The TRUST-CARE Model:
A Market-Based Architecture for Universal Health Coverage

Marcus A. Thomas
mtallis314@gmail.com

May 11, 2026

Working paper. Comments welcome.

Abstract

This paper argues that the central failure of United States healthcare financing is not only inadequate risk pooling, but inadequate credit architecture, and derives a universal coverage system from that premise. The current system's instability—expiring ACA subsidies, fragmented multi-payer administration, and the absence of any structured mechanism for resolving catastrophic medical debt—makes structural rethinking timely rather than merely academic.

Standard responses each inherit the risk-pooling frame's limitations. Premium-based insurance embeds opaque cross-subsidies, severs price signals at the point of service, and generates adverse selection pressures that require increasingly elaborate regulatory patches. Consumer-directed health plans restore cost-sharing but, as the empirical literature documents, cut high-value and low-value care in roughly equal proportions. Single-payer systems achieve price compression through monopsony bargaining power but replace market complexity with centralized administrative complexity and supply-side rationing when reimbursement falls below market-clearing levels.

TRUST-CARE treats healthcare financing as a credit problem. Patients pay for routine care at the point of service, preserving price signals and competitive pressure on providers. For costs that exceed immediate ability to pay, patients access care through an insurance credit mechanism with repayment obligations indexed continuously to verified wealth and income, and automatic forgiveness after a plan-determined horizon. Risk pooling occurs through a government backstop calibrated to insurer loss experience rather than through premium cross-subsidies, making redistribution from healthy to sick explicit and fiscally transparent rather than embedded in premium structures. A preventive care safe harbor protects high-value services from the indiscriminate utilization reductions documented in consumer-directed plan research. A retrospective price dispersion surcharge disciplines within-system price discrimination without imposing price controls or mandated fee schedules.

Two structural implications follow from the credit reframe without additional machinery. Employment lock dissolves when the compensation currently delivered as employer premium payments is redirected as wages into a permanently health-restricted capital account the employee owns directly—converting a status-based benefit into portable, individually controlled health capital without mandates or income-threshold subsidies. Elder and chronically ill populations need not require a separate categorical program: life-tenure plan structures with favorable forgiveness parameters emerge from competitive market dynamics under the existing parameter space, with the government backstop absorbing residual costs as designed.

A static calculation estimates net government cost at \$400–\$404 billion annually, approximately 19 percent of current combined federal and state public health expenditure of \$2.1 trillion. Under a combined adverse stress scenario, including removal of fee-surplus offsetting and higher enrollment-fee financing, public exposure rises to approximately 56 percent of current public health expenditure. These estimates do not model behavioral responses to point-of-service payment, provider pricing adjustments under increased price transparency, or the cost reductions plausibly attributable to preventive care utilization. Nor do they depend on assuming that these dynamic effects materialize. The calculations are therefore best read as fiscal plausibility bounds rather than as forecasts. Whether the design's competitive dynamics, incentive mechanisms, and verification infrastructure perform as specified are empirical questions the paper identifies and structures but cannot resolve analytically.

AI Disclosure. The author used AI-based tools for drafting support, structural review, and copyediting. The author is responsible for the argument, analysis, citations, and final text.

Keywords: universal health coverage; health insurance; medical debt; health credit; catastrophic care; healthcare financing; adverse selection; price transparency; public backstop; Medicare; Affordable Care Act; health economics

1 Introduction

The central failure of United States healthcare financing is not inadequate risk pooling. It is inadequate credit architecture. This paper develops a universal coverage system from that premise.

The diagnostic claim has four parts.

First, premium-based insurance severs the price signal at the point of service. Patients facing copayments do not observe prices. Providers facing negotiated rates do not compete on posted prices. Insurers and providers together set effective prices through network design and benefit tiers rather than through transparent posted schedules. No participant has both the information and the incentive to discipline prices. The Papanicolas finding that the United States pays roughly twice what comparable countries pay for similar services (Papanicolas et al., 2018) is a symptom of this market failure rather than an independent problem. Prior US attempts to reintroduce price signals through deductibles and copayments have not addressed the underlying structure: when insurers still set effective prices through negotiated rates, network rules, and benefit design, patients face a weak and noisy marginal signal, and providers are not competing on posted prices that patients can compare at the point of decision.

Second, risk pooling through premiums embeds redistribution in an instrument that cannot bear it. Premium-based pooling requires healthy enrollees to subsidize sick enrollees through a payment they make voluntarily. This generates adverse selection, requires regulatory patches in the form of community rating, individual mandates, risk adjustment, and open enrollment windows, and ties the redistribution to employment status and plan choice rather than to ability to pay. The redistribution is real but opaque, fiscally unauditable, and politically unstable. Expiring ACA subsidies are the current instance of this instability.

Third, existing cost-sharing instruments are too blunt to discriminate high-value from low-value care. The empirical anchor is Brot-Goldberg et al. (2017), who study a large employer that switched all employees to a high-deductible plan. Enrollees reduced spending by about 12 percent but achieved this reduction by cutting high-value and low-value care in roughly equal proportions. The deductible cliff eliminates marginal cost sensitivity at the point of highest aggregate consumption, once the out-of-pocket maximum is reached. Any system that reintroduces price signals must explain why this attempt failed and what it does differently.

Fourth, the system has no structured mechanism for resolving catastrophic medical debt. Medical debt is the largest source of consumer collections in the United States. The current resolution paths are bankruptcy, charity care, hospital write-offs negotiated case by case, and accumulation. None of these is a designed mechanism. The absence of a structured forgiveness path means catastrophic illness produces financial ruin even for the insured, and the public ultimately absorbs much of the cost through uncompensated care subsidies, Medicaid spend-down, and bankruptcy externalities. The cost is borne; it is borne through the least efficient channels available.

Standard alternatives address subsets of these four problems and create new ones. Premium-based insurance with regulatory patches, exemplified by the Affordable Care Act, addresses the redistribution channel through subsidies but accepts the price-signal failure and the absence of structured debt resolution. Consumer-directed health plans restore cost-sharing but operate within the same blunt-instrument constraint that Brot-Goldberg et al. (2017) document. Single-payer systems achieve price compression through monopsony bargaining power, but they replace market price discovery with administrative price-setting and produce supply-side rationing when reimbursement falls below market-clearing levels.

TRUST-CARE (Transparent, Rational, Universal, Sustainable Treatment—Choice and Access with Responsible Economics) is a fourth option. Its four design moves correspond to the four problems above. Patients pay for routine care at the point of service, restoring the price signal. Risk pooling occurs through a government backstop calibrated to insurer loss experience and funded through general taxation, making redistribution explicit and fiscally auditable. A preventive care safe harbor and a retrospective price dispersion surcharge address the discrimination problem in cost-sharing by protecting high-value services and shifting competitive pressure onto provider pricing behavior. For costs that exceed immediate ability to pay, patients access care through an insurance credit mechanism with repayment obligations indexed continuously to verified wealth and income, and automatic forgiveness after a plan-determined horizon. The credit mechanism is the structured debt resolution that the current system lacks.

Two evaluative criteria constrain the design. *Universal access* means that any person has access to the most plausibly effective health services relevant to their condition. *Individual* *agency* means that a person's decision-making authority over their own or their dependents' health is not arbitrarily restricted. Neither is treated as a hard constraint. In a world of finite resources and differing preferences, no policy can satisfy them simultaneously and continuously for every member of society. They function as criteria against which the proposed system should compare favorably to available alternatives.

The paper is organized as follows. Section 2 reviews the empirical and theoretical literature on each of the four problems. Section 3 describes the basic financing structure and develops the ability-to-pay repayment formula. Sections 4–8 detail the transparency requirements, the government backstop, and the health capital account. Sections 9–11 describe the incentive properties, plan competition, and comparison to alternative systems.

2 Related Literature

The four problems identified in the introduction each have a corresponding body of empirical and theoretical work. This section reviews that work in the order the problems were stated.

Price signals and patient utilization. The RAND Health Insurance Experiment (Newhouse et al., 1993) remains the most comprehensive study of how cost-sharing affects healthcare utilization. Families assigned to free care used roughly 30 percent more services than those facing cost-sharing, with modest effects on health outcomes for most of the population. The primary exceptions were low-income individuals with specific conditions, for whom reduced care had measurable health consequences. The RAND results establish that substantial price sensitivity exists in the patient population, and that for most patients it can be activated without compromising health outcomes.

The price-transparency literature speaks to a separate channel: whether patients act on price information when it is available. Reference pricing programs, in which insurers pay a fixed amount for a procedure and patients pay the remainder, reduce spending on targeted services and shift volume toward lower-cost providers (Whaley, 2019). Most pure transparency mandates produce more modest effects. Patients access price information infrequently and face search costs at the point of care. The implication is that transparency works when paired with a financial stake in provider selection but not as a standalone intervention.

Premium-based pooling and adverse selection. Rothschild and Stiglitz (1976) establish that competitive insurance markets unravel when individuals have private information about their own risk. The standard policy responses are community rating mandates, open enrollment windows, and risk adjustment transfers between insurers. Enthoven (1993) proposed organizing these elements under a managed competition framework with standardized benefit packages and risk-adjusted subsidies. This framework is the intellectual basis for ACA exchange design. Handel (2013) documents that switching costs and plan inertia are large in practice: a substantial fraction of health plan choices are inconsistent with the plans' posted terms, suggesting enrollment decisions are driven partly by inertia rather than optimization. Tebaldi (2022) estimates equilibrium effects of premium subsidy design under the ACA and documents the notch problems generated by threshold-based subsidy schedules tied to the federal poverty level.

The literature on managed competition is the literature on patching premium-based pooling. The patches address subsets of the adverse selection problem at the cost of introducing complexity, fiscal opacity, and threshold-based discontinuities. The present paper takes the reading that the patches are evidence of the underlying instrument's unsuitability rather than evidence of progress toward a stable equilibrium.

Cost-sharing as a blunt instrument. The blunt-instrument problem is most cleanly documented in Brot-Goldberg et al. (2017). Patients facing a high-deductible plan reduced spending by approximately 12 percent and reduced high-value and low-value care in roughly equal proportions. Patients do not reliably distinguish necessary from discretionary services when facing a uniform cost-sharing threshold. The proximate causes are clinical knowledge gaps at the point of decision and opaque or unavailable price information.

The value-based insurance design literature (Chernew et al., 2007) addresses this problem by varying cost-sharing across services according to clinical value. Services with strong evidence of net benefit face lower cost-sharing; lower-value services face higher cost-sharing. The empirical evidence indicates that targeted cost-sharing reductions on high-value services increase utilization of those services, though the magnitude of the effect varies across study designs and target populations.

Catastrophic medical debt. Himmelstein et al. (2009) estimate that medical problems contributed to a majority of personal bankruptcies in the United States during the period studied, with a large share of medical bankruptcy filers carrying private health insurance at the time of the medical event that precipitated the filing.[1] Dobkin et al. (2018) use detailed administrative data to estimate the financial consequences of hospitalization for non-elderly adults, finding that hospitalization causes substantial declines in earnings and increases in unpaid medical bills, bankruptcy, and access to credit. The Consumer Financial Protection Bureau has documented that medical debt is the largest single category on consumer credit reports.

The structural feature of this literature is that catastrophic medical financing in the United States operates through unstructured channels: bankruptcy, hospital write-offs negotiated case by case, charity care, accumulated household debt, and Medicaid spend-down. None of these channels was designed for the purpose. The cost is absorbed; it is absorbed through institutional pathways that produce deadweight loss in the form of bankruptcy administration, financial distress effects on labor supply and consumption, and reduced access to credit for affected households. The absence of a designed forgiveness mechanism is the gap that the TRUST-CARE credit architecture is built to fill.

Universal coverage architectures. Papanicolas et al. (2018) document that the United

[1] The magnitude of the medical contribution to bankruptcy is contested. Dranove and Millenson (2006) reanalyze earlier data from the same research program and argue that medical bills are a contributing factor in a substantially smaller share of personal bankruptcies than the headline figure suggests, with the difference driven largely by how broadly "medical" causes are defined. Dobkin et al. (2018) use linked hospitalization and credit-report data and find that hospital admissions trigger fewer than 5 percent of bankruptcies in their sample. The structural point in the text—that catastrophic medical financing operates through unstructured channels—does not depend on the larger estimate.

States spends approximately twice the average of comparable high-income countries on healthcare with outcomes that are broadly similar or worse. Anderson et al. (2019) argue that prices, not utilization, account for most of the cost difference, and that administrative costs in the US multi-payer system are substantially higher than in single-payer systems. Single-payer systems achieve lower prices primarily through monopsony bargaining power, but they produce rationing through waiting times and provider supply constraints when reimbursement rates fall below market-clearing levels.

