WINPEC Working Paper Series No. E26010
July 2026

Anticipation and Adaptation: Cost-Sharing Responses and Mortality Among the Oldest-Old

Rong Fu Masato Oikawa Akira Kawamura
Haruko Noguchi

Waseda INstitute of Political EConomy
Waseda University
Tokyo, Japan

Anticipation and Adaptation: Cost-Sharing Responses and Mortality Among the Oldest-Old

Rong Fu* Masato Oikawa† Akira Kawamura‡ Haruko Noguchi§

July 16, 2026

Abstract

We study a Japanese reform that doubled coinsurance rates for higher-income enrollees aged 75+, using administrative data linking insurance claims to tax records for nearly the entire oldest-old population over 26 months. Patients exhibit a last-minute anticipatory surge, then persistently reduce utilization by cutting visit frequency rather than treatment intensity. Responses follow a gradient of service discretionarity, and a binding expenditure cap shields the most medically intensive patients from the reform's financial impact. Despite utilization reductions, we find no detectable adverse effect on 15-month mortality. Cost-sharing increases, paired with effective safety nets, can restrain utilization without detectable mortality consequences.

Keywords: cost-sharing; healthcare utilization; oldest-old; value-based insurance design

JEL Codes: I11; I13; J14; H51

*Corresponding author. Faculty of Commerce, Waseda University, and Waseda Institute of Social & Human Capital Studies (WISH), Tokyo, Japan. Email: nataliefu@waseda.jp. This research received support from the Health Labour Sciences Research Grant (Research on Policy Planning and Evaluation) provided by the Ministry of Health, Labour, and Welfare, under Grant Number 22AA1002, with Haruko Noguchi serving as the Principal Investigator. This study utilized anonymized secondary data, for which the Ethics Review Committee for Research Involving Human Subjects at Waseda University waived the informed consent requirement (Approval No.: 2022-HN038).

†Faculty of Education and Integrated Arts and Sciences, Waseda University, and WISH. Email: m.oikawa@waseda.jp

‡Faculty of Human Sciences, Waseda University, Saitama, Japan, and WISH. Email: akawamura@waseda.jp

§Faculty of Political Science and Economics, Waseda University, Tokyo, Japan, and WISH. Email: h.noguchi@waseda.jp

1 Introduction

Understanding how individuals anticipate and adapt to price changes has long been central to economics, with evidence spanning retail purchasing (Hendel and Nevo, 2004, 2006), housing transactions (Best and Kleven, 2018), durable goods programs (Mian and Sufi, 2012), and fiscal policy reforms (Poterba, 1988; Goolsbee, 2000). Healthcare offers a particularly well-suited setting for studying these mechanisms: the deferability of care varies sharply across service types, from life-threatening emergencies that cannot wait to elective procedures that patients can readily reschedule, providing rich within-domain variation in how much room individuals have to time their consumption around anticipated price changes (Manning et al., 1987; Chandra et al., 2010; Duarte, 2012; Einav et al., 2018). Yet the existing literature has focused overwhelmingly on working-age and younger elderly populations, leaving open three questions: how the oldest-old—the fastest-growing demographic globally and the heaviest users of healthcare—respond to cost increases, which margins of adjustment they employ, and whether the resulting reductions in care harm their health.

Healthcare cost-sharing theory, established by Arrow (1963) and Grossman (1972), demonstrates how price signals influence consumption decisions. The RAND Health Insurance Experiment found healthcare demand price elasticities of approximately -0.2 (Manning et al., 1987; Newhouse, 1993), with subsequent research revealing substantial heterogeneity across populations and contexts (Aron-Dine et al., 2013). Building on this foundation, the literature has documented behavioral responses along three dimensions: (i) anticipatory behavior where individuals strategically time healthcare consumption around policy changes (Card et al., 2008; Einav et al., 2015; Alpert, 2016), (ii) adaptation pathways as individuals adjust to new cost regimes following implementation (Brot-Goldberg et al., 2017), and (iii) systematic heterogeneity across service types, with essential services showing lower price elasticity than discretionary care (Chandra et al., 2010; Duarte, 2012). This last pattern also provides empirical support for value-based insurance design, which holds that cost-sharing should be calibrated to the clinical value of each service rather than applied uniformly.[1] (Chernew et al., 2007; Fendrick et al., 2012; Baicker and Levy, 2015; Gruber et al., 2020).

These strands of research, however, leave a coherent set of open questions. First, existing studies face two interrelated methodological limitations that prevent a complete picture of behavioral responses. Most examine either anticipatory behavior (Einav et al., 2015) or post-implementation adaptation (Brot-Goldberg et al., 2017) in isolation, capturing only part of the behavioral arc—yet the full trajectory carries distinct implications: a sharp change that gradually attenuates suggests a temporary adjustment; a persistent stable shift implies a genuine new utilization equilibrium;

[1] Value-based insurance design principles recognize that medical services differ in their contribution to health, with clinical benefits that vary by patient, timing, and context (Chernew et al., 2007; Fendrick et al., 2012). Inpatient and emergency care show low price sensitivity owing to their urgent nature (Ellis and McGuire, 2007; Chandra et al., 2010), dental care demonstrates markedly higher responsiveness (Manning et al., 1987; Newhouse, 1993), and pharmaceutical utilization reveals strategic patient responses to coverage changes (Einav et al., 2015; Abaluck et al., 2018).

and a progressively deepening response implies long-run adjustment remains incomplete. Clean identification compounds this challenge, as eligibility criteria are often measured with error in cost-sharing studies (Bound et al., 2001), making it difficult to isolate behavioral responses from correlated individual characteristics. Second, the oldest-old represent the fastest-growing yet least-studied demographic in this literature. Unlike younger cohorts who retain flexibility to substitute or delay care, the oldest-old face a qualitatively different clinical profile: frailty affects one in four, multimorbidity affects 70% of those over 80, and emergency utilization rises sharply with age (Fried et al., 2001; Clegg et al., 2013; Salive, 2013; Gruneir et al., 2011), all of which constrain the scope for discretionary reduction without health consequences. Whether and how this population responds to cost increases, and on which margins, remains incompletely understood. Third, direct evidence that cost-sharing harms health—the question that ultimately motivates the policy debate—remains scarce (Rice and Matsuoka, 2004; Chandra et al., 2010), and is especially absent where the stakes are highest: aging societies facing acute fiscal pressure (Ministry of Health, Labour and Welfare, 2022; Eurostat, 2024; OECD, 2023) are expanding cost-sharing precisely for populations whose health may be most vulnerable to such changes.

Japan's healthcare system provides an ideal context to address these gaps. Its comprehensive universal coverage, standardized fee schedules, and transparent cost-sharing structure minimize the provider and pricing variation that complicates identification in many other settings (Iizuka and Shigeoka, 2022). Its rapidly aging population further makes insights from this context directly relevant to other societies facing similar fiscal pressures: as of 2025, oldest-old adults aged 75 and above (75+) account for 17.2% of Japan's population yet 35.8% of national medical expenditure, so even modest behavioral responses to cost-sharing carry outsized budgetary consequences (Statistics Bureau, Ministry of Internal Affairs and Communications, 2025; Ministry of Health, Labour and Welfare, 2025). The existing Japanese literature has extensively examined behavioral responses to cost-sharing *reductions* (Ando and Takaku, 2016; Shigeoka, 2014; Fukushima et al., 2016; Nishi et al., 2012), but responses to *increases* remain far less studied, despite evidence of potential asymmetry (Iizuka and Shigeoka, 2023). The closest antecedent to our study is Komura and Bessho (2025), who examine coinsurance rate increases for patients aged 70–74 and document persistent utilization reductions without measurable adverse health impacts. Their analysis, however, relies on semi-aggregated data, which precludes the individual-level temporal analysis needed to trace the complete behavioral arc from anticipation through adaptation, and uses indirect proxies for health stock rather than direct outcomes, leaving open whether utilization reductions translate into adverse health consequences. Our study addresses these constraints and extends the setting in two further respects: individual-level monthly claims linked to income-tax records let us trace the full trajectory of behavioral response for essentially the entire 75+ population, and administrative panel attrition validated against official death records provides a direct individual-level mortality measure. Our setting also differs in two ways of independent interest. First, we study the oldest-old (75+) rather than the near-old (70–74)—a population with markedly higher frailty, multimorbidity, and essential-care needs, for whom the scope to absorb cost-sharing without harm is a priori far more limited and the health stakes correspondingly higher. Second, the reform we study pairs the coinsurance increase with a binding monthly expenditure cap (*Kōryō*, see Section 2), allowing us to identify how a safety net and a price increase jointly determine who bears the burden of a reform—an interaction that has received little attention in the literature

We study a policy implemented in Japan in October 2022 that doubled coinsurance rates from 10% to 20% for individuals aged 75+ with income above specified thresholds. Our central advantage is the data: we leverage a novel dataset linking national health insurance claims to individual income tax records for essentially the *entire* population aged 75+—over 18 million enrollees observed monthly across 26 months. Because eligibility turns on that threshold, our estimates identify local average treatment effects for the middle-income enrollees near the cutoff: a group of direct policy relevance, since it is precisely the population the reform moved to the higher rate. Crucially, eligibility is determined by each individual's income in the *previous* fiscal year, reported well before the policy was announced, rendering income manipulation essentially impossible. Furthermore, the reform was announced two months before implementation, creating a well-defined anticipation window during which affected enrollees knew their new rates but had not yet faced them (Figure 1). Near-perfect compliance at the eligibility cutoff further confirms that treatment assignment is exogenously determined: the probability of paying the 20% rate jumps from near zero to 1.0 at the threshold, approximating random assignment within the local neighborhood of the cutoff (Figure 2).

