugs-we-already-have

Lundh, A., Lexchin, J., Mintzes, B., Schroll, J. B., & Bero, L. (2017). Industry sponsorship and research outcome. *Cochrane Database of Systematic Reviews*, (2). https://doi.org/10.1002/14651858.MR000033.pub3

Michaeli, D. T., Mills, M., & Kanavos, P. (2022). Value and price of multi-indication cancer drugs in the USA, Germany, France, England, Canada, Australia, and Scotland. *Applied Health Economics and Health Policy*, *20*(5), 757–768. https://doi.org/10.1007/s40258-022-00737-w

Mills, M., & Kanavos, P. (2023). Healthcare payer perspectives on the assessment and pricing of oncology multi-indication products: Evidence from nine OECD countries. *PharmacoEconomics Open*, *7*(4), 553–565. https://doi.org/10.1007/s41669-023-00406-1

Mills, M., Michaeli, D., Miracolo, A., & Kanavos, P. (2023). Launch sequencing of pharmaceuticals with multiple therapeutic indications: Evidence from seven countries. *BMC Health Services Research*, *23*(1), 150. https://doi.org/10.1186/s12913-023-09095-2

Nordon, C., Sanchez, B., Zhang, M., Wang, X., Hunt, P., Belger, M., ... GetReal Initiative (2023). Testing the "RCT augmentation" methodology: A trial simulation study to guide the broadening of trials eligibility criteria and inform on effectiveness. *Contemporary Clinical Trials Communications*, *33*, 101142. https://doi.org/10.1016/j.conctc.2023.101142

Preckler, V., & Espín, J. (2022). The role of indication-based pricing in future pricing and reimbursement policies: A systematic review. *Value in Health*, *25*(4), 666–675. https://doi.org/10.1016/j.jval.2021.11.1376

Rossini, E. E., Galeone, C., Lucchetti, C., & Jommi, C. (2024). From indication-based pricing to blended approach: evidence on the price and reimbursement negotiation in italy. *PharmacoEconomics-Open*, *8*(2), 251–261. https://doi.org/10.1007/s41669-023-00467-2

Schmidt, A. F., Groenwold, R. H. H., Van Delden, J. J. M., Van Der Does, Y., Klungel, O. H., Roes, K. C. B., ... Van Der Graaf, R. (2014). Justification of exclusion criteria was underreported in a review of cardiovascular trials. *Journal of Clinical Epidemiology*, *67*(6), 635–644. https://doi.org/10.1016/j.jclinepi.2013.12.005

Siena, L. M., Papamanolis, L., Siebert, M. J., Bellomo, R. K., & Ioannidis, J. P. A. (2023). Industry involvement and transparency in the most cited clinical trials, 2019–2022. *JAMA Network*

Open, *6*(11), e2343425. https://doi.org/10.1001/jamanetworkopen.2023.43425

Van Spall, H. G. C., Toren, A., Kiss, A., & Fowler, R. A. (2007). Eligibility criteria of randomized controlled trials published in high-impact general medical journals: A systematic sampling review. *JAMA*, *297*(11), 1233–1240. https://doi.org/10.1001/jama.297.11.1233

Appendix: Generality of the Equivalence Between Manufacturer Profits and Total Therapeutic Benefit

This appendix demonstrates that our main results on the equivalence between manufacturer profits and total therapeutic benefit hold for general therapeutic benefit functions $b_i(x_i)$, not just the linear specification used in the main text.

Consider a general therapeutic benefit function $b_i(x_i)$ with the following properties: (i) $b_i'(x_i) < 0$ for all $x_i \in [0,1]$ (benefit decreases with therapeutic mismatch); (ii) $b_i(0) > 0$ (maximum benefit at perfect match); and (iii) there exists $\bar{x}_i \in (0,1)$ such that $b_i(\bar{x}_i) = 0$ (some patients have negative benefit).

We show that under our three efficient mechanisms, manufacturer profit maximization leads to the selection of clinically eligible populations that maximize total therapeutic benefit.