Gruber (2008) surveys mechanisms for expanding coverage and documents the importance of subsidies calibrated to income. Income-related health financing has been proposed in various forms, but income-only measures systematically understate resources for asset-rich, income-poor households. The use of wealth alongside income distinguishes TRUST-CARE from prior progressive financing proposals.

3 Basic Structure

This section operationalizes the diagnostic claims from Section 2 into a financing mechanism. Two design moves are developed here. Routine care is paid at the point of service, addressing the price-signal failure documented in Section 2. Costs that exceed immediate ability to pay are financed through an insurance credit mechanism with repayment indexed to verified wealth and income, addressing the catastrophic debt resolution gap documented in Himmelstein et al. (2009) and Dobkin et al. (2018). The remaining two design moves—the preventive care safe harbor and the price dispersion surcharge—are developed in Section 9. The government backstop that replaces premium-based pooling is developed in Section 5.

3.1 Point-of-Service Payment for Routine Care

Patients pay for medical care at the point of service. The payment mechanism is direct: cash, debit, or any standard payment method. Price information is visible to the patient at the moment of decision. The patient evaluates the necessity, expected benefit, and cost of the service before incurring the obligation.

The reintroduction of price signals is meaningful only if the surrounding contracting, transparency, and repayment mechanics make the patient the true marginal payer. Prior US attempts to add price signals through copayments and deductibles failed this test. When insurers continue to set effective prices through negotiated rates and network rules, patients face a weak and noisy marginal signal: the relevant choice is about copays and coverage tiers, not the underlying price level. Providers do not compete on posted prices, and patients cannot reliably compare like-for-like offers at the point of decision. TRUST-CARE treats price visibility and point-of-service payment as a coupled design constraint. Mandatory price posting, documented in Section 4, supplies the transparency requirement. Direct payment supplies the marginal-payer requirement.

3.2 The Insurance Credit Mechanism

For costs that a patient does not wish to settle fully at the point of service, the patient may choose to enroll in an insurance plan. Upon enrollment, the patient obtains an insurance credit instrument that can be used to pay for medical services. The instrument is usable for any medical expense, including routine care, but financing creates an origination fee obligation that direct payment does not. For that reason, credit use is most attractive when the patient judges that paying some or all of the charge immediately would impose excessive short-term financial strain.

The boundary between direct payment and credit use is therefore not the clinical category of the service. Nor is it mechanically defined by the repayment-capacity formula, since Π_t depends on plan parameters that apply only once a charge has entered the credit system. The patient's decision is an ordinary financing decision made under personal circumstances. A patient may pay the full bill directly, finance the full bill through TRUST-CARE credit, or pay part of the bill directly and finance the remaining balance. This split-settlement structure preserves the point-of-service price signal while allowing patients to smooth costs that exceed current liquidity, compete with other household obligations, or would otherwise be difficult to absorb immediately.

Catastrophic care is therefore not the eligibility boundary for TRUST-CARE credit. It is the clearest case in which direct payment is infeasible and structured, ability-to-pay-indexed financing becomes socially valuable.

When a patient incurs medical costs through the credit mechanism, the total debt associated with the financed portion of a charge is

$$\text{Debt} = \text{Financed medical cost} \cdot (1 + P) \tag{1}$$

where P is an origination fee parameter set by the insurer. This amount must be paid by the end of the billing month unless it exceeds the patient's monthly ability-to-pay repayment capacity. The capacity formula is the central instrument that distinguishes this credit mechanism from ordinary consumer credit. It is developed below.

3.3 Two Quantities, Two Roles

The wealth and income measure used by TRUST-CARE serves two distinct roles. The first is the determination of monthly repayment capacity, which is a flow concept: how much the patient can pay this month against outstanding obligations. The second is the determination of the forgiveness horizon, which is a stock-adjusted commitment concept: how many months the obligation remains active before the remaining balance is forgiven. The same underlying wealth measure feeds both, but the two outputs serve different functions in the design and are calibrated separately by plan parameters.

3.3.1 Wealth Definition

$$W(t) = \sum_{\text{properties}} \frac{C^{(i)}_{\text{housing}}(t)}{N^{(i)}_{\text{household}}(t)} + I_{\text{annual}}(t) + A_{\text{investments}}(t) + \sum_{\tau=t-10}^{t} G(\tau) \tag{2}$$

$C^{(i)}_{\text{housing}}(t)$ is the annual cost of housing for property i in year t. For renters, this is annual rent. For owners, it is the greater of annual mortgage payments plus property taxes, or market rent for a comparable property. The sum includes all properties occupied or owned by the person. $N^{(i)}_{\text{household}}(t)$ is the number of adults aged 18 and over residing at property i. $I_{\text{annual}}(t)$ is annual income. $A_{\text{investments}}(t)$ is net investment holdings. $G(\tau)$ is realized capital gains or losses in year τ, summed over the preceding ten years.

The use of wealth alongside income addresses a documented failure of income-only financing measures: asset-rich, income-poor households appear poorer than they are. Gruber (2008) discusses this in the context of subsidy design. The TRUST-CARE measure incorporates housing consumption, financial assets, and a ten-year realized-gains tail to produce a stock measure that tracks resources available for healthcare repayment over a multi-year horizon.

3.3.2 Monthly Repayment Capacity

The patient's monthly ability-to-pay repayment capacity is

$$\Pi_t = \frac{D \cdot I_{\text{discretionary}}(t) + X \cdot W(t)}{2} \tag{3}$$

where D and X are parameters set by the insurer and $I_{\text{discretionary}}(t)$ is discretionary monthly income. If total outstanding debt in month t is less than Π_t, the patient repays only the outstanding amount. Payments are prorated across active charges according to the plan's allocation rule, developed in Section 5. Π_t is a maximum payment capacity, not a minimum payment or a fixed fraction of outstanding debt as in traditional revolving credit. A patient with no outstanding debt pays nothing.

3.3.3 Forgiveness Horizon

The maximum repayment duration assigned to a financed charge is a function of the patient's verified wealth at the time of service:

$$M(W) = M_{\min} + \big(M_{\text{base}} - M_{\min}\big)\phi\left(\frac{W}{W_{\text{ref}}}\right) \tag{4}$$

$M(W)$ is the default number of monthly repayment windows associated with the charge. Each window is an accounting period during which any unpaid balance remains eligible for repayment.

$M_{\min}$ is set by regulatory authorities and establishes a lower bound on the number of repayment windows. Its role is to ensure that every financed charge is treated as a finite obligation rather than as an immediately forgiven transfer. If the outstanding balance is less than the patient's repayment capacity, the charge may still be repaid in a single payment. $M_{\min}$ bounds the duration over which any unpaid balance remains eligible before contractual forgiveness.

M_{base} is plan-specific and specifies the number of repayment windows faced by a reference-wealth patient with $W = W_{\text{ref}}$. Plans with lower values of M_{base} are more generous. The reference wealth W_{ref} serves only to normalize wealth into a dimensionless quantity.

$\phi(\cdot)$ is publicly specified and monotone increasing, with $\phi(0) = 0$ and $\lim_{x\to\infty} \phi(x) = 1$. Repayment windows increase smoothly with wealth, without discontinuities or threshold effects. Low-wealth patients approach $M_{\min}$. Higher-wealth patients approach but never exceed M_{base}.

This structure produces progressive cost-sharing along two dimensions. Lower-wealth patients reach forgiveness more quickly through the horizon function. Lower-income patients face lower monthly repayment obligations through Π_t. The two dimensions are independent: a low-income, high-wealth retiree faces low monthly payments but a long forgiveness horizon, while a low-income, low-wealth patient faces low monthly payments and a short forgiveness horizon. The finite horizon limits the risk of permanent financial burden for any patient at any wealth level.

4 Plan Acceptance and Repayment Tiering

This section addresses the second diagnostic problem from Section 2: adverse selection in premium-based pooling. TRUST-CARE replaces the standard managed-competition response (community rating, individual mandates, risk adjustment, open enrollment) with two instruments. The first is universal acceptance. The second is repayment-history tiering. Together they prevent both insurer cream-skimming and patient moral hazard without requiring the regulatory apparatus that managed competition assembles to patch premium-based pooling.

4.1 Universal Acceptance and Public Plan Terms

Insurers must accept all applicants regardless of health status or pre-existing conditions. The terms offered to patients for each plan—the values of D, X, M_{base}, P, χ, the allocation rule, and any other parameters—must be public. This requirement is structural rather than informational: plan parameters are inputs to the patient's plan choice, and the patient cannot make that choice without observing them.

4.2 Cross-Insurer Repayment History

Plan terms may vary based on the patient's prior repayment history across all previous insurers, not just the current one. This cross-insurer history-sharing rule is essential to the design. Without it, a patient who fails to meet repayment obligations under one insurer could switch to a new insurer and obtain standard terms on the next medical event, eliminating the consequence of non-payment. The rule treats repayment history as a portable record that follows the patient across plan changes, in the same way that consumer credit history is portable across lenders in ordinary credit markets.

The history record is limited to repayment behavior on financed charges. It does not include diagnosis, utilization patterns, or any other clinical information.

4.3 Tiered Terms for Patients with Poor Repayment History

Patients with poor repayment history face more restrictive terms, but retain access to medically necessary care through the credit mechanism. Insurers may apply graduated cost-sharing (temporarily higher values of D, X, and M_{base}, adjusting smoothly without changing the structure of the repayment rule) and credit limit caps (maximum amounts chargeable to the credit system, requiring patients to find alternative financing for amounts above the cap). Insurers retain discretion over additional enforcement instruments—deposits, waiting periods for elective care, co-signer requirements—subject to regulatory caps and to the universal acceptance rule for medically necessary care.

The tiering operates as a continuous adjustment along the existing parameter space rather than as categorical exclusion. A patient with poor history faces tighter terms but the same repayment formula and the same forgiveness horizon structure as any other patient at the same wealth level.

5 Government Backstop and Risk Pooling

The government backstop is the instrument that replaces premium-based pooling. It addresses the second diagnostic problem from Section 2 by relocating redistribution from healthy to sick from a cross-subsidy embedded in voluntary premium payments to a fiscal transfer funded through general taxation. The relocation is what makes the redistribution explicit, auditable, and stable against adverse selection: a patient who never uses the credit system contributes through taxes, not through a premium that they can decline by leaving the pool.

Insurers use accrual accounting to recognize losses on medical costs, excluding origination fees from loss calculations. Origination fees are booked as revenue. The government reimburses a portion of accumulated losses based on the insurer's overall profit margin. The remainder of this section develops the accounting mechanics that make backstop activation calibratable and auditable.

5.1 Charge-Local Amortization and Capacity Allocation

A central accounting objective is to recognize expected losses at the time services are financed, while ensuring that the timing of care does not create artificial cliff effects in recoverability. Each financed medical charge is treated as its own amortized obligation with a finite horizon, while drawing repayment from a single verified stream of ability-to-pay capacity.

A patient incurs a sequence of financed charges indexed by i, with medical cost C_i incurred at time t_i measured in months and anchored to the date of service rather than the date of billing. Charge i generates total debt $C_i(1+P)$ and origination fee revenue C_iP.

At origination, define the charge-specific maximum repayment duration

$$M_i = M\big(W(t_i)\big), \tag{5}$$

and the contemporaneous monthly repayment capacity

$$\Pi_i = \Pi(t_i) = \frac{D \cdot I_{\text{discretionary}}(t_i) + X \cdot W(t_i)}{2}. \tag{6}$$

The patient's repayment capacity over future months is a stream $\{\Pi_t\}_{t>t_i}$, verified and updated according to the tiered schedule in Section 5.7. For accounting at origination, insurers project the stream using currently verified values and apply conservative estimates for components of $W(t)$ pending their next scheduled update. Subsequent verification updates the stream and triggers loss adjustments.

Capacity allocation across multiple outstanding charges depends on the plan's allocation rule, developed below.

5.1.1 The Clustering Window Parameter χ and Allocation Rules

A naïve first-in-first-out allocation rule applied across all outstanding charges creates a timing distortion. Charges arriving in close succession share the same pool of horizon capacity: the second charge in a pair receives little or no allocation within its horizon because the first charge consumes the available months. A patient whose identical charges are spaced many months apart activates additional horizon capacity for each charge and faces larger total repayment. This asymmetry creates a structural incentive for patients with clinical flexibility to cluster elective procedures, and means patients with chronic conditions face higher aggregate repayment than patients incurring equivalent costs as a single acute episode.

To address this, TRUST-CARE introduces the plan parameter $\chi \geq 1$ measured in months, governing how capacity is allocated across charges that originate in temporal proximity. χ is plan-specific, publicly disclosed alongside D, X, M_{base}, and P, and subject to a regulatory minimum $\chi_{\min} \geq 1$ that ensures all charges originating within the same month are eligible for joint treatment. Plans must also specify an allocation rule. Two rules are permitted: *transitive clustering* and *sliding-endpoint pairwise*. Both distribute monthly capacity Π_t across the charges they identify as co-occurring, without exceeding Π_t in any month.