These two features—manipulation-proof income assignment and near-perfect institutional compliance —anchor our primary design: a difference-in-differences event study comparing enrollees just above and below the income threshold, before and after the reform, to trace the full temporal path of response that a static comparison would miss: pre-notification, anticipation, immediate adjustment, and steady state. The identifying assumption of parallel trends is supported by flat, precisely estimated pre-notification coefficients (Figure 3).[2] By tracing behavioral responses over 26 months and examining 15-month mortality outcomes, we provide a comprehensive picture of how the oldest-old navigate increased cost-sharing, what margins of adjustment they use, and whether care reductions carry detectable health consequences.

We document three sequential behavioral phases and a distinct margin of adjustment. First, a significant anticipatory last-minute surge (+5.9% in costs, +2.3% in utilization rates) concentrated entirely in the final month before implementation, with no detectable response in the earlier notification month—consistent with deadline salience rather than gradual intertemporal optimization.

[2]Cross-sectional regression discontinuity estimates at the threshold, which do not rely on parallel trends, reproduce the same dynamics as an independent check.

This concentration in a single month, rather than a response spread across the two-month notification window, is itself a distinct behavioral finding with direct implications for how advance-notification policies should be designed. Second, an immediate decrease (-6.3% in costs, -3.0% in utilization rates) upon implementation. Third, persistent but modest reductions ($\sim$3% in costs, $\sim$1% in utilization rates) throughout the subsequent year. The implied steady-state price elasticities (-0.03 to -0.06) are smaller than existing estimates for younger elderly populations (Shigeoka, 2014; Fukushima et al., 2016), reflecting both the more essential nature of care for the oldest-old (Clegg et al., 2013; Dent et al., 2019) and a methodological point of broader relevance: because the immediate response is more than twice the steady-state reduction, studies with shorter observation windows may conflate the two and overstate long-run price sensitivity. A decomposition of the intensive margin reveals visit frequency as the dominant adjustment mechanism: total spending moves almost one-for-one with days of service across the post-implementation period, while costs per day show only minimal change, confirming that patients reduce how often they seek care rather than how much care they receive per visit.

Responses vary systematically across service types, following a gradient aligned with the discretionary nature of each service. Dental care—the most discretionary category—exhibits the sharpest extensive margin response: visit rates surge in the final pre-implementation month and fall substantially afterward, while costs conditional on utilization show little persistent change, implying that patients adjust the visit-or-not decision rather than the scope of treatment within each visit. Outpatient and pharmaceutical services occupy an intermediate position, with persistent reductions across both margins that stabilize over the post-implementation period, indicating a new utilization equilibrium. Hospital admissions, by contrast, follow a qualitatively different trajectory: admission rates decline progressively before partially recovering, while inpatient costs conditional on admission remain essentially unchanged throughout, consistent with deferral of elective or semi-elective procedures rather than foregone acute care.

A parallel gradient emerges along another dimension: patients' underlying medical needs. Among tertiles defined by pre-reform utilization frequency, the margin of adjustment shifts with medical need: lower-need patients curtail visit frequency most sharply, while higher-need patients, with limited room to reduce visits further, show larger reductions in cost intensity. The highest-need group, patients covered by Japan's High-Cost Medical Expense Benefit (*Kōryō*, see Section 2), presents a striking contrast: despite having the greatest medical needs in the sample, they are the least price-sensitive of all. This is precisely what the institutional design predicts, because the *Kōryō* cap binds independently of the coinsurance rate, doubling the rate leaves their effective out-of-pocket costs largely unchanged. The interaction between a coinsurance rate increase and a binding expenditure cap, which together determine who bears the burden of a reform, has received little attention in the literature and carries a transferable implication for instrument design: cost-sharing policy cannot be evaluated by the rate alone, but only in combination with the safety nets that bound it.

Despite reductions in both outpatient visits and hospital admissions over the 15-month post-implementation period, we find no statistically detectable evidence of adverse effects on survival. The point estimate indicates that 15-month mortality is *lower* among those assigned the higher coinsurance rate by 0.2 percentage points, but the estimate is statistically insignificant. We read this result with caution: it provides no evidence that the reform harmed survival over the period we observe, but it is equally not positive evidence of safety, and the 15-month horizon leaves open the possibility of longer-run consequences. Interpreted alongside the service-type and diagnosis-level evidence—reductions concentrated in dental care and elective inpatient admissions rather than acute care—the mortality finding is consistent with, though it cannot by itself establish, the value-based insurance design hypothesis that the care foregone was disproportionately lower-value (Chernew et al., 2007; Fendrick et al., 2012).

The three sets of findings map directly onto the three questions posed above. On how the oldest-old respond, they do respond, but primarily through a last-minute deadline effect rather than gradual intertemporal optimization, and at a smaller steady-state magnitude than younger populations. On which margins, visit frequency dominates treatment intensity, and utilization reductions concentrate on discretionary and deferrable services while a binding expenditure cap shields the highest-need patients. On health consequences, we find no detectable adverse effect on 15-month mortality, consistent with the self-triage interpretation that the care reduced was disproportionately lower-value. For aging policy, the results suggest that carefully calibrated cost-sharing increases, when paired with effective safety nets, can restrain utilization without detectable adverse survival consequences over the horizon we observe.

2 Institutional Background

2.1 Japan's Healthcare System for the Oldest-Old Population

Japan's universal health insurance system offers an ideal setting for studying cost-sharing responses among the oldest-old. The system covers all residents through mandatory enrollment with standardized fee schedules, eliminating the provider and pricing variation that complicates identification in many other settings (Ikegami et al., 2011). Cost-sharing is structured around transparent age-based coinsurance tiers, with income serving as the sole determinant of differences within each age group—creating the clean variation for causal inference that we exploit below.

The coinsurance rate structure follows a graduated age-based reduction: 30% for individuals under age 70, decreasing to 20% at age 70, and further to 10% at age 75. Since 2006, individuals with income comparable to working-age earners pay 30% regardless of age, preserving income-based equity (Ministry of Health, Labour and Welfare, 2022); prior to the 2022 reform that is the focus of our study, the oldest-old (75+) therefore faced either 10% or 30% depending solely on their income level. In 2008, Japan established a dedicated financing system for this age group—the Late-Stage Elderly Healthcare System (LSEH)—to address their unique healthcare needs and enhance fiscal sustainability through a separate financing pool with its own intergenerational funding structure.[3] This restructuring affected the financing mechanism only; coinsurance rates for the oldest-old remained unchanged. By 2022, LSEH covered approximately 18.52 million individuals—about 98.6% of Japan's oldest-old population—with approximately 93% paying the 10% rate and only the highest-income 7% paying 30% (Ministry of Health, Labour and Welfare, 2022).

LSEH also incorporates an important expenditure safeguard: the High-Cost Medical Expense Benefit—known in Japanese as *Kōryō*—which caps monthly out-of-pocket expenditures at 57,600 JPY regardless of the coinsurance rate.[4] This cap plays a central role in shielding the most medically intensive patients from the financial impact of the 2022 reform—a feature we return to in detail below.

2.2 The 2022 Cost-Sharing Reform

Japan's demographic transformation created acute fiscal pressure on LSEH as the ratio of working-age contributors to elderly beneficiaries declined. In response, an additional tier of income-based cost-sharing was introduced in October 2022, increasing the coinsurance rate from 10% to 20% for middle-income LSEH enrollees while maintaining the existing 10% and 30% rates for lower- and higher-income groups (Ministry of Health, Labour and Welfare, 2022). Of the approximately 93% of enrollees previously paying 10%, the reform moved roughly one in five to the new 20% tier, creating a three-tier structure in which approximately 73% of enrollees continued to pay 10%, approximately 20% moved to the new 20% tier, and 7% maintained the pre-existing 30% rate.