Indication-Based Pricing

Under indication-based pricing with expected therapeutic benefit coverage criteria, the constraint is

$$p_i \leq \mathbb{E}[b_i(x_i)|x_i \leq \hat{x}_i] = \frac{\int_0^{\hat{x}_i} b_i(x_i) f_i(x_i) dx_i}{F_i(\hat{x}_i)}. \tag{61}$$

At the optimum, this constraint binds (otherwise the manufacturer could increase p_i without affecting demand). The manufacturer's profit becomes:

$$\begin{aligned} \Pi &= \sum_{i=1}^{n} p_i F_i(\hat{x}_i) \\ &= \sum_{i=1}^{n} \mathbb{E}[b_i(x_i)|x_i \leq \hat{x}_i] \cdot F_i(\hat{x}_i) \\ &= \sum_{i=1}^{n} \int_0^{\hat{x}_i} b_i(x_i) f_i(x_i) dx_i. \end{aligned} \tag{62}$$

This equals the total therapeutic benefit across all indications. The first-order condition for indi-

cation i is:

$$\frac{\partial \Pi}{\partial \hat{x}_i} = b_i(\hat{x}_i) f_i(\hat{x}_i) = 0, \tag{63}$$

yielding $\hat{x}_i^* = \bar{x}_i$, where $b_i(\bar{x}_i) = 0$. The manufacturer includes exactly those patients with non-negative therapeutic benefit.

Population-Weighted Uniform Pricing

Under population-weighted uniform pricing, the constraint is

$$p \leq \sum_{i=1}^{n} \omega_i \mathbb{E}[b_i(x_i) | x_i \leq \hat{x}_i], \quad \text{where} \quad \omega_i = \frac{F_i(\hat{x}_i)}{\sum_{j=1}^{n} F_j(\hat{x}_j)}. \tag{64}$$

With the constraint binding at optimum, the manufacturer's profit becomes:

$$\begin{aligned} \Pi &= p \sum_{i=1}^{n} F_i(\hat{x}_i) \\ &= \sum_{i=1}^{n} \omega_i \mathbb{E}[b_i(x_i) | x_i \leq \hat{x}_i] \cdot \sum_{j=1}^{n} F_j(\hat{x}_j) \\ &= \sum_{i=1}^{n} \frac{F_i(\hat{x}_i)}{\sum_{j=1}^{n} F_j(\hat{x}_j)} \cdot \frac{\int_0^{\hat{x}_i} b_i(x_i) f_i(x_i) dx_i}{F_i(\hat{x}_i)} \cdot \sum_{j=1}^{n} F_j(\hat{x}_j) \\ &= \sum_{i=1}^{n} \int_0^{\hat{x}_i} b_i(x_i) f_i(x_i) dx_i. \end{aligned} \tag{65}$$

Again, profit equals total therapeutic benefit, and the first-order conditions yield $\hat{x}_i^* = \bar{x}_i$ for all i.

Two-Part Tariffs

Under two-part tariffs, the manufacturer's problem is:

$$\max_{\{\hat{x}_i\}_{i=1}^n, \{p_i\}_{i=1}^n, T} \Pi = \sum_{i=1}^{n} p_i F_i(\hat{x}_i) + T \tag{66}$$

$$\text{subject to} \quad p_i \leq \overline{p}_i \quad \forall i \tag{67}$$

$$T \leq \sum_{i=1}^{n} \int_0^{\hat{x}_i} b_i(x_i) f_i(x_i) dx_i - \sum_{i=1}^{n} p_i F_i(\hat{x}_i). \tag{68}$$

As shown in the main text (Proposition 3), the lump-sum constraint binds at optimum. Substituting this binding constraint into the objective function:

$$\begin{aligned} \Pi &= \sum_{i=1}^{n} p_i F_i(\hat{x}_i) + \sum_{i=1}^{n} \int_0^{\hat{x}_i} b_i(x_i) f_i(x_i) dx_i - \sum_{i=1}^{n} p_i F_i(\hat{x}_i) \\ &= \sum_{i=1}^{n} \int_0^{\hat{x}_i} b_i(x_i) f_i(x_i) dx_i. \end{aligned} \tag{69}$$

The unit prices $\{p_i\}_{i=1}^n$ cancel out, confirming they do not affect the optimal clinically eligible populations. The first-order conditions again yield $\hat{x}_i^* = \bar{x}_i$ for all i.

For any therapeutic benefit function $b_i(x_i)$ satisfying our basic assumptions, all three efficient mechanisms lead to the same outcome: the manufacturer sets $\hat{x}_i^* = \bar{x}_i$, where $b_i(\bar{x}_i) = 0$, thereby including exactly those patients with non-negative therapeutic benefit. This maximizes total therapeutic benefit across all indications.

The linear specification $b_i(x_i) = v_i - \tau_i x_i$ used in the main text is thus without loss of generality for our efficiency results, with $\bar{x}_i = v_i/\tau_i$.