Transitive clustering. Define a graph in which each financed charge is a node and an edge connects charges i and k if $|t_k - t_i| \leq \chi$. The cluster $\mathcal{C}_t$ active in month t is the connected component containing all charges currently within their repayment horizons. All charges in the cluster share monthly capacity pro-rata by outstanding balance:

$$a_{i,t} = \Pi_t \cdot \frac{B_i(t)}{\sum\limits_{k \in \mathcal{C}_t} B_k(t)}, \tag{7}$$

where $B_i(t)$ denotes the outstanding balance of charge i in month t. A chain of charges, each within χ months of the next, forms one cluster regardless of total span.

Sliding-endpoint pairwise. In month t, the cluster $\mathcal{C}_t$ contains the most recent active charge r together with all active charges within χ months of r:

$$\mathcal{C}_t = \{k : k \text{ active in month } t,\ 0 \leq t_r - t_k \leq \chi\}. \tag{8}$$

Capacity is split pro-rata within $\mathcal{C}_t$. As new charges originate, older charges drop out of the cluster while their horizons continue to elapse.

Date-of-service anchoring. Both rules anchor the comparison $|t_k - t_i| \leq \chi$ to dates of service rather than dates of billing. If windows were anchored to billing dates, insurers could delay submitting charges to manipulate cluster composition. Anchoring to service dates removes this incentive: the insurer controls billing timing but not service timing.

Bounded timing asymmetry. Neither rule produces aggregate recovery equal to the single-charge case. A chain of n charges spaced within χ months of one another typically receives more total capacity than a single charge of equivalent total cost, because the union of horizons for a chain exceeds any single charge's horizon. The excess is bounded and computable from charge timestamps, $M(W)$, χ, and the rule. For a chain of n equal charges spaced $\delta \leq \chi$ months apart under transitive clustering, aggregate chain capacity exceeds single-charge capacity by $(n-1)\delta\Pi/(1+P)$. Total capacity claimed in any month sums to at most Π_t by construction, so neither rule permits double-counting across overlapping charges.

How the rules differ. Transitive clustering allocates capacity to every charge in a chain throughout its full horizon at diluted rates: every active charge draws something every month until forgiven. Sliding-endpoint pairwise concentrates capacity on the most recent charges. Once an older charge falls outside the window of the most recent active charge, it receives zero allocation for the remainder of its horizon and is forgiven with whatever balance remains. Sliding-endpoint pairwise produces fewer active-charge-months per chain than transitive clustering, and aggregate patient repayment is correspondingly lower.

Patient preference and selection pressure. The repayment-burden difference between the rules is monotone in expected chain length. Patients expecting only acute, isolated episodes are indifferent: with no chained care, both rules collapse to the same single-charge allocation. Patients expecting chained care from chronic conditions strictly prefer sliding-endpoint pairwise, since older charges are forgiven faster as new charges originate.

This creates a sorting pressure across plans. Holding D, X, M_{base}, P, and χ fixed, plans offering sliding-endpoint pairwise attract chronic patients disproportionately. The mechanism is structurally identical to adverse selection in premium-based systems—patients with private information about their expected utilization sort toward the option that minimizes their expected cost—but the consequence here is not premium spirals. It is concentrated insurer loss on sliding-endpoint plans and, absent rate adjustment, concentrated backstop expenditure on those plans. The reduced backstop reimbursement rate $\rho_{\text{sep}} \leq 1$ specified in Section 5.9 neutralizes this selection pressure by transferring the differential exposure back to the insurer, which prices it through higher origination fees, higher X, or shorter M_{base}. Chronic patients pay less in repayment under sliding-endpoint plans but more in plan terms, with the regulator adjusting ρ_{sep} to maintain aggregate fiscal neutrality across rules.

The two rules are alternatives, not a tier structure. Patients who do not sort by expected utilization pay no penalty for choosing transitive clustering.

Plan-level rule selection. Plans select one allocation rule and disclose it alongside the other plan parameters. The rule applies prospectively to charges originated under the plan. Charges transferred from a prior plan retain the rule under which they originated. Insurers maintain rule-tagged charge records to support backstop reimbursement and audit.

5.1.2 Capacity Allocation Constraint

Total payment allocated across all outstanding charges in any month cannot exceed the patient's verified repayment capacity:

$$\sum_{i:\, t_i < t \le t_i + M_i} a_{i,t} \le \Pi_t. \tag{9}$$

For each charge i, the maximum recoverable principal is

$$R_i = \sum_{t=t_i+1}^{t_i+M_i} a_{i,t}, \qquad \text{Recoverable principal for charge } i = \frac{R_i}{1+P}. \tag{10}$$

5.2 Expected Recovery at Charge Origination

When charge i is originated, the insurer computes an initial feasible allocation for that charge within its horizon, subject to the patient's existing obligations and the cluster membership and allocation rule that apply in each future month. This yields an initial allocated recovery $R_i^{(0)}$ and an initial recoverable principal

$$\widehat{C}_i = \min\left(C_i,\ \frac{R_i^{(0)}}{1+P} \right). \tag{11}$$

The insurer books a receivable of $\widehat{C}_i$ and recognizes immediate loss on the remainder:

$$L_i^{(0)} = \max\left(0,\ C_i - \widehat{C}_i\right). \tag{12}$$

5.3 Loss Recognition Under Updated Verification

The allocation $\{a_{i,t}\}$ is recomputed at each verification event, with frequency determined by the tiered schedule in Section 5.7. If verified capacity decreases, the feasible allocable recovery for some charges may shrink. Let $R_i^{(\text{new})}$ denote the newly feasible allocated recovery remaining for charge i. If the remaining booked receivable exceeds the newly allocable recoverable principal:

$$\Delta L_i = \max\left(0,\ \text{Receivable}_i^{\text{remaining}} - \frac{R_i^{(\text{new})}}{1+P} \right). \tag{13}$$

If verified capacity increases, the improved capacity raises the likelihood that remaining receivables are fully recovered and reduces expected losses on subsequent charges. Previously recognized losses are not reversed, by conservative-accounting convention.

The arrival of a new charge can trigger loss recognition on existing charges even when verified capacity is unchanged. Under transitive clustering, a new charge joining an existing cluster reduces every cluster member's pro-rata share. Under sliding-endpoint pairwise, a new charge can push older active charges out of the current cluster entirely. In both cases, existing receivables are reassessed against the updated allocable recovery and additional losses are recognized as needed.

5.4 Payment Application and Charge Forgiveness

Each month the patient makes a payment up to Π_t. The insurer applies the payment across outstanding charges according to the plan's allocation rule. Payments reduce receivables and outstanding balances. At the end of month $t_i + M_i$, any remaining balance on charge i is forgiven. No new loss is recognized at forgiveness because any unrecoverable portion was already recognized either at origination or during subsequent verification-based adjustments.

5.5 Properties of the Charge-Local Framework

The framework above has four properties that matter for the system as a whole. Each charge is an independent obligation with its own horizon, supporting linear aggregation and modular policy analysis. Aggregate recovery deviates from the single-charge baseline by a bounded amount computable ex ante from public parameters, so neither patients nor insurers can manipulate timing to extract favorable outcomes. Losses are recognized as soon as charges exceed allocable capacity and at each verification event, preventing the accumulation of hidden liabilities that surface only at forgiveness. Each charge produces a tractable audit trail—origination amount, fee revenue, booked receivable, recognized loss, allocated payments, allocation rule, and forgiveness outcome—reconstructible from publicly specified rules without discretionary judgment.

5.6 Required Data Access

Enrollment requires patients to authorize ongoing access to wealth and income records throughout the repayment period. Patients who decline cannot enroll. This is the condition under which society, via the government backstop, finances part or all of a patient's care.

At enrollment, patients authorize insurers to access IRS records for annual income and realized capital gains, Social Security Administration wage data for employment income, property records for housing costs and ownership status, and financial institution reports for investment account balances. Authorizations remain active while any financed charges remain within their repayment horizons or until all outstanding balances are fully repaid, whichever comes first. Authorizations terminate automatically when the final charge obligation ends. Insurers face criminal penalty for unauthorized sharing or sale of patient data.

5.7 Verification Frequency and Data Infrastructure

Verification frequency is tiered by the materiality of outstanding financed balances. Annual verification applies as a baseline for all enrolled patients; verification increases to quarterly and then monthly when outstanding balances on Backstop-Eligible Care, the subset of charges eligible for public loss sharing defined in Section 5.11, exceed regulated thresholds. The tiering preserves the conservative public exposure property for the cases that matter fiscally while avoiding disproportionate administrative burden for low-balance patients.

Components of $W(t)$ update at different natural frequencies. Wage data through the Social Security Administration updates continuously; federal tax returns are annual; property assessments are typically annual or biennial; investment account balances are available in near real-time. The composite capacity measure Π_t is therefore a blend of data sources with varying staleness. Insurers use the most recently available value for each component and apply conservative estimates for components pending their next scheduled update. Annual tax reconciliation triggers a mandatory receivable reassessment for all active charges.

Monthly verification at scale requires data infrastructure that does not currently exist in the United States. Integrated real-time reporting across IRS, SSA, and financial institutions is a precondition for implementation at the verification frequencies specified here, not an assumed capability. Countries with more integrated financial reporting systems provide reference points for implementation sequencing. The tiered structure concentrates infrastructure demands on the patient population where verification frequency has the greatest fiscal consequence, reducing the minimum viable infrastructure threshold relative to a uniform monthly requirement.

Patients must report material changes in wealth or income within 30 days. Material changes include job loss, inheritance, property acquisition or sale, liquidation of investments exceeding $10,000, and income changes exceeding 20 percent annually. Enrollment documents specify which data sources are accessed, how often verification occurs, what constitutes a material change, what documentation is required, and consequences of noncompliance. These terms are non-negotiable and identical across all plans offered by an insurer.

5.8 Default and Forgiveness

If a patient stops making payments despite verified ability to pay, remaining booked receivables are written off as additional loss. The failure to repay enters the patient's repayment history and may affect future plan terms under Section 4.

After a charge's assigned repayment duration elapses, the remaining balance is forgiven regardless of amount. Losses were recognized progressively as circumstances changed, so forgiveness creates no additional loss.

5.9 Backstop Activation

Annual recognized losses for an insurer equal the sum of initial losses at charge origination plus adjustments during repayment from changed circumstances or defaults. Losses are tracked separately by allocation rule.

The government provides backstop funding when annual losses exceed net insurer fee surplus. The reimbursement rate is rule-aware: plans using transitive clustering receive the standard rate; plans using sliding-endpoint pairwise receive a reduced rate calibrated by the regulator.

$$\text{Backstop payment} = \max(0,\ \rho_{\text{trans}}(1 - P_{\text{eff}})L_{\text{trans}} + \rho_{\text{sep}}(1 - P_{\text{eff}})L_{\text{sep}} - S_{\text{fees}}) \tag{14}$$

L_{trans} and L_{sep} are annual recognized losses under transitive clustering and sliding-endpoint pairwise. P_{eff} is total origination fees collected divided by total medical costs financed during the year. S_{fees} is net insurer surplus from enrollment and cancellation fees. ρ_{trans} and ρ_{sep} are reimbursement rates with $\rho_{\text{trans}} = 1$ as the baseline.

The differential rate $\rho_{\text{sep}} \leq 1$ exists because of the patient selection pressure identified in Section 5.1.1. Holding plan terms fixed, patients expecting chained care prefer sliding-endpoint pairwise, and absent rate adjustment this preference would concentrate chronic-patient losses on sliding-endpoint plans and on the public fisc. Setting $\rho_{\text{sep}} < 1$ transfers that differential exposure back to the insurer, which prices it through higher origination fees, X, or shorter M_{base} on its sliding-endpoint plans rather than through differential public subsidy. The regulator's calibration target is fiscal neutrality across rules: a patient with a given expected cost trajectory should generate similar government exposure regardless of which rule their plan uses. The calibrations in Section 10 suggest ρ_{sep} near 1 is appropriate under the population mix and parameter ranges examined there. The regulator monitors actual exposure by rule and adjusts ρ_{sep} if observed sorting produces wider divergence than the static calibrations imply.

The max operator applies at the aggregate level. Fee surplus generated by healthy enrollees offsets losses generated by chronically ill enrollees before any government payment is owed. Insurers with sufficiently large healthy populations may generate fee surplus large enough to fully offset losses, producing zero backstop activation. The government bears only residual losses after both origination fee revenue and enrollment fee surplus are exhausted.

This structure ensures insurers bear losses proportional to their profit margins. The government does not subsidize the origination fee portion. An insurer charging a higher P receives less government support per dollar of loss, creating competitive pressure to minimize origination fees while preventing moral hazard from full government reimbursement.