Eligibility for the increased rate required two conditions: classification as an ordinary income earner with individual annual taxable income exceeding 0.28 million JPY, and total household income exceeding thresholds of 2 million JPY for single-person households and 3.2 million JPY for multi-person households.[5] Figure A.1 illustrates the complete eligibility determination process. These thresholds generate sharp discontinuities in cost-sharing: individuals just above faced a doubling of their coinsurance rate while those just below maintained their original rate. Importantly, the income measure encompasses pension income and adjusts other income from the *previous* fiscal year—a backward-looking rule that is central to our identification strategy, since eligibility is determined

[3] LSEH operates under an intergenerational financing model with 10% funded by enrollee premiums, 40% by working-age contributions, and 50% by government subsidies.

[4] The *Kōryō* benefit applies across Japan's healthcare system and is not specific to LSEH. The cap amount varies by income tier; 57,600 JPY corresponds to the ordinary income category relevant to our analytical sample. Within the analytical sample, *Kōryō*-eligible individuals constitute 2% of the population, with no significant difference between treatment and control groups (Table 1).

[5] The policy applied these thresholds separately by household type, as the income pooling available to multi-person households means a given threshold implies different per-capita burdens across household compositions.

by income already earned and reported before the policy was announced, making manipulation of treatment status essentially impossible. Note, however, that coinsurance status is reassessed annually based on the prior year's income, with updates typically implemented in July or August; some individuals may therefore experience a change in status partway through our panel. We verify in Section 5.6 that this does not drive our results.

To mitigate the burden of the rate doubling, the reform incorporated two safeguards. For outpatient care, increases in monthly cost-sharing were capped at 3,000 JPY—a transitional relief measure in effect until September 2025, covering the entire study period.[6] As Appendix D shows, however, this cap meaningfully softened the price increase only for a small share of high-utilization patients: for the median treated outpatient user, the actual-to-nominal price increase ratio was 98.2%, indicating that the cap left the effective price increase nearly intact for most patients. In addition, the *Kōryō* cap continued to bind for the highest-spending enrollees independently of the coinsurance rate, shielding the most medically intensive patients from sharp increases in out-of-pocket costs. As we show in Section 5.4, this combination of safeguards effectively protected high-expenditure patients from the financial impact of the reform—an important feature for interpreting its distributional effects.

Reflecting the combined effect of both safeguards, realized out-of-pocket burden relative to baseline income (Figure D.1) rose sharply by 57–60% upon implementation and stabilized at 44–56% in steady state, confirming that the reform generated a substantial and economically meaningful financial shock for most affected individuals despite the protective measures in place.

Figure 1 Here

The implementation process followed a structured timeline, illustrated in Figure 1. We index all months relative to October 2022 ($k = 0$), the implementation month; notifications of new coinsurance rates were issued in August 2022 ($k = -2$), two months before implementation. This timing creates three well-defined periods: a pre-notification period ($k \leq -3$), a two-month anticipation window ($k \in \{-2, -1\}$) during which affected individuals knew their new rates but had not yet faced them, and a post-implementation period ($k \geq 0$). Combined with the essentially perfect compliance at income thresholds (see Figure 2 below), this temporal structure allows us to separately identify anticipatory and post-implementation behavioral responses within a difference-in-differences event study framework, as we describe in Section 4.

[6] For example, if an individual's monthly outpatient expenses were 50,000 JPY, the coinsurance payment would normally increase from 5,000 JPY (10%) to 10,000 JPY (20%), but the relief measure would limit the actual increase to 3,000 JPY, resulting in an actual payment of 8,000 JPY.

3 Data and Measurements

We employ the Medical Claims Data with Income Tax Information for the Oldest-Old in Japan (MCD-Tx), a novel administrative dataset developed by the Health Insurance Bureau of the Ministry of Health, Labour and Welfare (MHLW). MCD-Tx directly links two comprehensive databases: national health insurance claims and individual income tax records for the population aged 75+. This linkage enables precise measurement of each individual's income-based eligibility for cost-sharing rates—a key advantage over studies relying on self-reported income or administrative proxies (Bound et al., 2001).

The dataset contains monthly observations from November 2021 to December 2023 (26 months) for all LSEH enrollees, spanning the full reform period from well before notification through 15 months post-implementation. Crucially, it includes all enrollees regardless of whether they utilized healthcare services in a given month, avoiding the selection bias common in claims-based studies that observe only active service users. For each individual, the data provide monthly healthcare utilization metrics (total costs and service frequency), income by source (pension, salary, business, and other income), demographic characteristics (residential municipality, sex, and age), and service utilization by disease category.

3.1 Analytical Sample

Our analytical sample covers all LSEH enrollees across both single-person and multi-person households, with income-based eligibility determined following the exact formula specified by MHLW (see Figure A.1).[7] To ensure comparability between treatment and control groups, we restrict the sample to individuals with income within 0.5 million JPY of the applicable threshold (i.e., 1.5–2.5 million JPY for single-person households; 2.7–3.7 million JPY for multi-person households). To pool both household types in a single specification, we normalize income relative to the applicable thresholds, so that the cutoff is at zero for all observations. The resulting sample comprises 1,549,144 individuals (34% single-person, 66% multi-person households), contributing more than 40 million person-month observations.

Table 1 Here

Table 1 (columns 1 and 2) compares the full LSEH population to the analytical sample. The analytical sample contains a higher share of males (57% vs. 40%) and a lower share of single-person households (34% vs. 58%), reflecting the income and household composition of individuals near the middle-income threshold. Baseline healthcare utilization rates, however, are broadly similar across the two groups, suggesting that care consumption patterns near the threshold are representative

[7] For single-person households, the income threshold is 2 million JPY; for multi-person households, 3.2 million JPY. In both cases, the income measure is pension income plus adjusted other income from the previous fiscal year.

of the broader population's service-use behavior. Our estimates should therefore be interpreted as local average treatment effects at the income threshold—relevant to middle-income enrollees near the eligibility cutoff—rather than as population-average effects.

3.2 Treatment and Control

Treatment status is determined by whether an individual's normalized income exceeds the applicable household-specific thresholds as of October 2022. The treatment group comprises individuals whose coinsurance rate increased from 10% to 20%; the control group comprises those who maintained the 10% rate. Within our bandwidth, the treatment group accounts for approximately 80% of the analytical sample (Table 1), reflecting the rightward skew of the income distribution near the threshold (see Figure B.1).

Figure 2 Here

Figure 2 confirms near-perfect compliance at the thresholds: the probability of paying the 20% coinsurance rate jumps from near zero to exactly 1.0 immediately above the respective cutoff, with no anomalous bins or partial compliance. This sharp discontinuity confirms that treatment assignment is exogenously determined by the institutional income formula rather than individual choice.

3.3 Healthcare Utilization Outcomes

We examine multiple outcome measures capturing different dimensions of healthcare utilization response. Our two primary outcomes correspond to distinct margins of consumption: healthcare utilization rates (binary indicators of whether an individual used any healthcare services in a given month) capture the extensive margin, while total monthly healthcare costs, measured in 1,000 JPY and log-transformed among positive-use person-months, capture the intensive margin. We further decompose the intensive margin into total days of healthcare service use and costs per day of service, distinguishing two mechanisms of adjustment: reducing visit frequency versus reducing treatment intensity conditional on a visit. We interpret all log-transformed intensive-margin estimates as descriptive rather than causal, and present a level specification (including zero costs) as a robustness check (see Section 5.6).

Beyond these aggregate measures, we disaggregate utilization and costs into four service categories: inpatient (hospital admissions), outpatient (clinic and hospital visits for non-emergency care), dental, and pharmaceutical services, spanning a gradient from the least discretionary (inpatient) to the most discretionary (dental). This allows us to test whether patients direct reductions toward lower-value, more discretionary care.

3.4 Mortality Outcome

Beyond utilization outcomes, we examine 15-month cumulative mortality (October 2022 through December 2023) as the most consequential health outcome available to us. We proxy mortality using sample attrition, defined as the event that an individual permanently exits the dataset. In the LSEH system, attrition can arise from three sources: death, transition to public assistance (which moves an individual out of LSEH), or relocation across prefectures accompanied by a change in individual identifiers. Our data do not record the reason for exit, so we cannot directly separate mortality from these two non-mortality routes. Two features of our setting nonetheless suggest that attrition predominantly reflects mortality and that the residual non-mortality component is unlikely to differ systematically across the income threshold on which our design relies.

First, transition to public assistance is a means-tested program for low-income households, whereas our analytical sample is restricted to enrollees near the middle-income eligibility threshold, well above the assistance eligibility line on both sides of the cutoff. This route is therefore negligible throughout the bandwidth. Moreover, because eligibility for public assistance is governed by a resource test far removed from our income threshold, there is no reason to expect it to generate a discontinuity in exit at the cutoff. Second, inter-prefectural migration among the elderly is rare, at approximately 0.3% annually for ages 75–79 and lower still at older ages;[8] because this route is small in absolute terms, even a slight imbalance at the income threshold could induce only a correspondingly small effect on our estimates.

Appendix E further validates this proxy along three complementary dimensions. First, attrition counts capture approximately 73–96% of official national death counts across age groups (Table E.1), indicating that attrition captures the large majority of deaths in this near-universally-covered population. Second, healthcare utilization accelerates sharply in the months preceding attrition: hospitalization probability more than triples in the final eleven months, while outpatient visits decline monotonically over the same window, a pattern consistent with terminal transfer from ambulatory to inpatient care rather than administrative exit. Third, disease-specific trajectories for the three leading causes of elderly death in Japan (cancer, circulatory disease, and respiratory disease) all exhibit accelerating inpatient intensity approaching exit, confirming that the aggregate pattern reflects genuine end-of-life disease progression. Together, these checks support sample attrition as a valid proxy for mortality.

[8] Statistics Bureau of Japan (2023) records 18,391 inter-prefectural moves among individuals aged 75–79 in 2022.

4 Identification Strategy

4.1 Main Specifications

Our estimand is the causal effect of the coinsurance increase on healthcare utilization among the oldest-old, traced across event time to separate anticipatory, immediate, and steady-state responses. Identification rests on parallel trends: absent the reform, utilization in the treated and control groups would have evolved in parallel. This assumption allows a fixed level difference between the groups, absorbed by individual fixed effects, and requires only that the groups would not have diverged differentially around implementation. We establish its credibility in Section 4.4; briefly, eligibility is determined by income earned and reported before the reform was announced (ruling out manipulation), and compliance at the threshold is near-perfect (Figure 2), so treated and control groups differ within the bandwidth essentially only in the coinsurance rate they face.

Our primary design is a difference-in-differences (DD) event study that compares enrollees just above the income threshold, whose coinsurance rate rose to 20%, with those just below it, who remained at 10%, before and after implementation, within a 0.5 million JPY bandwidth around the threshold. We assess parallel trends directly through the pre-notification coefficients reported below in Figure 3, and corroborate the event-study estimates with a regression discontinuity design that rests on a different identifying assumption—continuity of potential outcomes at the threshold—rather than parallel trends (Section 4.4). Formally, letting $k = 0$ denote the implementation month (October 2022) and normalizing income relative to the applicable threshold, we estimate:

$$y_{imt} = \alpha_0 + \sum_{k \neq -3} \delta_k \, T_i \times \mathbf{1}\{t = k\} + \theta_i + \phi_t + \eta_{mt} + u_{it} \tag{1}$$

for individual i in municipality m and month t. The dependent variable y_{imt} represents measures of healthcare utilization (binary indicators) or costs (log-transformed values defined only for person-months with positive spending). The treatment indicator T_i equals one if the individual's coinsurance rate is 20% at $k = 0$ and zero if it remains at 10%. The terms $\mathbf{1}\{t = k\}$ are time indicators, with $k = -3$ (July 2022) as the reference period, the last month before individuals were notified of their new coinsurance rates (Figure 1). This choice provides a clean baseline unaffected by anticipatory behavior while preserving the full temporal structure needed to separately identify anticipatory and post-implementation responses. We include individual fixed effects (θ_i) to absorb time-invariant characteristics, time fixed effects (ϕ_t) to absorb secular trends and seasonality, and municipality-by-month fixed effects (η_{mt}) to absorb local factors. Standard errors are clustered at the individual level.

The coefficients of interest, δ_k, capture the differential effect of treatment in each month relative to $k = -3$, tracing the complete behavioral arc from pre-notification ($k = -11$ through $k = -3$) through the anticipation window to post-implementation adjustment: $k = -2$ (August 2022) and $k = -1$ (September 2022) are the two notification months, $k = 0$ marks implementation (October 2022), and positive values denote subsequent post-implementation months.

For binary outcomes, we transform estimated coefficients to percentage changes relative to the pre-treatment mean of the treatment group: percentage change $= (\delta_k/\bar{y}_{\text{pre}}) \times 100$, where $\bar{y}_{\text{pre}}$ is the treatment group mean from $k = -11$ to $k = -3$, so that percentage changes are expressed relative to the treated individuals' own pre-reform baseline. For log-transformed costs, coefficients directly approximate percentage changes. Equation (1) thus estimates a local average treatment effect for middle-income enrollees near the eligibility cutoff.

4.2 Heterogeneity Analysis

We examine behavioral heterogeneity along two dimensions. First, for service-type heterogeneity, we estimate equation (1) separately for inpatient, outpatient, dental, and pharmaceutical services, examining both the extensive and intensive margins. These four categories span a gradient from the least discretionary (inpatient) to the most discretionary (dental), allowing us to test whether patients direct reductions toward lower-value care.

Second, we examine whether price sensitivity varies with patients' underlying medical needs—a dimension conceptually distinct from service-type discretionarity. Whereas discretionarity is an intrinsic property of a service category, medical needs reflect individual patients' health status and care dependency.[9] We classify individuals into four groups based on pre-notification utilization patterns ($k = -11$ to $k = -3$): individuals who were *Kōryō*-eligible at least once during this period form the "*Kōryō*" group,[10] and the remaining patients are divided into three tertiles—low (T1), middle (T2), and high (T3)—based on average outpatient visit days prior to $k = -3$.

The *Kōryō* group warrants particular attention for two reasons. First, because the benefit cap binds independently of the coinsurance rate (see Section 2), this group's effective out-of-pocket costs remain largely unchanged by the reform, providing a built-in falsification test: we should observe little or no behavioral response among them if our estimated effects elsewhere truly reflect price incentives rather than other confounds. Second, the *Kōryō* group represents the most medically intensive patients in our sample, so whether the existing safety net shields them from adverse consequences of the coinsurance increase is of direct relevance to benefit design.

[9] As a concrete illustration, a high-need patient's dental visit can remain discretionary, and a low-need patient's emergency admission can remain non-discretionary.

[10] An individual is classified as *Kōryō*-eligible in a given month if their total healthcare costs exceed 576,000 JPY—the level at which the 10% coinsurance rate would generate out-of-pocket costs equal to the *Kōryō* monthly cap of 57,600 JPY (576,000 $\times$ 0.10 = 57,600). At this cost level or above, the cap binds regardless of whether the coinsurance rate is 10% or 20%, leaving effective out-of-pocket costs unchanged by the reform.

4.3 Mortality Analysis

Beyond these utilization margins, we examine whether cost-sharing affects survival by estimating a cross-sectional regression discontinuity design (RDD) using cumulative mortality over the 15-month post-implementation period ($k = 0$ through $k = 14$) as the outcome. This RDD serves as corroborating evidence by providing a static estimate of the mortality effect that does not rely on the parallel trends assumption. Specifically, we estimate:

$$D_i = \alpha_0 + \tau \cdot T_i + f(X_i) + \varepsilon_i \tag{2}$$

where D_i is an indicator for mortality over the 15-month post-implementation window, proxied by permanent attrition from the panel. $T_i = \mathbf{1}\{X_i \geq 0\}$ is the treatment indicator and X_i denotes normalized income (income relative to the applicable household-specific threshold). The function $f(\cdot)$ is estimated using local linear regression separately on each side of the cutoff, with the optimal bandwidth selected by Calonico et al. (2014). We report the conventional point estimate alongside robust bias-corrected standard errors. The coefficient of interest τ identifies the local average treatment effect on 15-month mortality at the threshold under the assumption that potential outcomes vary smoothly through the cutoff.

Because mortality is proxied by attrition, this continuity assumption additionally requires that any gap between attrition and true mortality not jump at the income threshold. Although we cannot test this directly, we address the concern institutionally instead: as detailed in Section 3.4, both non-mortality exit routes (transition to public assistance and inter-prefectural migration) are rare near the middle-income threshold and neither has a reason to vary systematically at the cutoff. The residual measurement error is therefore unlikely to be differential across the threshold, and we expect it to constitute, at most, a minor source of noise in the estimate. We interpret the result as bearing on a 15-month horizon only, and return to its substantive interpretation and limitations in Section 5.5.