5.10 Fee-Surplus Pooling: A Design Choice

The backstop formula subtracts net fee surplus S_{fees} before public reimbursement activates. This embeds a within-pool redistribution: enrollment fee surplus generated by healthy enrollees offsets losses generated by chronically ill enrollees within the same insurer's book before any government funding is owed. The choice deserves explicit defense, since it sits in tension with the introduction's commitment to redistribution operating through fiscal policy rather than through cross-subsidies embedded in insurance instruments.

The philosophically cleaner alternative. A strict implementation of the redistribution principle yields a simpler formula:

$$\text{Backstop payment} = \max(0,\ \rho_{\text{trans}}(1 - P_{\text{eff}})L_{\text{trans}} + \rho_{\text{sep}}(1 - P_{\text{eff}})L_{\text{sep}})\,. \tag{15}$$

Under this formula, all redistribution from healthy to sick occurs through public taxation funding the backstop. Within-insurer fee pooling plays no role. Insurers compete on enrollment and cancellation fees as pure pricing instruments rather than as implicit cross-subsidization mechanisms. The cleaner formula is more consistent with the diagnostic framing in Section 2.

The case for retaining within-pool pooling. Despite its philosophical clarity, equation (15) weakens an incentive channel that managed-competition theory (Enthoven, 1993) identifies as material even when premium discrimination is prohibited. With universal acceptance and uniform fees within plans, insurers cannot cream-skim through underwriting, but they retain selection instruments that operate through plan design. An insurer can attract healthier enrollees by offering plans with attributes correlated with low expected utilization (specialty network composition optimized for routine rather than chronic care, marketing targeted at younger enrollees, customer service infrastructure designed for episodic engagement) without violating any acceptance or pricing rule. These selection pressures are subtler than underwriting but, in the empirical literature on managed competition, persistent.

The S_{fees} subtraction gives insurers a direct fiscal stake in their realized enrollee mix. An insurer whose marketing and plan design attract a healthier-than-average population generates fee surplus that reduces backstop reliance, but the surplus must absorb that insurer's own losses before the public reimburses anything. Under equation (15), the surplus is pure profit and the public absorbs losses regardless of mix. The within-pool channel partially offsets soft selection pressures that universal acceptance alone cannot eliminate.

The pooling parameter β. A fee-pooling parameter $\beta \in [0, 1]$ makes this design choice tunable:

$$\text{Backstop payment} = \max(0,\ \rho_{\text{trans}}(1 - P_{\text{eff}})L_{\text{trans}} + \rho_{\text{sep}}(1 - P_{\text{eff}})L_{\text{sep}} - \beta \cdot S_{\text{fees}})\,. \tag{16}$$

$\beta = 1$ recovers the formula above: full within-pool pooling. $\beta = 0$ recovers equation (15): no within-pool pooling.

The base case throughout the paper is $\beta = 1$, including in the fiscal calculation of Section 10. This reflects a conservative posture toward selection pressures the design has not yet been empirically shown to neutralize. The parameter is a regulatory instrument, not a fixed feature of the model. Once the system is deployed and selection dynamics are observable, the regulator may adjust β toward zero. Conditions favoring lower β include empirical evidence that plan-design selection is small or absent, political-economy preference for taxation-funded redistribution over within-pool pooling, or simplification pressure that prefers a cleaner formula.

Fiscal implications of β variation. Section 10 reports per-person net government cost under $\beta = 1$ at \$1,200–\$1,212, bounded across the two allocation rules. Under $\beta = 0$, fee surplus would not offset losses, raising per-person cost by the full \$1,067 fee revenue figure to \$2,267–\$2,279. Aggregate exposure at US population would scale to \$755–\$760 billion, compared to \$400–\$404 billion at $\beta = 1$. Both figures remain substantially below combined federal and state public health expenditure of approximately \$2.1 trillion (Centers for Medicare & Medicaid Services, 2023), so the choice of β does not threaten fiscal plausibility.

Administrative cost. The backstop formula does not separately reimburse insurer administrative cost. Origination fees P are intended to fund administration, credit risk on outstanding receivables, and insurer profit margin; competitive pressure across plans disciplines P toward levels that cover these costs without excess. This treatment is implicit in the $(1 - P_{\text{eff}})$ factor, which excludes the origination-fee portion from public reimbursement.

This works only if administrative costs scale roughly with charge volume. To the extent costs are per-enrollee rather than per-charge—wealth and income verification, dispersion-surcharge maintenance, BEC classification documentation—they are recovered through enrollment fees rather than through P. Under $\beta = 1$, per-enrollee admin recovery is indirectly funded by the same fee structure that performs within-pool pooling. Under $\beta = 0$, fees become pure insurer revenue and admin recovery operates through unrestricted profit. In both cases, competition is the discipline mechanism. A future extension might introduce explicit admin reimbursement through a separate parameter, but doing so would create the centralized cost-definition and audit infrastructure that single-payer systems require and that TRUST-CARE deliberately avoids.

5.11 Backstop-Eligible Care

The government backstop defines the boundary of social risk-sharing. TRUST-CARE preserves universal access to treatment through the credit mechanism, but not all health-related consumption is subject to public loss sharing. Backstop-Eligible Care (BEC) defines this boundary.

A financed medical charge is classified as BEC if failure to provide the service within a clinically reasonable time window would foreseeably result in irreversible morbidity or mortality, in substantial future health expenditure likely to be borne by public risk pools, or in the transmission of significant health risk to others. The standard is interpretive rather than computational: it does not require numerical probability estimates at the patient level, but anchors eligibility in the foreseeable consequences of delay within ordinary clinical reasoning.

Eligibility is determined at the level of clinical context rather than procedure code alone. The same procedure may be BEC in one context and not in another, depending on indication, degree of functional impairment, and expected consequences of deferral.

TRUST-CARE establishes a Backstop Exclusion Registry specifying clinical contexts in which financed services are presumptively not eligible for public loss sharing. Exclusions are defined in terms of indication and functional context rather than specialty labels. The Registry is reviewed at regular intervals and updated through a transparent public process. Unless a service-context appears in the Registry as excluded, financed care remains eligible for BEC classification subject to audit.

Insurers classify charges as BEC at origination, subject to ex post audit. Regulatory oversight evaluates alignment with the interpretive standard and the current Registry. Where systematic misclassification is identified, backstop reimbursements are subject to clawback and financial penalty. The audit-based structure avoids centralized preauthorization bottlenecks while maintaining fiscal discipline.

The Registry does not encode static clinical guidelines into statute. Medicine evolves, and the boundary of public obligation evolves with it. The principled standard paired with a revisable exclusion framework produces a system adaptable to scientific progress and resistant to gradual expansion of public subsidy into preference-driven health consumption.

5.12 Financing of Non-Backstop-Eligible Services

Services not classified as BEC may still be financed through the credit infrastructure. Denial of public loss sharing does not imply denial of access or prohibition of financing. Individuals remain free to obtain and finance preference-driven services, but the associated credit exposure is private rather than social.

Non-BEC charges are excluded from backstop calculations. Losses arising from such charges are borne entirely by the originating insurer. Insurers segregate BEC and non-BEC receivables and report them separately.

Non-BEC charges are placed in a distinct repayment queue. Monthly ability-to-pay capacity allocated to BEC obligations is determined without reference to non-BEC balances. Only residual discretionary capacity, after satisfying the BEC schedule, is applied to non-BEC debt. The clustering window and allocation rules apply uniformly within each queue: BEC charges cluster among themselves, non-BEC charges cluster among themselves, and the two queues do not share clusters or capacity.

Insurers retain discretion over underwriting standards, origination fees, repayment durations, and pricing terms for non-BEC financing. Because these charges are not supported by the backstop, insurers price and manage their credit risk accordingly.

5.13 Risk Pooling Through Taxation

Risk pooling occurs through the government backstop. When insurer losses require government funding, these costs are covered through tax revenues. Healthier and wealthier individuals who pay cash for routine care and never use the credit system contribute to catastrophic costs through taxation. Redistribution from healthy to sick is explicit and operates through fiscal policy, not through cross-subsidies embedded in insurance instruments. This is the contrast with premium-based pooling that motivates the design choice in Section 2.

6 Plan Competition Under Continuous Enrollment

The TRUST-CARE goal is not to prevent sorting, it is to prevent hidden subsidy extraction. The system allows patients to switch between insurance plans at any time without open enrollment periods. The design constraints are:

1. Patients should face low barriers when switching plans for legitimate reasons.

2. Patients should be able to switch at any time rather than waiting for annual enrollment windows.

3. Different plans should offer different coverage terms (M_{base}, D, X, P, χ, allocation rule) to create meaningful market competition. Plans with larger χ extend cluster membership across a broader window. Plans selecting transitive clustering distribute

capacity across all charges in a chain throughout their full horizons; plans selecting sliding-endpoint pairwise concentrate capacity on the most recent charges and forgive older charges faster. The two rules trade off differently for chronic patients versus acute patients, with implications for plan selection developed in Section 6.3.

4. Plans may charge enrollment and early termination fees as part of their revenue model. Fee caps are regulated.

5. Plans may not alter repayment terms retroactively for any existing balances, including balances transferred from previous plans. Terms apply prospectively only to new medical charges incurred after enrollment. The allocation rule applicable to a charge is the rule of the plan under which that charge originated, regardless of subsequent plan switches.

6.1 Gaming Prevention Through Switching Costs

Consider a patient enrolled in Plan A who receives a cancer diagnosis and evaluates switching to Plan C. Switching costs include a cancellation fee from Plan A (up to \$5,000) and an enrollment fee for Plan C (up to \$5,000), for a total of up to \$10,000. The benefit of switching is the reduction in repayment obligations, which depends on the patient's income and wealth under the distinct ability-to-pay formulas. For many patients—particularly those with modest incomes for whom ability-to-pay obligations remain small regardless of plan parameters—the switching cost exceeds any plausible repayment savings. For wealthier patients who would face substantial repayment obligations, the switching cost still represents meaningful friction.

This creates a stable equilibrium where plans can maintain differentiated coverage terms without experiencing adverse selection spirals. Patients select plans based on ex ante assessment of catastrophic risk, and switching costs are calibrated such that ex post optimization, while possible, is sufficiently expensive to preserve plan differentiation.

6.2 Plan Stratification and Fee Competition

Under this structure, plans naturally stratify according to the relationship between enrollment costs and coverage generosity. Three illustrative plans, deliberately differing on multiple parameters to convey the design space rather than to isolate any single parameter's effect:

Plan A charges a \$1,000 enrollment fee and \$2,000 cancellation fee while offering harsh terms ($M_{\text{base}} = 48$, $D = 15\%$, $X = 5\%$, $\chi = 1$, transitive clustering). The minimal χ means only charges originating in the same month share capacity. The transitive rule has little to operate on at this χ. This plan attracts young, healthy individuals who assess their risk as low and expect at most isolated acute episodes.

Plan B charges a \$5,000 enrollment fee and \$5,000 cancellation fee with moderate terms ($M_{\text{base}} = 24$, $D = 10\%$, $X = 3\%$, $\chi = 4$, transitive clustering). The four-month window groups charges arising from a single clinical episode or short course of treatment. This plan serves risk-neutral individuals seeking balanced protection.

Plan C charges a \$10,000 enrollment fee and \$10,000 cancellation fee with generous terms ($M_{\text{base}} = 12$, $D = 5\%$, $X = 1\%$, $\chi = 12$, sliding-endpoint pairwise). The twelve-month window combined with the sliding-endpoint rule means older charges in a chronic trajectory are forgiven faster as new charges originate. This plan attracts patients expecting chained care: those with chronic conditions, family histories of illness, or older enrollees.

Plans can also compete on the structure of fees rather than just their level. A plan offering a \$3,000 enrollment fee with an \$8,000 cancellation fee creates an easy-entry, committed-stay structure, while a plan with an \$8,000 enrollment fee and \$3,000 cancellation fee requires upfront commitment but offers later flexibility.

The illustrative plans bundle the allocation rule with other parameters (Plan C is more generous on D, X, M_{base}, and χ, in addition to using sliding-endpoint pairwise), reflecting how plans are likely to position themselves rather than the contribution of any single parameter. Section 6.3 examines the allocation-rule axis on its own.

6.3 Allocation Rule as a Competitive Parameter

Holding D, X, M_{base}, P, and χ fixed, the two allocation rules produce different repayment burdens for the same patient. Patients expecting only acute, isolated episodes are indifferent: with no chained care, both rules collapse to the same single-charge allocation. Patients expecting chained care from chronic conditions strictly prefer sliding-endpoint pairwise, since older charges drop out of the cluster as new charges originate and are forgiven with whatever balance remains, reducing total repayment. The strength of preference grows with expected utilization frequency and chain duration.