4.4 Identification Validity

The DD event study and the RDD rest on distinct identifying assumptions—parallel trends and continuity of potential outcomes at the threshold, respectively—but share a common threat: manipulation of income around the eligibility cutoff, which could render treated and control groups non-comparable and undermine both designs simultaneously. We therefore begin by ruling out manipulation. Income in the relevant reporting period was reported well before the policy was announced, making strategic adjustment essentially impossible. Two tests confirm this. Figure B.1 shows no bunching below the threshold in the income density, with smooth distributions for the full sample and for single-person and multi-person households separately. Figure B.2 shows that five demographic and socioeconomic covariates (age, gender, single-person household ratio, employment ratio, and pension amount) are smooth through the cutoff without detectable discontinuities, supporting quasi-random assignment within the bandwidth.[11]

With manipulation ruled out, each design is validated on its own terms. For the DD event study, the identifying assumption is parallel trends: absent the reform, utilization would have evolved similarly across the threshold. The pre-notification coefficients are statistically insignificant and tightly clustered around zero across all outcomes, directly supporting this assumption. For the RDD, the identifying assumption is continuity of potential outcomes at the cutoff, which requires no parallel-trends condition; cross-sectional RDD estimates for individual months (Figure B.3) reproduce the direction, broad magnitude, and statistical significance of the event-study findings throughout. Because the two designs rely on different assumptions yet reach the same conclusions, our findings do not depend on either assumption alone.

As a final check on the price transmission mechanism underlying our interpretation, Figure D.1 (Appendix D) confirms that the reform generated a sharp, exogenously timed increase in realized out-of-pocket burden, consistent with the estimated behavioral responses reflecting genuine price incentives rather than other concurrent changes. The absence of a differential pre-trend in this series further corroborates the parallel-trends assumption.

5 Results

5.1 Descriptive Statistics

Table 1 (columns 3–4) presents descriptive statistics for the analytical sample at $k = -3$ (July 2022), the reference month. The population is predominantly male (57%) with a mean age of 81.65 years; 34% live in single-person households. Healthcare utilization is high: 83% of individuals accessed any services, with outpatient care most commonly used (78%), followed by pharmaceutical (64%) and dental services (22%). Average monthly healthcare costs total 60,590 JPY, with inpatient services accounting for the largest share (25,420 JPY) despite being utilized by only 4% of individuals, reflecting the high per-episode cost of hospital admissions. The *Kōryō* share is 2% with no difference between treatment and control groups.

Treatment and control groups differ modestly on observable characteristics by construction: the treatment group is slightly older (81.69 vs. 81.47 years), more likely to live alone (37% vs. 24%),

[11]We note that the single-person household ratio exhibits a pronounced dip in the region just below the threshold before recovering above it. This pattern is unlikely to reflect strategic income manipulation: as shown in Figure B.1, there is no evidence of bunching immediately below the cutoff in the overall income density. Instead, the non-monotone pattern in household composition likely reflects exogenous institutional features of Japan's pension benefit structure and tax deduction system (e.g., the basic income deduction), which tend to concentrate non-single-person households in the income range just above 2 million JPY. Crucially, there is no detectable *discontinuity* in the single-person household ratio at the cutoff itself, satisfying the key requirement for identification.

and has higher income (3.03 vs. 2.76 million JPY). These between-group differences are absorbed by individual fixed effects in our main specification, which identifies treatment effects from within-person variation over time. As discussed in Section 4.1, parallel trends rather than identical levels is the relevant requirement, so a level gap of this magnitude is immaterial to the estimates.

5.2 Effects on Healthcare Utilization and Costs

Figure 3 Here

Figure 3 presents event study estimates of the effect of the increased coinsurance rate on healthcare utilization and costs over the 26-month observation period, with $k = 0$ marking implementation and $k = -3$ the omitted reference period. Across both panels, pre-notification estimates ($k \leq -3$) are statistically insignificant and clustered tightly around zero, consistent with the parallel trends assumption.

Panel (a) reveals three sequential patterns in healthcare utilization rates. First, a statistically significant anticipatory surge at $k = -1$ shows a 2.3% increase among the treatment group. Notably, the coefficient at $k = -2$, the first notification month, is indistinguishable from zero, indicating that the anticipatory response was concentrated entirely in the final month before implementation rather than spread across the two-month notification window. This concentration is consistent with deadline salience rather than gradual forward-looking optimization: patients appear to have responded not upon receiving their notification letters, but only under the immediate pressure of an impending price change (Laibson, 1997; O'Donoghue and Rabin, 1999).[12] Second, at $k = 0$, utilization falls sharply by 3.0%. Third, the post-implementation period ($k = 1$ through $k = 14$) shows persistent negative effects ranging from -0.6% to -1.2%, consistent with a sustained reduction of approximately 1%.

Panel (b) examines log healthcare costs (defined on positive-use person-months) and reveals a similar temporal structure with larger magnitudes. The anticipatory increase at $k = -1$ is 5.9%, while the coefficient at $k = -2$ is again near zero, reinforcing the last-minute character of the surge. At $k = 0$, costs fall immediately by 6.3%, followed by a persistent steady-state decline of approximately 3% throughout the post-implementation period, with point estimates ranging from -2.7% to -4.4%.[13]

[12] The asymmetry between $k = -2$ and $k = -1$ could reflect several mechanisms. Present-biased preferences predict precisely this pattern: individuals systematically underweight future consequences until they become imminent, then respond sharply as the deadline arrives (O'Donoghue and Rabin, 1999). Demand-side inertia may also play a role, as patients may require the salience of an imminent price change, one month away, to overcome scheduling frictions. Provider-side constraints may further compress feasible reorganization into the final month: booking substantially extra appointments two months in advance is difficult in Japan's busy outpatient system, whereas accelerating already-planned visits within a single month is more feasible (Komura and Bessho, 2025). Our data do not allow us to disentangle these demand-side and supply-side explanations, but the concentration of the strategic response at $k = -1$ is itself informative: it implies that the effective window for anticipatory care-seeking is considerably shorter than the formal notification period, with direct implications for how advance notification periods should be designed in future reforms.

[13] Figure B.3 and Table B.1 present cross-sectional RDD estimates for four key months ($k = -3, -1, 0, +5$), corroborating the event-study findings. Each panel shows binned averages (dots) and local linear fits (solid lines) estimated separately on each side of the cutoff. The bandwidth is selected by `rdrobust` with a local linear specification. At $k = -3$ (Panel a), the discontinuity is $+0.4\%$ and statistically insignificant, confirming no pre-existing differences. At $k = -1$ (Panel b), a positive

The near-equal magnitude and opposite sign of the anticipation surge (+5.9% at $k = -1$) and the immediate drop (-6.3% at $k = 0$) indicate that much of the implementation-month decline reflects the reversal of pulled-forward utilization rather than an additional reduction in care: the two effects nearly offset, so the immediate drop should not be read as the magnitude of the persistent price response. That the reduction stabilizes at approximately 3% well into the post-implementation period confirms that the steady-state decline is a genuine new equilibrium rather than lingering intertemporal substitution. Taken together, the 3% steady-state cost reduction against a 1% reduction in utilization rates implies a roughly 3:1 ratio of intensive to extensive margin effects, indicating that the cost response is driven more by reductions in the amount of care consumed than by the decision of whether to seek care at all. These findings yield price elasticity estimates of -0.03 to -0.06,[14] substantially smaller than the -0.1 to -0.3 range documented for younger elderly populations (Shigeoka, 2014; Fukushima et al., 2016), reflecting the more essential nature of healthcare needs among the oldest-old.

Figure 4 Here

The disproportionately large intensive margin response raises an important question about the underlying mechanism. Figure 4 decomposes the intensive margin into days of service use and costs per day. Days of service use and total costs move almost identically throughout the post-implementation period, both declining by approximately 3% in steady state, while costs per day shows only a minimal decline of 0.5–1%. We interpret the cost-per-day series descriptively, as it is defined only among continuing users. Patients thus adjust almost entirely by reducing visit frequency rather than by economizing on treatment intensity within each visit. We examine this visit-frequency mechanism further across service types in the next section.

5.3 Heterogeneity Across Service Types

Figures 5 and 6 Here

We examine heterogeneity across four service categories: inpatient, outpatient, dental, and pharmaceutical. Given inpatient care's distinct nature—lower utilization rates, higher per-episode costs,

discontinuity of $+3.6\%$ corroborates the last-minute surge. At $k = 0$ (Panel c), the discontinuity reverses sharply to -5.0%, and by $k = 5$ (Panel d) it has partially attenuated to -3.0%, consistent with the steady-state decline. All three post-baseline estimates are significant at the 1% level. The full-sample specification in levels, which includes zero-spending months and is therefore free of the conditional-on-positive selection discussed in Section 5.6, reproduces this path (Figure C.6), so the aggregate cost response does not depend on conditioning on use.