This produces a sorting equilibrium across plans. With other terms equal, sliding-endpoint pairwise plans attract chronic patients disproportionately. The mechanism is structurally identical to adverse selection in premium-based systems — patients with private information about their expected utilization sort toward the option that minimizes their expected cost — but the consequence is concentrated insurer loss on sliding-endpoint plans rather than premium spirals. The differential backstop reimbursement rate $\rho_{\text{sep}} \leq 1$ specified in Section 5.9 addresses this selection by transferring differential exposure back to the insurer.

In practice, insurers pricing this exposure adjust the surrounding plan terms. A sliding-endpoint plan can price its selection exposure through any combination of higher origination fees P, higher X, shorter M_{base}, lower χ, or higher enrollment fees. The combined effect is that chronic patients pay less in repayment under sliding-endpoint plans but more in plan terms. The regulator calibrates ρ_{sep} to maintain fiscal neutrality across rules: a patient with a given expected cost trajectory generates similar government exposure regardless of which rule their plan uses, with the difference absorbed by plan terms rather than public subsidy.

The empirical analysis in Section 10 indicates aggregate exposure under the two rules is similar at base-case parameters, supporting ρ_{sep} near 1. If realized sorting concentrates chronic-patient losses on sliding-endpoint plans more strongly than the static calibration implies, the regulator reduces ρ_{sep} accordingly. Plan competition on the rule axis is thus disciplined by the backstop rate, not by direct regulation of which rule any plan may offer.

A patient with no clear utilization expectation faces no penalty for choosing transitive clustering, since both rules deliver the same allocation in the absence of chained care. The two rules are alternatives, not a tier structure.

6.4 Financing Enrollment Fees

Enrollment fees may be financed through the credit system at the insurer's discretion. A patient who experiences a health event while uninsured may enroll in a plan by charging the enrollment fee to the insurance credit card, subject to the same ability-to-pay repayment formula. This ensures that catastrophic protection remains available even to patients who cannot pay enrollment fees in cash, while the repayment obligation preserves the switching deterrent.

A financed enrollment fee is not a medical charge and does not meet the Backstop-Eligible Care standard in Section 5.11. It is therefore placed in the non-BEC queue and serviced only after BEC obligations are met, and any losses arising from forgiveness of unpaid enrollment fees are borne entirely by the originating insurer rather than through the government backstop. Insurers extending enrollment-fee financing therefore bear the associated credit risk directly. This is consistent with the broader principle that public loss sharing applies only to medical care; financing of plan-entry costs, like financing of non-BEC medical services, remains a private credit transaction between insurer and patient.

7 Elder Care and Life-Tenure Plans

Elderly and permanently disabled patients represent the highest-cost segment of any universal coverage architecture. Their utilization profiles differ qualitatively from the working-age episodic model that motivates much of the preceding design: charges arrive continuously rather than episodically, income flows are fixed and limited, and wealth is largely illiquid. The question is whether TRUST-CARE handles this population through its existing mechanisms or requires categorical additions. This section argues that no new machinery is required. Life-tenure plans emerge as a natural market equilibrium, the government backstop absorbs chronic high-cost populations as designed, and the illiquid wealth problem is addressable within the existing parameter space.

7.1 Parameter Adaptation for Chronic High-Utilization Populations

The plan parameter space described in Section 6 spans the full range of coverage terms relevant to elderly patients. An insurer serving this population would set X near zero, removing the weight on asset holdings in Π_t and grounding repayment capacity primarily in discretionary income from Social Security and pension flows. A short M_{base} accelerates forgiveness for patients generating continuous charges, reducing the risk that accumulated obligations outpace a fixed income stream. A large χ combined with sliding-endpoint pairwise allocation means that older charges in a chronic trajectory are forgiven continuously as new charges originate, so the repayment horizon does not expand indefinitely with utilization. A low or zero origination fee P reduces per-charge burden for patients whose costs are certain to recur.

No modification to the core architecture is needed to produce these terms. They are combinations of parameters already present in the system, and an insurer pricing a chronic high-utilization population would be led to them by the same competitive logic that governs working-age plan design.

To make this concrete, consider a patient generating a financed charge of C every month indefinitely under sliding-endpoint pairwise with $\chi \geq 1$. In any given month, the cluster contains the most recent charge together with all active charges within χ months of it. As each new charge originates, the oldest active charge falls outside the cluster window and receives zero further allocation, exhausting its horizon with whatever balance remains and then being forgiven. The patient's repayment obligation at any point reflects only the charges within the current cluster window, not the full accumulated history of care. Forgiveness operates continuously rather than as a terminal event. The government backstop absorbs the unrecovered balances at each forgiveness event, exactly as it does for acute high-cost episodes in working-age patients.

7.2 The Life-Tenure Equilibrium

The conditions that make elderly-tailored plan parameters appropriate are permanent and observable to both insurer and patient. Age, chronic burden, fixed income, and illiquid wealth are verifiable at enrollment and do not change in ways that alter the fundamental character of the plan relationship. This is not the environment that motivates the switching cost analysis in Section 6. There, entry and exit fees are calibrated to preserve the patient's ability to switch when circumstances genuinely change, while making ex post optimization around a known acute episode sufficiently costly to be unattractive. For elderly patients in a life-tenure plan, circumstances have already changed. The plan terms reflect a known, permanent utilization profile. The patient has no reason to switch because no alternative plan offers more favorable terms for their situation, and the insurer is pricing a relationship that will persist until death.

The fee structure on elderly-tailored plans is asymmetric by design. The entry fee is low or zero: the target population is asset-rich and income-poor, and a large upfront cost would block access for precisely the patients the plan is designed to serve. The exit fee is large. Its function is to make enrollment unattractive for patients who are not the target population. A young or otherwise healthy patient enrolling in a plan with short M_{base}, low X, large χ, and sliding-endpoint pairwise allocation would obtain favorable forgiveness terms at low expected cost to themselves, since they are unlikely to generate the continuous charge stream those terms are designed to handle. The large exit fee prices away that option. Early exit by the target population, including death, is not a problem the insurer needs to protect against: the plan was priced knowing that tenure is finite and that the enrolled population is elderly. The exit fee is not a penalty on the target population leaving but a deterrent to the wrong population entering.

A secondary consideration reinforces this structure. An insurer serving an elderly population accumulates provider pricing intelligence relevant to the chronic conditions and service patterns of that population: dispersion surcharge histories, BEC classification patterns, and price comparison data for the procedures and specialist networks its enrollees use most. This knowledge has value for the enrolled population and is not transferable to a patient with a fundamentally different utilization profile. A large exit fee partially recovers the fixed cost of building this infrastructure if a non-target enrollee exits after a short tenure without contributing to the charge history the insurer's pricing knowledge is built around.

The result is a life-tenure plan segment that emerges from market dynamics without regulatory mandate. Insurers competing for elderly enrollees converge on a low-entry, high-exit fee structure with favorable repayment parameters, and the fee asymmetry naturally limits enrollment to the population the plan is designed for. This is not a distinct product category requiring separate regulation. It is a region of the plan parameter space that the market occupies in response to the demographic and clinical reality of aging.

7.3 Backstop Implications

Elderly-tailored plans will generate higher backstop activation rates than working-age plans. This is by design. The government backstop is the explicit residual bearer of costs that exceed patient repayment capacity, and elderly patients with limited income flows and continuous utilization will exhaust their repayment capacity more frequently and more predictably than younger patients with episodic care needs. Concentration of elderly patients in tailored plans concentrates backstop activation, but this is structurally identical to the function Medicare serves in the current system. The difference is that activation occurs through the same transparent, parameterized formula that governs the rest of the system rather than through a separate categorical program with its own pricing and eligibility rules.

The regulator monitors backstop activation rates by plan demographic composition and retains the existing instruments for adjustment. The ρ parameters in the backstop formula can be calibrated to reflect differential loss experience across plan types if observed activation diverges from projections. No new reimbursement category is required. The fiscal plausibility calculation in Section 10 does not separately model this population, and a full accounting would require cell definitions that reflect the elderly income and wealth distribution rather than the general population quintiles used there. That extension is a natural direction for subsequent work. The qualitative implication is clear: per-person backstop costs for elderly enrollees will exceed the population mean, and aggregate fiscal exposure scales accordingly. Whether this remains within the fiscal envelope established in Section 10 depends on the fraction of the population in the life-tenure segment and on the plan parameters that market competition produces for that segment.

7.4 The Illiquid Wealth Problem

The wealth definition in equation (2) includes home equity and retirement account balances. For working-age patients, these holdings represent genuine resources that can in principle be liquidated to meet repayment obligations. For elderly patients on fixed incomes, this assumption is less defensible. A patient with \$300,000 in home equity and \$18,000 in annual Social Security income does not have repayment capacity proportional to the asset value: the home is not a monthly cash flow, and liquidating it may be infeasible or incompatible with the patient's basic housing security.

The primary resolution is through plan design. An insurer setting X near zero effectively removes asset holdings from the capacity calculation. This is the natural competitive response for an insurer whose enrolled population consists predominantly of asset-rich, income-poor patients. The plan parameter absorbs the illiquidity problem without requiring a modification to the wealth formula.

A supplementary regulatory convention is also available within the existing architecture. Insurers may be permitted or required to exclude primary residence equity from $A_{\text{investments}}(t)$ for patients above a specified age or below a specified income threshold, or to apply a liquidity discount to retirement account balances that are subject to withdrawal penalties or required minimum distribution constraints. Either convention preserves the formula's structure and its progressive properties while acknowledging that asset illiquidity is a systematic feature of elderly finances rather than an idiosyncratic edge case. The appropriate parameterization of such a convention is an empirical question that can be addressed once the system is deployed and insurer pricing of elderly-tailored plans is observable.

The illiquid wealth problem does not require a categorical solution. It is a calibration question that the plan parameter space and, if necessary, a narrow regulatory convention can address within the existing framework.

8 Health Capital and Wage Unbundling

The TRUST-CARE financing architecture eliminates the actuarial function of insurance premiums. Repayment obligations arise at service and scale to verified ability to pay. This creates a question about what happens to the household resources currently committed to premium payments.

A common objection to premium-replacing designs is that individuals will face new costs that employers previously absorbed. This framing is misleading. Employer premium contributions are a component of total compensation. When an employer pays a premium on an employee's behalf, the employee receives correspondingly less in wages. The relevant question is not whether the employee bears a cost—they do under both arrangements—but whether the resource commitment is routed through an instrument the employee owns.

8.1 The Tax-Free TRUST-CARE Repayment Account

To redirect this resource stream, TRUST-CARE introduces a dedicated tax-advantaged account, the Tax-Free TRUST-CARE Repayment Account (TFRA), used exclusively for TRUST-CARE repayments. The TFRA is optional but designed to absorb a portion of the compensation that many households already commit to healthcare. Funds in the TFRA cannot be withdrawn for non-health consumption and cannot be converted into ordinary wealth. This distinguishes the TFRA from a general savings account: it is a health-restricted balance sheet item whose purpose is to smooth and reduce the ex post repayment obligations that TRUST-CARE creates.

The TFRA strengthens the system without being a prerequisite for access. Access is guaranteed by the financing rules themselves. The TFRA is a voluntary buffer that reduces stress on the government backstop and increases individual control over repayment timing. This distinction matters for low-income individuals: under TRUST-CARE, a person with minimal TFRA balances can still obtain care and repay according to verified income and wealth, with no compounding interest and with plan-specified forgiveness. The TFRA is helpful when present, but absence of TFRA savings does not create a denial mechanism.

8.2 Wage Unbundling

A central implication of TRUST-CARE paired with TFRA is that employers no longer need to administer healthcare benefits as a distinct compensation category. The desired outcome is that compensation previously delivered as premium payments is instead delivered as wages, with employees directing a portion to TFRA contributions if they choose. This eliminates job lock, removes differential access based on employer size and bargaining power, and restores transparency to total compensation. Reintroducing employer contributions as a privileged channel would recreate the status-based inequities the model is designed to eliminate.

8.3 Inheritance and the Transfer Tax

The TFRA is intended to roll into retirement, providing continuity of health capital when earned income declines. It is also inheritable. Unused TFRA balances represent foregone healthcare consumption and preserved health capacity. Permitting transfer respects ownership while maintaining the account's restricted purpose.

To prevent TFRA from functioning as an unbounded tax-free intergenerational compounding vehicle, balances are subject to a TFRA-specific flat transfer tax at death. This tax is conceptually distinct from general estate taxation: TFRA balances are not part of the ordinary estate tax base. The transfer tax is a maintenance levy whose revenue replenishes the universal backstop. A flat rate is motivated by simplicity and by the framing of the instrument: the tax is not a redistributive mechanism but a device for returning accumulated health capital surplus to the system that made the capital formation possible.