[14] Appendix D provides a more detailed analysis accounting for the outpatient relief cap and *Kōryō* benefit; the median treated outpatient user faced an actual-to-nominal price increase ratio of 98.2%, confirming that the price signal reached most patients in full. Elasticities are computed as the steady-state percentage change in spending divided by the percentage change in the patient price. The 10%-to-20% coinsurance increase is a 100% price increase; steady-state cost reductions of roughly 3–6% therefore imply elasticities of -0.03 to -0.06. Because the outpatient relief cap left the effective increase marginally below nominal for some high spending patients, these are conservative in absolute value.

and correspondingly wider confidence intervals—we present its results separately in Figure 5, while Figure 6 displays the remaining three service types together to facilitate comparison.

Inpatient care. Hospital admission rates (Panel a) show a progressive decline over the post-implementation period, deepening from approximately -2.5% in the first months to around -5% to -7% by months 8–10 before partially recovering. Among those admitted, however, inpatient costs (Panel b) show no systematic change throughout, fluctuating around zero. This divergence—declining admission rates alongside stable per-admission costs—is consistent with deferral of elective or semi-elective episodes: if patients are postponing admissions rather than forgoing acute care, the composition and intensity of care within each admitted episode would remain unchanged, as we observe. Diagnosis-level evidence from Appendix F further supports this interpretation: the post-implementation decline is concentrated among elective and deferrable diagnoses such as cataract and arthrosis, averaging approximately 14–15% below pre-treatment levels, while non-deferrable acute-sensitive diagnoses (ischemic heart disease, stroke) show no comparable systematic decline—though the latter estimates carry wide confidence intervals given the rarity of these events. We return to the health consequences of these reductions in Section 5.5, where the mortality evidence provides further context.

Dental care. Dental services exhibit the most pronounced extensive margin response among all service types. Dental visit rates surge by approximately 6–7% at $k = -1$ before falling by a similar magnitude post-implementation, with a persistent steady-state reduction of 3–5%—much larger than the aggregate 1% reduction in overall utilization rates. This pattern reflects the discretionary nature of dental care: unlike outpatient or inpatient services, some dental visits—cleanings, check-ups, elective restorations—can be postponed or brought forward with minimal health risk, making the visit-or-not decision more responsive to price signals. The intensive margin tells a contrasting story: monthly dental costs conditional on utilization show only modest immediate changes and revert toward zero in steady state—markedly weaker than the aggregate 3% intensive margin reduction—implying that when patients do attend, the scope of treatment is largely unchanged. Dental care is thus the most discretionary at the visit margin but not at the within-visit treatment margin.

Outpatient care and pharmaceutical services. Outpatient and pharmaceutical services exhibit an intermediate response pattern. Both show anticipatory behavior at $k = -1$, immediate reductions at $k = 0$, and persistent reductions of approximately 2–3% in both the extensive and intensive margins throughout the post-implementation period. Given that outpatient care is by far the most prevalent service type—utilized by 78% of individuals in the pre-reform period (Table 1)— these reductions in both margins account for the bulk of the aggregate utilization and cost declines documented in the main results. The stability of both margins also implies a genuine new utilization equilibrium rather than a temporary adjustment.

Taken together, these patterns reveal a clear gradient of price responsiveness aligned with service discretionarity: dental care shows the largest visit-frequency reductions but no persistent change in within-visit costs; outpatient and pharmaceutical care show moderate, stable reductions across both margins; and inpatient care shows progressively deepening admission rate declines concentrated in deferrable diagnoses, alongside unchanged per-admission costs. As in the aggregate results, visit frequency is the dominant margin of adjustment across service types, with costs conditional on attendance remaining broadly stable where separable—most clearly for inpatient care and dental care. We next examine whether this pattern varies further with patients' underlying medical needs, and whether the existing safety net shields the most vulnerable.

5.4 Heterogeneity by Underlying Medical Needs

Figure 7 Here

Figure 7 presents results for the four groups defined in Section 4.2, ordered by increasing medical needs. Panel (a) reports extensive margin responses alongside each group's baseline utilization rate; Panel (b) reports intensive margin responses alongside baseline average monthly costs.

Among the non-*Kōryō* groups, price sensitivity follows opposite gradients across the two margins. At the extensive margin, sensitivity varies *inversely* with medical needs: lowest-need users (T1, baseline utilization rate 56%) exhibit the largest anticipatory surge (+6.7% at $k = -1$) and the sharpest immediate reduction (−5.3% at $k = 0$), while highest-need users (T3, baseline 99%) show minimal fluctuation. This reflects a mechanical ceiling on extensive margin adjustment, as patients who already use services in nearly every month have little room to reduce visit frequency further regardless of price incentives (Chandra et al., 2010). At the intensive margin, the gradient reverses: higher medical needs are associated with larger cost reductions in both the immediate and steady-state periods, with T3 showing the largest swings.[15] Having exhausted the extensive margin as an adjustment channel, high-need patients express their price response through whatever intensive margin flexibility remains, while low-need patients who can freely reduce visit frequency have less need to economize within each visit.

The *Kōryō* group stands in sharp contrast to all three non-*Kōryō* tertiles. Despite being the highest-need group in the sample (average monthly costs of JPY 317,070, more than five times those of T3), they are the *least* price-sensitive across both margins: effects on visit rates and costs are small and

[15] The outpatient transitional relief cap (Section 2; Appendix D) attenuates the effective price increase in the highest-spending outpatient months, which are concentrated among higher-need patients. Any such attenuation works against finding a response among T3, so the larger intensive-margin reductions we observe for this group are, if anything, a lower bound on their true price sensitivity. This stands in contrast to the *Kōryō* group, whose null response reflects the expenditure cap binding fully and independently of the coinsurance rate.

statistically insignificant throughout the post-implementation period. This null result is precisely what the institutional design predicts: patients whose monthly expenses regularly exceed the *Kōryō* cap face only negligible changes in effective out-of-pocket costs when the coinsurance rate doubles, because the cap binds independently of the rate. The *Kōryō* benefit thus shields the most medically vulnerable patients from the financial consequences of the reform at precisely the point where unmet need would be most consequential. The absence of behavioral response among this group also confirms the falsification test described in Section 4.2: our estimated effects elsewhere reflect genuine price incentives rather than other confounds.

The distributional implications are therefore more benign than a uniform cost-sharing increase might suggest. Utilization reductions are concentrated among patients with sufficient room to adjust their care-seeking, while the existing safety net shields those for whom foregone care would be most consequential.

5.5 Effects on 15-Month Cumulative Mortality

Figure 8 Here

The utilization reductions documented above, particularly the progressive decline in hospital admissions, raise a central policy concern: does the foregone care adversely affect health? We address this question directly by examining 15-month cumulative mortality.

Figure 8 presents RDD estimates of the effect of the increased coinsurance rate on 15-month cumulative mortality ($k = 0$ through $k = 14$). Panel (a) plots the raw relationship between normalized income and mortality across the full income distribution; a negative income–mortality gradient is apparent overall, with no obvious discontinuity near the threshold.[16] Panel (b) presents the formal RDD estimate. Despite the reductions in hospital admissions and outpatient visits documented above, we find no statistically detectable evidence that the increased coinsurance rate adversely affected survival: the point estimate indicates that 15-month mortality is slightly lower among individuals just above the threshold by 0.2 percentage points (a 2.5% relative reduction against a baseline mortality rate of approximately 8%), but it does not reach conventional levels of statistical significance.[17] We therefore read this result as the absence of evidence of harm over a 15-month horizon rather than as positive evidence of safety.[18] The estimate is nonetheless informative about the effects the data can

[16] The greater dispersion in mortality among lower-income individuals likely reflects the compounding of multiple disadvantages at the lower end of the income distribution: greater heterogeneity in health status, living conditions, social isolation, and access to informal care networks (Liang et al., 2002). These individuals fall outside our analytical bandwidth and are therefore unaffected by the reform examined here; this systematic heterogeneity nonetheless represents an important avenue for future research.

[17] See Table C.1 for the point estimate with statistical inference. Because mortality is proxied by panel attrition, the estimate may carry measurement error from non-mortality exits; as discussed in Sections 3.4 and 4.3, these exits are rare and unlikely to differ systematically at the threshold, which bounds the resulting bias.

[18] Our data cannot confirm that the deferred care carried no health cost beyond the 15-month window. Particularly, the progressive deepening of hospital admission reductions is worth monitoring with extended follow-up, as adverse effects beyond our observation window cannot be ruled out.

exclude: the 95% confidence interval spans roughly -0.6 to $+0.2$ percentage points, ruling out any increase in 15-month mortality beyond about 0.2 percentage points.