8.4 Life-Cycle Interpretation

During working years, employer premium spending is displaced by wages. Workers optionally contribute to TFRA, converting premium-like flows into owned health capital. When care is needed, access is guaranteed and repayment is scaled to means, with TFRA available as a voluntary buffer. In retirement, TFRA remains available as health capital. At death, balances transfer after a flat levy that replenishes the backstop. Across the full arc, the system moves from premium-based intermediation toward a combination of individual health capital and structured ex post financing, without requiring ex ante risk pooling as the organizing principle.

9 Incentive Structure

Several features of this system create specific incentives. Patients face the full cost of routine medical care, which encourages price comparison and evaluation of medical necessity for non-catastrophic services. The availability of the credit system for catastrophic events means that no patient faces complete financial ruin from illness, but the repayment obligation creates meaningful cost-sharing. The tiering system based on repayment history provides motivation to meet payment obligations, since failure to do so results in more restricted coverage should future medical needs arise.

9.1 Preventive Care Safe Harbor

A central concern with point-of-service payment is that patients may reduce not only low-value care but also high-value preventive and early diagnostic services. The evidence from Brot-Goldberg et al. (2017) makes this concern concrete. To address it, the TRUST-CARE framework includes a *preventive care safe harbor*: a narrowly defined set of services for which financing terms are strictly more favorable than for general elective consumption.

Safe-harbor services include evidence-based preventive interventions and screenings with strong clinical consensus regarding net benefit—routine vaccinations, well-child visits, pregnancy-related preventive care, age- and risk-appropriate cancer screenings, hypertension and diabetes screening, and other guideline-supported preventive services. For these services, insurers must offer standardized preferred terms that reduce or eliminate patient-facing financing friction. Preferred terms may include:

1. **Zero origination fee:** For safe-harbor claims, $P = 0$.
2. **More favorable repayment parameters:** Safe-harbor claims use lower values of D and X, and a lower base repayment duration M_{base}, than would otherwise apply.
3. **First-dollar credit access:** Safe-harbor services may be charged to the insurance credit system regardless of credit-limit caps, preventing prior repayment history from blocking preventive access.
4. **Non-punitive classification:** Use of the safe harbor does not worsen a patient's repayment-history bucket, provided the patient follows the applicable ability-to-pay repayment schedule.

Safe-harbor eligibility must be determined by a publicly specified, periodically updated list tied to clinical guidelines and administered with transparent criteria. Services outside the safe harbor remain subject to standard cash payment norms for routine care and standard financing terms for catastrophic events.

The wealth-adjusted repayment period $M(W)$ plays a central role in balancing incentives across wealth levels. Higher-wealth patients face longer repayment obligations, creating progressive cost-sharing that aligns payment capacity with payment duration. Lower-wealth patients reach debt forgiveness more quickly, reducing the risk that one medical event creates permanent financial burden. The function $M(W)$ must be calibrated such that even wealthy patients face finite repayment periods while low-wealth patients are not overwhelmed by extended obligations.

10 Fiscal Plausibility: A Static Backstop Calculation

This section establishes the order-of-magnitude fiscal exposure of the backstop mechanism. Under base-case parameters, aggregate net government cost lies between \$1,200 and \$1,212 per person, scaling to \$400–\$404 billion at US population. This is approximately 19 percent of current combined federal and state public health expenditure of \$2.1 trillion (Centers for Medicare & Medicaid Services, 2023).

The section reports four sets of results. Section 10.1 develops the base case. Section 10.2 reports sensitivity to the plan parameters X and M_{base}. Section 10.3 reports adverse-scenario stress tests under unfavorable assumptions about the world. Section 10.4 reports design-feature scenarios that incorporate elements developed elsewhere in the paper but omitted from the base calculation. Section 10.5 states remaining limitations.

The four categories are kept distinct because they answer different questions. Parameter sensitivity asks what happens if plans choose differently. Stress tests ask what happens if the world is worse than the base case assumes. Design-feature scenarios ask what happens when mechanisms the paper develops are incorporated into the calculation. Limitations identify what the calculation cannot resolve.

The calculation is static. It does not model behavioral responses to point-of-service payment, provider pricing adjustments under increased price transparency, or the cost reductions plausibly attributable to preventive care utilization. Each of these would generally reduce government cost relative to the figures reported here.

10.1 Base Case

10.1.1 Cell Definitions, Repayment Horizons, and Ability-to-Pay Capacity

Households are partitioned into five income cells corresponding to approximate quintile boundaries, with median annual income and net wealth drawn from the Current Population Survey (U.S. Census Bureau, 2022) and the Survey of Consumer Finances (Bhutta et al., 2020) respectively. Discretionary monthly income is gross annual income minus fixed annual costs for housing, food, and transportation, divided by twelve. Table 1 summarizes the cells.

Table 1: Household cell definitions

Cell	Annual income (median)	Disc. income/mo	Net wealth	Pop. share
1 (Q1)	\$20,000	\$400	\$7,000	20%
2 (Q2)	\$42,000	\$1,000	\$43,000	20%
3 (Q3)	\$67,000	\$1,833	\$104,000	20%
4 (Q4)	\$105,000	\$3,750	\$280,000	20%
5 (Q5)	\$220,000	\$9,167	\$790,000	20%

Base-case plan parameters are $D = 10\%$, $X = 3\%$, $M_{\text{base}} = 24$, $M_{\min} = 6$, $W_c = \$50{,}000$, $P = 5\%$, and $\chi = 4$ months. Monthly ability-to-pay capacity is

$$\Pi_j = \frac{D \cdot I_{\text{discretionary},j} + X \cdot W_j}{2}.$$

The repayment horizon for a charge originating in cell j uses the exponential specification:

$$M(W_j) = M_{\min} + \left(M_{\text{base}} - M_{\min}\right)\left(1 - e^{-W_j/W_c}\right).$$

The single-charge recovery ceiling is

$$R_j^{\text{acute}} = \frac{M(W_j) \cdot \Pi_j}{1+P}.$$

Table 2 reports these quantities by cell.

Table 2: Monthly repayment capacity and single-charge recovery ceiling by cell

Cell	Π_j	$M(W_j)$	R_j^{acute}
1 (Q1)	\$125	8.35	\$994
2 (Q2)	\$695	16.38	\$10,844
3 (Q3)	\$1,652	21.75	\$34,215
4 (Q4)	\$4,388	23.93	\$100,008
5 (Q5)	\$12,308	24.00	\$281,334

10.1.2 High-Cost Trajectory Modeling

High-cost patients incur charges across multiple clinical episodes. The fiscal calculation models high-cost trajectories as three equal charges spaced four months apart, calibrated to MEPS high-utilization episodes. The spacing $\delta = 4$ months coincides with the base-case clustering window $\chi = 4$, placing consecutive charges at the cluster boundary. This is the most fiscally adverse spacing under both rules: it maximizes the chain's extended union of horizons under transitive clustering and produces the largest pairwise overlap under sliding-endpoint pairwise. Trajectories with $\delta < \chi$ produce shorter chain spans and smaller per-cluster overhang. Trajectories with $\delta > \chi$ produce no clustering. The $\delta = \chi$ choice is conservative for fiscal exposure.

Under both rules with $\delta = \chi = 4$, the three charges form a connected chain spanning months 1 through $M(W_j) + 2\delta$. Monthly capacity Π_j is distributed pro-rata among active charges in any given month. The rules differ only in which charges are considered active members of the cluster.

Under transitive clustering, the per-charge ceilings for the three positions are:

$$R_j^{\text{TR},(1)} = R_j^{\text{TR},(3)} = \frac{\Pi_j[\delta + \delta/2 + (M(W_j) - 2\delta)/3]}{1+P}$$

$$R_j^{\text{TR},(2)} = \frac{\Pi_j[\delta/2 + (M(W_j) - 2\delta)/3 + \delta/2]}{1+P}$$

The first and third charges receive equal allocation by symmetry. Each retains δ solo months at the chain boundary, δ pair months when one neighbor is active, and $M(W_j) - 2\delta$ triple months when all three are active. The middle charge never has solo months and receives less.

Under sliding-endpoint pairwise, the cluster contains only the most recent charge and charges within χ months of it:

$$R_j^{\mathrm{SE},(1)} = \frac{\Pi_j[\delta + \delta/2]}{1+P}$$
$$R_j^{\mathrm{SE},(2)} = \frac{\Pi_j \cdot M(W_j)/2}{1+P}$$
$$R_j^{\mathrm{SE},(3)} = \frac{\Pi_j[(M(W_j) - 2\delta)/2 + \delta/2 + \delta]}{1+P}$$

The first charge receives solo months at the start and pair months with the second charge, then drops out of the cluster when the third charge arrives. The third charge enters in a pair, remains in the cluster through its full horizon, and receives solo months at the chain's tail.

Table 3 reports per-charge ceilings by cell.

Table 3: Per-charge recovery ceilings, base-case parameters.

	Single	Transitive clustering			Sliding-endpoint			Chain
Cell	charge	First	Middle	Last	First	Middle	Last	total
Q1	\$994	\$728	\$490	\$728	\$714	\$497	\$735	\$1,946
Q2	\$10,844	\$5,821	\$4,497	\$5,821	\$3,971	\$5,422	\$6,746	\$16,139
Q3	\$34,215	\$16,648	\$13,502	\$16,648	\$9,438	\$17,107	\$20,253	\$46,798
Q4	\$100,008	\$47,264	\$38,907	\$47,264	\$25,071	\$50,004	\$58,361	\$133,435
Q5	\$281,334	\$132,852	\$109,408	\$132,852	\$70,333	\$140,667	\$164,111	\$375,112

Note. The single-charge column applies to buckets 1–3 (one charge of cost C_b). The chain columns apply to buckets 4–6 (three charges of cost $C_b/3$ spaced four months apart). The two scenarios apply to different bucket sizes, not alternative representations of the same total cost. Chain totals match across rules because monthly capacity Π_j is fixed and the chain spans $M(W_j) + 2\delta$ months under both rules.

Chain totals are identical across rules. Both rules deliver Π_j per month to the cluster whenever any charge is active. Total chain capacity equals Π_j times the number of months at least one charge is active, which is $M(W_j) + 2\delta$ for a three-charge chain spaced δ apart. The rules distribute this fixed total differently across charges.

The Q1 cell sits at the formula boundary, with $M(W_j) = 8.35$ just above $2\delta = 8$. Q1 entries should be read as boundary-regime values rather than as on-the-formula values.

10.1.3 Medical Cost Distribution

Annual per-person healthcare spending is approximated using six cost buckets calibrated to published Medical Expenditure Panel Survey (MEPS, Agency for Healthcare Research and Quality, 2020) percentiles, with a population mean of approximately \$6,876. The distribution is heavily right-skewed: the top two buckets account for roughly 5 percent of individuals but approximately 55 percent of total spending, consistent with published MEPS summary statistics.

The charge structure follows from the cost distribution. Buckets 1–3 represent routine and moderately acute episodes and are modeled as single financed charges, using R_j^{acute} as

Table 4: Stylized annual medical cost distribution

Bucket	Annual cost	Probability
1	\$200	0.40
2	\$1,200	0.28
3	\$5,000	0.20
4	\$18,000	0.07
5	\$60,000	0.04
6	\$180,000	0.01

the recovery ceiling. Buckets 4–6 represent high-cost trajectories consistent with multiple clinical episodes; each is modeled as three equal charges within a χ-window, with per-charge ceilings $R_j^{(k)}$ for $k = 1, 2, 3$ specified by the active allocation rule.

This bucket approximation applies a uniform cost distribution across income cells. In practice, healthcare spending correlates positively with income through age composition and chronic disease prevalence. The uniform distribution overstates spending in low-wealth cells, which drive most backstop activation. A cell-specific distribution calibrated to empirical income-spending correlations would reduce aggregate government cost below the figures reported here. The uniform distribution is retained as the conservative baseline.

10.1.4 Enrollment Fee Revenue

The backstop formula in Section 5.9 subtracts net insurer fee surplus before backstop activation. The loss calculations above account for origination fees through the $(1 - P)$ factor but omit enrollment fee revenue. This subsection adds it.

Enrollment fee revenue depends on the distribution of patients across fee tiers, which is governed by four offsetting forces. *Chronic burden* pulls high-utilization patients toward favorable repayment terms (which are typically attached to higher-fee plans), and chronic prevalence is higher at lower incomes (Boersma et al., 2020), so this force concentrates high-fee enrollment at the bottom. *Wealth* pulls high-wealth patients toward the same favorable terms because they treat enrollment fees as immaterial, concentrating high-fee enrollment at the top. *Fee avoidance* pulls patients across all incomes toward low-fee plans regardless of health status. *Rule sorting*, identified in Section 6.3, pulls chronic patients toward sliding-endpoint plans, which insurers price through whichever plan terms the selection exposure requires, including potentially higher enrollment fees. The first three forces operate in offsetting directions; the fourth compounds with chronic burden. The net distribution is not well-identified from available data.