The absence of a detectable mortality effect, read alongside the utilization evidence, points to a consistent interpretation: the care foregone in response to higher cost-sharing was likely lower-value and deferrable. Three strands of evidence support this reading. First, *Kōryō*-eligible patients, those with the most intensive and presumably most essential healthcare needs, experienced no significant utilization reductions, implying that the safety net ring-fenced care for the highest-need patients at precisely the point where foregone care would have been most consequential. Second, the largest utilization reductions occurred in dental care, highly elective at the visit margin, with the progressive deepening and partial recovery of inpatient admission rates consistent with deferral rather than permanent foregone care. Third, diagnosis-level evidence from Appendix F shows that the admission decline is concentrated in elective and deferrable procedures rather than acute conditions, suggesting that the marginal admissions foregone were clinically deferrable. Taken together, these patterns suggest that the oldest-old engage in self-triage that concentrates reductions on more deferrable care, producing an outcome consistent with what value-based insurance design aims to achieve: reductions concentrated in lower-value utilization without detectable compromise to survival (Chernew et al., 2007; Fendrick et al., 2012).[19] Though our data cannot confirm that the deferred care carried no health cost beyond the 15-month window, the convergence of evidence across service types, diagnosis groups, and medical need strata is consistent with this interpretation.

5.6 Robustness Checks

We conduct five additional checks addressing potential threats to our main estimates, all of which leave our main findings unchanged. First, we re-estimate with an alternative reference period, using $k = -4$ as the omitted baseline instead of $k = -3$; estimates are nearly identical throughout (Figure C.1).

Second, because coinsurance rates are reassessed annually based on updated income, some individuals—particularly those whose pension income declined—may have reverted from the 20% to the 10% rate in the final months of our panel; the reverse is unlikely in this elderly, largely pension-dependent population, since pension income rarely rises. If anything, this would attenuate our late-period estimates toward zero, since some individuals coded as treated would in fact face the lower rate. To verify this does not drive our results, we exclude individuals whose coinsurance status changed following the annual reassessment, which affected approximately 16% of individuals observed in the month of the update; excluding them leaves the main results essentially unchanged

[19] We emphasize that the reform itself was not a value-based design: it raised the coinsurance rate uniformly for all eligible enrollees rather than tiering it by the clinical value of each service. The alignment with value-based design principles documented here thus reflects patient self-triage and the *Kōryō* expenditure cap—responses to a blunt instrument—rather than deliberate policy calibration.

(Figure C.2).

Third, and related, we guard against post-implementation income manipulation, whereby individuals just above the threshold might selectively reduce their income in subsequent years to escape the higher rate, by examining the income distribution and income changes roughly a year after implementation. We find no bunching below the threshold and no discontinuity in income changes at the cutoff, providing no evidence of strategic manipulation (Figures C.3 and C.4).

Fourth, we conduct placebo threshold tests at income levels 0.5 million JPY below and above the actual cutoff, where no reform-induced rate change occurred; neither placebo produces the anticipatory surge, implementation drop, or persistent decline observed at the true threshold (Figure C.5).

Fifth, we address the interpretation of intensive-margin cost outcomes. Log total monthly healthcare costs is defined only for positive-use person-months, so the regression sample is itself affected by treatment; these coefficients cannot be read as clean causal effects on spending per user (Buntin and Zaslavsky, 2004; Chen and Roth, 2024). A level specification inclusive of zero-spending months reproduces the same anticipation, implementation, and steady-state pattern (Figure C.6), bounding this concern. We therefore retain log spending as our primary intensive-margin outcome for comparability and precision, and treat conditional-on-use quantities (cost per day, days of service, and per-service costs conditional on a visit or admission) as descriptive of the composition of care among continuing users rather than as causal per-visit effects. Our conclusion that adjustment operates through visit frequency rather than intensity per visit rests on the full-sample extensive-margin response (Figure 3), which is free of any conditioning; the decomposition in Figure 4 provides consistent descriptive corroboration among continuing users.

6 Discussion and Conclusion

Using data on essentially the entire population aged 75+ in Japan, this study traces the behavioral arc of the 2022 reform that doubled coinsurance from 10% to 20%, and examines its mortality consequences. Our findings yield three principal contributions: documenting a last-minute anticipatory surge driven by deadline salience rather than gradual forward-looking optimization; identifying visit frequency rather than treatment intensity as the dominant margin of adjustment, with heterogeneous adaptation paths across service types and medical need groups; and showing that carefully calibrated cost-sharing, paired with effective safety nets, restrained utilization without detectable adverse mortality consequences over a 15-month horizon among the oldest-old.

We first document a last-minute anticipatory surge concentrated in the final month before implementation, with no detectable response in the earlier notification month, a pattern that holds uniformly across service types and medical need groups. This concentration is consistent with a deadline effect (Laibson, 1997; O'Donoghue and Rabin, 1999): patients respond to the immediate salience of an impending price change rather than engaging in gradual intertemporal optimization upon receiving written notification two months earlier. This contrasts with the classic anticipatory response literature, which has typically found behavioral responses spread across notification windows (Einav et al., 2015; Card et al., 2008). The implication is that the effective anticipation window for this population is considerably shorter than the formal notification period, reflecting both demand-side inertia and supply-side scheduling constraints. That even the oldest-old engage in strategic healthcare timing challenges any assumption that essential-need populations are behaviorally unresponsive to cost signals.

Post-implementation adaptation unfolds in three sequential phases: a sharp immediate reduction, a period of partial attenuation, and convergence to a persistent new equilibrium stable throughout our 15-month post-implementation window. The immediate response is more than twice the magnitude of the steady-state reduction, a gap with an important methodological implication for the broader literature. Prior studies have documented price elasticities for elderly populations substantially larger than our estimates (Shigeoka, 2014; Fukushima et al., 2016; Ando and Takaku, 2016; Nishi et al., 2012). While part of this difference reflects the more essential nature of healthcare needs among the oldest-old (Clegg et al., 2013; Dent et al., 2019), our results suggest an additional explanation: studies with shorter observation windows may capture the immediate response rather than the steady-state equilibrium, conflating the two and potentially overstating long-run price sensitivity. Our post-implementation window is, to our knowledge, among the longest in this literature, and the stability of steady-state estimates supports their interpretation as mid- to long-run elasticities.

Understanding how the steady-state equilibrium is achieved requires examining the underlying mechanism. The evidence points to visit frequency rather than treatment intensity as the operative margin. The clearest support is the extensive-margin response itself: a full-sample, causal reduction in whether any care is sought at all, which is by definition a frequency response. Consistent with this, a decomposition among continuing users shows total spending and days of service declining almost identically while cost per day is essentially flat; we read this as descriptive corroboration, since its components are defined only on positive-use person-months. The relevant margin for cost-sharing policy among the oldest-old thus appears to be the visit-or-not decision rather than the treatment-intensity decision.

This visit-frequency mechanism manifests differently across service types according to how much discretion patients have over the visit decision. The gradient runs from dental care, where visit rates move sharply but the care delivered per visit is little changed, through outpatient and pharmaceutical services at an intermediate position, to hospital admissions, where a progressive decline concentrated in elective and deferrable diagnoses (Appendix F), alongside unchanged cost per admission, indicates deferral of elective procedures rather than foregone acute care.

Heterogeneity by medical needs reveals a complementary dimension. Among non-*Kōryō* patients, price sensitivity follows opposite gradients across margins: lowest-need patients reduce visit frequency most sharply, while highest-need patients, facing a ceiling on extensive margin adjustment, express their response through the intensive margin instead. The *Kōryō*-eligible group, the most medically intensive in our sample, shows no sustained behavioral response across either margin, precisely because the expenditure cap binds independently of the coinsurance rate and leaves their effective out-of-pocket costs largely unchanged. The two gradients together imply that reductions are systematically concentrated among patients with lower medical needs consuming more discretionary services, a pattern that, read alongside the null mortality finding, is consistent with the interpretation that foregone care was disproportionately lower-value.

Finally, we find no detectable adverse effect on 15-month cumulative mortality. Despite reductions in outpatient visits and hospital admissions, the point estimate is slightly lower among those assigned the higher coinsurance rate. Although statistically insignificant, it is precise enough to exclude any increase in mortality beyond a small margin (roughly 0.2 percentage points) over this horizon. We read this as an absence of evidence of harm rather than as positive evidence of safety. The reading is reinforced by the service-type, diagnosis-level, and medical-need gradients: reductions concentrated in dental and elective or deferrable inpatient care, with the *Kōryō* cap ring-fencing care for the highest-need patients. Together these patterns are consistent with the value-based insurance design hypothesis (Chernew et al., 2007; Fendrick et al., 2012): although the reform applied a uniform rather than value-tiered increase, the oldest-old appear to engage in meaningful self-triage, concentrating reductions on discretionary and deferrable care while protecting less deferrable care.