The calculation uses a uniform distribution across fee tiers, one-third of enrollees in each tier in every cell, as a neutral baseline. Enrollment fees are paid at plan entry and amortized over an assumed mean enrollment interval of five years to convert them to an annual flow. Under a uniform distribution, each enrollee is equally likely to hold a low-fee (\$1,000), mid-fee (\$5,000), or high-fee (\$10,000) plan, yielding a mean enrollment fee of \$5,333 and an annualized fee revenue of \$1,067 per person, applied uniformly across all cells.

10.1.5 Expected Government Cost Under Each Rule

The dual-rule design creates a calibration question that this calculation addresses through bounding rather than estimation. Patient sorting across rules under continuous enrollment cannot be predicted without empirical data on a system that does not yet exist. Rather than guess the realized distribution, the calculation reports aggregate exposure under two pure-case scenarios: every plan uses transitive clustering, and every plan uses sliding-endpoint pairwise. The pure cases bound aggregate fiscal exposure regardless of how patients sort.

For each cell j and each rule $\xi \in \{\mathrm{TR}, \mathrm{SE}\}$, the expected insurer loss per person is

$$\mathbb{E}[L_j^\xi] = \sum_{b=1}^{3} \max\big(0,\, C_b - R_j^{\text{acute}}\big)\ \Pr(C_b) + \sum_{b=4}^{6} \left[\sum_{k=1}^{3} \max\Big(0,\, \tfrac{C_b}{3} - R_j^{\xi,(k)}\Big) \right] \Pr(C_b).$$

Applying the backstop formula, the gross government cost per person in cell j before enrollment fee offsets is $(1-P) \cdot \mathbb{E}[L_j^\xi]$. Table 5 reports the result by cell.

Table 5: Gross government cost per person by cell, both rules, base-case parameters

Cell	Gross govt cost (transitive)	Gross govt cost (sliding-endpoint)
Q1	\$5,781	\$5,781
Q2	\$3,347	\$3,397
Q3	\$1,767	\$1,777
Q4	\$442	\$442
Q5	\$0	\$0
Mean (population-weighted)	\$2,267	\$2,279

The two rules produce nearly identical gross exposure, with population-mean difference of \$12 per person, less than 1 percent of either figure. This reflects the chain-capacity identity above: aggregate chain economics are unchanged across rules, and per-charge differences matter only when redistribution moves a charge across the loss-clipping threshold.

10.1.6 Net Government Cost: Aggregate Bounds

The backstop formula applies the $\max(0, \cdot)$ operator at the aggregate level. Fee surplus generated by healthy enrollees offsets losses generated by chronically ill enrollees before any government payment is owed. Aggregate net government cost per person under rule ξ is

$$\text{Net cost}^\xi = \max\Big(0,\, \overline{(1-P) \cdot L^\xi} - \overline{F}\Big).$$

Applying this to the figures in Table 5 and the \$1,067 fee offset:

$$\text{Net cost}^{\mathrm{TR}} = \max(0,\ \$2{,}267 - \$1{,}067) = \$1{,}200$$
$$\text{Net cost}^{\mathrm{SE}} = \max(0,\ \$2{,}279 - \$1{,}067) = \$1{,}212.$$

Multiplying per-person net cost by the US population of approximately 333 million yields aggregate net government cost in the range \$400–\$404 billion, approximately 19 percent of current combined federal and state public health expenditure of \$2.1 trillion (Centers for Medicare & Medicaid Services, 2023).

10.2 Parameter Sensitivity

Plan parameters X and M_{base} are chosen by insurers under competitive pressure. This subsection reports how aggregate exposure varies across plausible parameter values. Table 6 reports population-weighted mean net government cost per person across the parameter grid for both allocation rules.

Table 6: Population-weighted mean net government cost per person (\$) under both allocation rules, varying X and M_{base}, $D = 10\%$, $\chi = 4$.

Panel A: Transitive clustering

	M_{base}			
X	12	24	36	48
1%	\$2,686	\$2,168	\$1,812	\$1,662
3%	\$1,578	\$1,200	\$907	\$803
5%	\$1,097	\$779	\$640	\$501
7%	\$807	\$582	\$390	\$207

Panel B: Sliding-endpoint pairwise

	M_{base}			
X	12	24	36	48
1%	\$2,686	\$2,196	\$1,986	\$1,845
3%	\$1,578	\$1,212	\$1,064	\$977
5%	\$1,094	\$850	\$711	\$572
7%	\$786	\$588	\$406	\$336

Three findings carry across both panels. First, X is the primary fiscal lever: moving from $X = 1\%$ to $X = 7\%$ at $M_{\text{base}} = 24$ reduces net government cost by roughly \$1,600 per person under either rule, compared with a smaller reduction from extending M_{base} from 12 to 48 at fixed X. Second, the differential between rules is small in absolute terms: maximum spread across the grid is approximately \$200 per person at the most generous parameter combinations. Third, the rule differential is monotone in horizon length: it vanishes at $M_{\text{base}} = 12$ and grows at longer horizons.

The narrow gap between rules supports a backstop reimbursement rate ρ_{sep} near 1 across the parameter range examined.

10.3 Adverse-Scenario Stress Tests

This subsection reports aggregate net government cost under five adverse scenarios and one combined scenario. Each scenario perturbs one base-case assumption in the unfavorable direction. The combined scenario applies three perturbations simultaneously. All figures use base-case plan parameters and the transitive clustering rule; sliding-endpoint figures are within $30 per person of the transitive figures in every scenario.

Scenario A: lower recovery. The recovery ceilings R_j^{acute} and $R_j^{(k)}$ are reduced by 25 percent across all cells. This models realized capacity below verified capacity due to verification lag, partial default, or non-payment despite ability to pay. Recovery ceilings in the base case assume full collection against allocated capacity.

Scenario B: higher utilization. The probability mass on buckets 4–6 increases by 50 percent, with buckets 1–3 rescaled to preserve total probability. This models population health deterioration or shifts in clinical practice that increase the frequency of high-cost events. Aggregate medical spending in this scenario is approximately 40 percent above the base case.

Scenario C: lower fee revenue. The enrollment fee offset is reduced from $1,067 to $533 per person. This models shorter mean enrollment tenure, lower realized fee tiers, or greater fee avoidance than the uniform-distribution baseline assumes.

Scenario D: no fee pooling ($\beta = 0$). The backstop formula is applied without subtracting fee surplus, as developed in Section 5.10. Public reimbursement activates on every dollar of insurer loss net of origination fees, with fee surplus retained by insurers as profit. This is the philosophically cleanest implementation of the redistribution principle and is the most consequential design parameter the regulator can adjust.

Scenario E: adverse provider price response. The cost distribution is inflated by 15 percent across all buckets. This models providers raising prices to absorb patient repayment capacity, the opposite of the price-compression dynamic the base case omits. The 15 percent figure is intentionally large given that providers face increased price transparency under the design; it serves as a stress floor.

Combined scenario: B + C + E. Higher utilization, lower fee revenue, and adverse provider pricing applied simultaneously. This scenario assumes correlated adverse outcomes across unrelated channels and represents a stress floor rather than a central estimate.

Table 7 reports results.

Four observations.

First, every individual adverse scenario remains below 40 percent of current combined federal and state public health expenditure. The worst single scenario is Scenario D at 36.0 percent.

Table 7: Adverse-scenario stress tests, transitive clustering rule, base-case plan parameters

Scenario	Description	Net/person	Aggregate	% of $2.1T
Base	–	$1,200	$400B	19.0%
A	Lower recovery (-25%)	$1,501	$500B	23.8%
B	Higher utilization (+50% high-bucket)	$2,242	$746B	35.5%
C	Lower fee revenue (halved)	$1,734	$577B	27.5%
D	No fee pooling ($\beta = 0$)	$2,267	$755B	36.0%
E	Provider prices (+15%)	$1,721	$573B	27.3%
B+C+E	Combined stress floor	$3,534	$1,177B	56.0%

Second, the combined stress floor remains below current public spending. At 56 percent of $2.1 trillion, the system absorbs simultaneous adverse outcomes on utilization, fee revenue, and provider pricing while still costing less than existing programs.

Third, Scenarios B and E produce similar aggregate exposure through different mechanisms. Both raise expected loss in the right tail of the cost distribution, where the recovery ceilings of low-wealth cells bind most tightly. Most aggregate loss is concentrated in this tail under any parameterization that varies cost or utilization, so the choice between the two mechanisms matters less than the magnitude of the perturbation.

Fourth, Scenario D produces the largest exposure among single-channel adverse scenarios. The fee-pooling design choice has roughly twice the fiscal implication of any of the world-state perturbations tested. This supports the characterization of β in Section 5.10 as a substantive regulatory instrument rather than a technical accounting convention.

The stress test does not exhaust adverse scenarios. The cell-specific cost distribution noted as conservative in the base case operates in the opposite direction and is not credited here. Behavioral responses to point-of-service payment, provider pricing adjustments under increased price transparency, and preventive care utilization effects would generally reduce exposure in all scenarios. The figures reported here are upper bounds in this restricted sense.

10.4 Design-Feature Scenarios

The base case omits design features developed elsewhere in the paper. This subsection incorporates one such feature: the Tax-Free TRUST-CARE Repayment Account (TFRA) developed in Section 8. TFRA balances are drawn down to pay TRUST-CARE obligations directly, reducing the financed amount on each charge and the associated insurer loss.

The static calculation cannot model lifecycle dynamics of contribution and drawdown. It assumes an average steady-state TFRA balance available at the time of a covered event, applied as direct reduction of bucket cost. The financed amount for bucket b becomes $\max(0, C_b - T)$ where T is the average available balance.

The TFRA balances assumed here are stylized estimates. They reflect a working-age population accumulating health-restricted capital across decades of wage unbundling. Realized balances depend on contribution rates, drawdown patterns, and lifecycle position that the design does not constrain ex ante. The scenarios demonstrate the directional effect and order of magnitude, not point estimates.

Table 8 reports results.

Table 8: TFRA drawdown scenarios, transitive clustering rule, base-case plan parameters

Avg TFRA balance	Net/person	Aggregate	% of \$2.1T
\$0 (base case)	\$1,200	\$400B	19.0%
\$5,000	\$794	\$264B	12.6%
\$10,000	\$576	\$192B	9.1%
\$15,000	\$382	\$127B	6.1%

A \$10,000 average TFRA balance cuts aggregate net cost roughly in half. The mechanism is mechanical: TFRA reduces the financed amount before the recovery ceiling applies, and most loss is concentrated on charges that exceed recovery ceilings, so reducing the financed amount has high leverage on aggregate exposure. A \$5,000 balance is sufficient to absorb most of buckets 1–3 entirely, eliminating the small-charge contribution to backstop activation that the base case retains.

The drawdown channel is distinct from the wealth-accumulation channel that TFRA balances would create if they entered $A_{\text{investments}}(t)$ in the wealth formula. The current paper does not specify TFRA balances as part of $W(t)$, so the design-feature scenario here captures only the direct drawdown effect. The wealth channel, if added, would further reduce backstop activation by raising Π_t and $M(W)$ for patients with TFRA accumulation.

10.5 Scope and Limitations

Five qualifications apply.

First, the high-cost charge model uses a stylized three-charge trajectory with equal four-month spacing. Empirical charge arrival distributions from MEPS episode data would refine this. The choice $\delta = \chi$ is conservative for fiscal exposure.

Second, the calculation is cross-sectional. It does not model capacity allocation dynamics across overlapping charges from multiple years.

Third, the calculation reports figures under each pure-case rule scenario and bounds aggregate realized exposure between them. It does not model the realized distribution of patients across rules. The Section 6.3 analysis indicates chronic patients will sort toward sliding-endpoint plans, but the aggregate fiscal envelope established by the bounds applies regardless of sorting.

Fourth, the enrollment fee revenue calculation uses a five-year mean enrollment interval that is a stylized assumption with no direct empirical basis. It is conservative relative to typical employer-sponsored plan tenure, and it excludes cancellation fee revenue. The net figures should be read as illustrative of the direction and order of magnitude of the fee offset, not as point estimates.

Fifth, the design-feature scenarios in Section 10.4 assume stylized average TFRA balances. Lifecycle dynamics of contribution, drawdown, and inheritance under steady-state operation would require multi-decade simulation that the static framework cannot support.

The bounding result, aggregate per-person net government cost between \$1,200 and \$1,212 under base-case parameters scaling to \$400–\$404 billion at US population, is robust to the dominant sources of uncertainty identified above. Under individual adverse scenarios, exposure rises to at most 36 percent of current public spending. Under a combined stress floor assuming correlated adverse outcomes, exposure reaches 56 percent of current public spending. Under TFRA drawdown scenarios incorporating a design feature the base case omits, exposure falls to between 6 and 13 percent. Subject to the limitations above, aggregate fiscal exposure is substantially below current combined federal and state public health expenditure across the full range of scenarios examined.

11 Comparison to Alternative Systems

This section compares TRUST-CARE to the two dominant alternatives: premium-based insurance with regulatory patches (the ACA model) and single-payer financing. The comparison follows the diagnostic structure of Section 2. Each of the four failures identified there is addressed in turn, followed by administrative complexity and political-economy properties.

11.1 Price Signals at the Point of Service

Premium-based insurance severs the price signal. Patients facing copayments observe out-of-pocket amounts, not underlying prices. Providers negotiate rates with insurers and do not compete on posted prices visible to patients. Insurers and providers jointly set effective prices through network design and benefit tiers. No participant has both the information and the incentive to discipline prices.

Single-payer systems compress prices through monopsony bargaining rather than through patient-level price signals. The compression is real but operates as administrative price-setting. When reimbursement falls below market-clearing levels, the system produces rationing through waiting times and provider supply constraints. The price signal at the point of service is suppressed entirely.

TRUST-CARE restores the price signal for routine care. Patients pay directly at the point of service and observe prices before incurring the obligation. Mandatory price posting (Section 4) supplies the transparency requirement. Direct payment makes the patient the marginal payer. For high-cost care, the retrospective price dispersion surcharge prices opacity in provider billing behavior without imposing fee schedules.

11.2 Risk Pooling and Redistribution

Premium-based pooling requires healthy enrollees to subsidize sick enrollees through voluntary premium payments. This generates adverse selection, necessitating community rating, individual mandates, risk adjustment, and open enrollment windows. Redistribution is real but opaque, fiscally unauditable, and tied to employment status and plan choice rather than ability to pay. Expiring ACA subsidies are the current instance of this instability.

Single-payer systems achieve redistribution through general taxation. This is fiscally transparent but requires finanzing the full medical economy which constrains both program design and revenue flexibility. The redistribution is divorced from utilization, so individual contributions bear no relationship to individual consumption.

TRUST-CARE relocates redistribution to the government backstop (Section 5). Risk pooling occurs through fiscal transfer calibrated to insurer loss experience and funded through general taxation. The β parameter (Section 5.10) tunes the share of redistribution operating within insurer fee pools versus through public taxation. $\beta = 0$ recovers the single-payer property of taxation-only redistribution; $\beta = 1$ retains within-pool pooling as a defense against soft selection through plan design. The choice is a regulatory instrument rather than a fixed feature of the model.

The redistribution is explicit in two senses. It is fiscally auditable: the backstop payment formula in Section 5.9 is a parameterized function of recognized losses, origination fees, and fee surplus. It is stable against adverse selection: a patient who never uses the credit system contributes through taxes, not through a premium they can decline by leaving the pool.

11.3 Cost-Sharing Discrimination

Cost-sharing in premium-based systems operates through deductibles, copayments, and coinsurance. Brot-Goldberg et al. (2017) document that high-deductible plans reduce high-value and low-value care in roughly equal proportions. The deductible cliff eliminates marginal cost sensitivity once the out-of-pocket maximum is reached. Value-based insurance design (Chernew et al., 2007) addresses this by varying cost-sharing across services according to clinical value, but remains within the premium-based architecture.

Single-payer systems eliminate point-of-service cost-sharing for covered services. This avoids the discrimination problem by removing the instrument entirely, at the cost of removing patient-level price sensitivity for both high-value and low-value care.

TRUST-CARE addresses cost-sharing discrimination through two mechanisms. The preventive care safe harbor (Section 9) protects evidence-based preventive services from financing friction. Safe-harbor services carry zero origination fee, favorable repayment parameters, and first-dollar credit access. The categorical protection avoids the indiscriminate utilization reductions documented in high-deductible plan research. The Backstop-Eligible Care standard (Section 5.11) performs an analogous function for catastrophic care: clinically necessary services receive public loss sharing; preference-driven services do not. The two mechanisms together produce graduated protection that discriminates by clinical value rather than by uniform cost-sharing thresholds.

11.4 Catastrophic Debt Resolution

Premium-based systems have no structured mechanism for resolving catastrophic medical debt. Current resolution paths are bankruptcy, charity care, hospital write-offs negotiated case by case, accumulation, and Medicaid spend-down. Himmelstein et al. (2009) and Dobkin et al. (2018) document the financial consequences. The public ultimately absorbs much of the cost through uncompensated care subsidies, Medicaid spend-down, and bankruptcy externalities, but does so through the least efficient channels available.

Single-payer systems eliminate catastrophic medical debt by eliminating patient-level financing. The cost is fully socialized at the point of service. This is structurally clean but forgoes any patient-level information about willingness or ability to contribute to care costs.

TRUST-CARE supplies a structured debt resolution mechanism. The credit instrument (Section 3) finances charges that exceed immediate ability to pay, with repayment indexed to verified wealth and income through Π_t and forgiveness horizon $M(W)$. The charge-local amortization accounting (Section 5) recognizes losses progressively as recovery capacity becomes infeasible. The government backstop absorbs residual losses. The patient faces a finite, formulaic, ability-indexed obligation rather than open-ended debt resolved through bankruptcy. The forgiveness is not a default; it is a designed termination of the obligation at a horizon set by the plan and known at origination.

This is the failure most distinctive to TRUST-CARE. Premium-based and single-payer systems handle catastrophic financing either through punitive unstructured channels or through full socialization. TRUST-CARE supplies a third option: structured patient-level obligation calibrated to means, with explicit public absorption of unrecoverable cost.

11.5 Administrative Complexity

A potential objection to TRUST-CARE is that it appears administratively complex relative to single-payer alternatives. The relevant comparison is the location of complexity rather than its presence. Healthcare financing necessarily entails continuous adjudication of heterogeneous services, evolving clinical practice, uncertain costs, patient heterogeneity, and strategic behavior. Any universal system must process this information.

Premium-based systems distribute complexity across insurers, providers, and regulators. Each insurer maintains separate networks, benefit structures, formularies, and claims processing. Providers maintain billing infrastructure for multiple payers. Regulators maintain rules for community rating, risk adjustment, and exchange operation. The aggregate administrative cost is high (Anderson et al., 2019), and the complexity creates friction at every interaction.

Single-payer systems concentrate complexity in a public authority that defines coverage rules, sets or negotiates prices across a large and changing service space, manages exceptions and appeals, and revises policy in response to technological change. Centralization reduces visible contractual variety but creates a single bottleneck for updating rules and correcting misestimation. When a centralized authority misprices or misclassifies a service, the error propagates system-wide and is often slow to correct because revisions are politically salient.

TRUST-CARE decomposes complexity into modular components distributed across competing insurers, while confining government's role to a parameterized backstop and selective auditor. Contract terms are public and formulaic, allowing fiscal exposure to be modeled explicitly. Plan competition operates as a parallel search process over repayment schedules. Misestimation remains local to particular plans and can be corrected through contract revision and competitive pressure.

The verification infrastructure required to implement ability-to-pay assessment at the frequencies the design specifies does not currently exist in the United States. The tiered verification schedule (Section 5.7) concentrates infrastructure demands on patients with material outstanding balances, reducing the implementation threshold relative to uniform monthly verification. This is the design's response to the verification feasibility question; the question remains an empirical one.

11.6 Political Economy

The three systems differ in how stable their financing arrangements are against political pressure.

Premium-based systems with regulatory patches are unstable because the redistribution is embedded in instruments that cannot bear it. Premium subsidies require annual appropriation and are subject to repeated renegotiation. Risk adjustment requires technical rule-making that is opaque to the public but salient to industry. The patches accumulate and require continuous maintenance.

Single-payer systems are politically stable once enacted but face high enactment thresholds because they require simultaneous transition of providers, patients, and revenue streams. The single-bottleneck property that concentrates operational complexity also concentrates political risk: revisions are high-stakes and infrequent.

TRUST-CARE separates the redistribution decision from the financing architecture. The backstop is the redistribution; plan competition is the financing. The redistribution decision is fiscally explicit and reducible to a small number of parameters ($\rho_{\text{trans}}, \rho_{\text{sep}}, \beta$, M_{min}, the safe-harbor list, the BEC standard). The financing architecture operates through plan competition under standardized rules. The two layers can be adjusted independently. A future administration that wished to expand or contract public exposure could do so through the backstop parameters without restructuring the underlying financing.

This separation also addresses the employment-lock problem. Wage unbundling (Section 8) converts employer premium contributions into wages, with the TFRA available as a voluntary buffer. Coverage detaches from employment status without mandates or income-threshold subsidies. The political economy of employer-sponsored insurance, which has stabilized the current system against major reform for decades, is replaced by individual ownership of health-restricted capital.

The model is not presented as a fully specified legislative proposal. It is a financing architecture. Its purpose is to identify a coherent alternative design space in which universal access, patient agency, explicit redistribution, price visibility, and structured medical-debt forgiveness can be made jointly compatible. The empirical question is not whether every parameter chosen here is correct, but whether the architecture defines a feasible class of systems worth simulating, piloting, and refining.

References

Anderson, G. F., Hussey, P., and Petrosyan, V. (2019). It's still the prices, stupid: Why the US spends so much on health care, and a tribute to Uwe Reinhardt. *Health Affairs*, 38(1), 87–95.

Brot-Goldberg, Z. C., Chandra, A., Handel, B. R., and Kolstad, J. T. (2017). What does a deductible do? The impact of cost-sharing on health care prices, quantities, and spending dynamics. *Quarterly Journal of Economics*, 132(3), 1261–1318.

Chernew, M. E., Rosen, A. B., and Fendrick, A. M. (2007). Value-based insurance design. *Health Affairs*, 26(2), w195–w203. doi:10.1377/hlthaff.26.2.w195

Enthoven, A. C. (1993). The history and principles of managed competition. *Health Affairs*, 12(Suppl 1), 24–48.

Gruber, J. (2008). Covering the uninsured in the U.S. *Journal of Economic Literature*, 46(3), 571–606.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7), 2643–2682.

Himmelstein, D. U., Thorne, D., Warren, E., and Woolhandler, S. (2009). Medical bankruptcy in the United States, 2007: Results of a national study. *The American Journal of Medicine*, 122(8), 741–746. doi:10.1016/j.amjmed.2009.04.012

Newhouse, J. P. and the Insurance Experiment Group (1993). *Free for All? Lessons from the RAND Health Insurance Experiment*. Harvard University Press, Cambridge, MA.

Papanicolas, I., Woskie, L. R., and Jha, A. K. (2018). Health care spending in the United States and other high-income countries. *JAMA*, 319(10), 1024–1039.

Rothschild, M. and Stiglitz, J. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics*, 90(4), 629–649.

Tebaldi, P. (2022). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the ACA. *Review of Economic Studies*, 89(4), 1471–1518.

Whaley, C. M. (2019). Searching for health: The effects of online price transparency. *Journal of Health Economics*, 67, 102200.

U.S. Census Bureau (2022). Current Population Survey, Annual Social and Economic Supplement. U.S. Census Bureau, Washington, DC. Available at `https://www.census.gov/data/datasets/2022/demo/cps/cps-asec-2022.html`.

Bhutta, N., Bricker, J., Chang, A. C., Dettling, L. J., Goodman, S., Hsu, J. W., Moore, K. B., Reber, S., Henriques Volz, A., and Windle, R. A. (2020). Changes in U.S. family finances from 2016 to 2019: Evidence from the Survey of Consumer Finances. *Federal Reserve Bulletin*, 106(5), 1–42. doi:10.17016/bulletin.2020.106

Boersma, P., Black, L. I., and Ward, B. W. (2020). Prevalence of multiple chronic conditions among US adults, 2018. *Preventing Chronic Disease*, 17, 200130. doi:10.5888/pcd17.200130

Dobkin, C., Finkelstein, A., Kluender, R., and Notowidigdo, M. J. (2018). The economic consequences of hospital admissions. *American Economic Review*, 108(2), 308–352. doi:10.1257/aer.20161038

Dranove, D. and Millenson, M. L. (2006). Medical bankruptcy: Myth versus fact. *Health Affairs*, 25(2), w74–w83. doi:10.1377/hlthaff.25.w74

Centers for Medicare & Medicaid Services (2023). National Health Expenditures 2023 Highlights. Available at `https://www.cms.gov/files/document/highlights.pdf`.

Agency for Healthcare Research and Quality (2020). Medical Expenditure Panel Survey, Household Component Summary Tables. Rockville, MD. Available at `https://meps.ahrq.gov/mepsweb/data_stats/quick_tables.jsp`.