Our study has several important limitations. First, the mortality analysis carries two related caveats: the 15-month window may not capture longer-term consequences of sustained care reductions, with the progressive deepening of hospital admission reductions particularly worth monitoring; and mortality is measured by panel attrition that our data cannot decompose into mortality and non-mortality components, so the estimate carries measurement error we can bound but not eliminate. Second, while mortality is the most consequential health outcome available to us, we cannot rule out adverse effects on health status, functional capacity, or quality of life that do not manifest as mortality within this period. Third, our analysis focuses on a relatively affluent subset near the income threshold, and responses may differ among lower-income populations (Finkelstein et al., 2019). Fourth, we cannot fully distinguish demand-side from supply-side forces driving the visit-frequency mechanism, an important avenue for future research.

Our findings nonetheless speak to a broader tension between fiscal sustainability and health protection (Rice and Matsuoka, 2004). The conventional worry is that coinsurance rate increases reduce essential care and harm vulnerable populations. Our evidence suggests a more nuanced picture: when safety nets shield the highest-need patients and the broader population retains sufficient discretionary care to absorb visit-frequency reductions, cost-sharing can achieve fiscal objectives without detectable mortality costs over our observation window.

Several implications follow for the design of cost-sharing among the oldest-old. The *Kōryō* experience suggests that binding expenditure caps are a critical complement to coinsurance rate increases: by shielding the highest-need patients from the financial consequences of reform at precisely the point where unmet need would be most consequential, such safety nets help cost-sharing achieve fiscal objectives while limiting the risk to care most essential to health. The dominance of the visit-frequency margin also has practical relevance for instrument choice: flat per-visit fees may achieve similar utilization goals more directly than proportional coinsurance rates, which in principle target treatment intensity but appear largely ineffective at that margin among the oldest-old. Finally, the concentration of the anticipatory response in a single month rather than across the two-month notification window suggests that longer advance notice alone may not spread behavioral adjustment smoothly; pairing notification with targeted provider outreach may be necessary to prevent last-minute scheduling surges.

References

Abaluck, J., Gruber, J., and Swanson, A. (2018). Prescription drug use under Medicare Part D: A linear model of nonlinear budget sets. *Journal of Public Economics*, 164:106–138.

Alpert, A. (2016). The anticipatory effects of Medicare Part D on drug utilization. *Journal of Health Economics*, 49:28–45.

Ando, M. and Takaku, R. (2016). Affordable false teeth: The effects of patient cost sharing on denture utilization and subjective chewing ability. *BE Journal of Economic Analysis & Policy*, 16(3):1387–1438.

Aron-Dine, A., Einav, L., and Finkelstein, A. (2013). The RAND health insurance experiment, three decades later. *Journal of Economic Perspectives*, 27(1):197–222.

Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review*, 53:941–973.

Baicker, K. and Levy, H. (2015). Cost sharing as a tool to drive higher-value care. *JAMA Internal Medicine*, 175(3):399–400.

Best, M. C. and Kleven, H. J. (2018). Housing market responses to transaction taxes: Evidence from notches and stimulus in the U.K. *The Review of Economic Studies*, 85(1):157–193.

Bound, J., Brown, C., and Mathiowetz, N. (2001). Measurement error in survey data. In *Handbook of Econometrics*, volume 5, pages 3705–3843. Elsevier, Amsterdam.

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.

Buntin, M. B. and Zaslavsky, A. M. (2004). Too much ado about two-part models and transformation? comparing methods of modeling health care costs. *Journal of Health Economics*, 23(3):525–542.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.

Card, D., Dobkin, C., and Maestas, N. (2008). The impact of nearly universal insurance coverage on health care utilization: Evidence from Medicare. *American Economic Review*, 98(5):2242–2258.

Chandra, A., Gruber, J., and McKnight, R. (2010). Patient cost-sharing and hospitalization offsets in the elderly. *American Economic Review*, 100(1):193–213.

Chen, J. and Roth, J. (2024). Logs with zeros? Some problems and solutions. *Quarterly Journal of Economics*, 139(2):891–936.

Chernew, M. E., Rosen, A. M., and Fendrick, A. M. (2007). Value-based insurance design. *Health Affairs*, 26(Supplement 2):w195–w203.

Clegg, A., Young, J., Iliffe, S., Rikkert, M. O., and Rockwood, K. (2013). Frailty in elderly people. *Lancet*, 381(9868):752–762.

Dent, E., Martin, F. C., Bergman, H., Woo, J., Romero-Ortuno, R., and Walston, J. D. (2019). Management of frailty: Opportunities, challenges, and future directions. *Lancet*, 394(10206):1376–1386.

Duarte, F. (2012). Price elasticity of expenditure across health care services. *Journal of Health Economics*, 31(6):824–841.

Einav, L., Finkelstein, A., and Polyakova, M. (2018). Private provision of social insurance: Drug-specific price elasticities and cost sharing in Medicare Part D. *American Economic Journal: Economic Policy*, 10(3):122–153.

Einav, L., Finkelstein, A., and Schrimpf, P. (2015). The response of drug expenditure to nonlinear contract design: Evidence from Medicare Part D. *Quarterly Journal of Economics*, 130(2):841–899.

Ellis, R. P. and McGuire, T. G. (2007). Predictability and predictiveness in health care spending. *Journal of Health Economics*, 26(1):25–48.

Eurostat (2024). Population structure and ageing. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Population_structure_and_ageing.

Fendrick, A. M., Martin, J. J., and Weiss, A. E. (2012). Value-based insurance design: More health at any price. *Health Services Research*, 47(1):404–413.

Finkelstein, A., Hendren, N., and Shepard, M. (2019). Subsidizing health insurance for low-income adults: Evidence from Massachusetts. *American Economic Review*, 109(4):1530–1567.

Fried, L. P., Tangen, C. M., Walston, J., Newman, A. B., Hirsch, C., Gottdiener, J., Seeman, T., Tracy, R., Kop, W. J., Burke, G., and McBurnie, M. A. (2001). Frailty in older adults: Evidence for a phenotype. *Journal of Gerontology Series A*, 56(3):M146–M156.

Fukushima, K., Mizuoka, S., Yamamoto, S., and Iizuka, T. (2016). Patient cost sharing and medical expenditures for the elderly. *Journal of Health Economics*, 45:115–130.

Goolsbee, A. (2000). What happens when you tax the rich? evidence from executive compensation. *Journal of Political Economy*, 108(2):352–378.

Grossman, M. (1972). On the concept of health capital and the demand for health. *Journal of Political Economy*, 80(2):223–255.

Gruber, J., Maclean, J. C., Wright, B., Wilkinson, E., and Volpp, K. G. (2020). The effect of increased cost-sharing on low-value service use. *Health Economics*, 29(10):1180–1201.

Gruneir, A., Silver, M. J., and Rochon, P. A. (2011). Emergency department use by older adults: A literature review on trends, appropriateness, and consequences of unmet health care needs. *Medical Care Research and Review*, 68(2):131–155.

Hendel, I. and Nevo, A. (2004). Intertemporal substitution and storable products. *Journal of the European Economic Association*, 2(2–3):536–547.

Hendel, I. and Nevo, A. (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673.

Iizuka, T. and Shigeoka, H. (2022). Is zero a special price? evidence from child health care. *American Economic Journal: Applied Economics*, 14(4):381–410.

Iizuka, T. and Shigeoka, H. (2023). Asymmetric demand response when prices increase and decrease: The case of child health care. *Review of Economics and Statistics*, 105(5):1325–1333.

Ikegami, N., Yoo, B.-K., Hashimoto, H., Matsumoto, M., Ogata, H., Babazono, A., Watanabe, R., Shioya, I., Ikezaki, M., Comura, H., Reich, M., and Kobayashi, Y. (2011). Japanese universal health coverage: Evolution, achievements, and challenges. *Lancet*, 378(9796):1106–1115.

Komura, N. and Bessho, S. (2025). Dynamics of consumer responses to medical price changes. *American Economic Review: Insights*, 7(4):534–549.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112(2):443–478.

Liang, J., Bennett, J., Krause, N., Kobayashi, E., Kim, H., Brown, J. W., Akiyama, H., Sugisawa, H., and Jain, A. (2002). Old age mortality in Japan: Does the socioeconomic gradient interact with gender and age? *Journals of Gerontology: Series B*, 57(5):S294–S307.

Manning, W. G., Newhouse, J. P., Duan, N., Keeler, E. B., and Leibowitz, A. (1987). Health insurance and the demand for medical care: Evidence from a randomized experiment. *American Economic Review*, 77(3):251–277.

Mian, A. and Sufi, A. (2012). The effects of fiscal stimulus: Evidence from the 2009 Cash for Clunkers program. *Quarterly Journal of Economics*, 127(3):1107–1142.

Ministry of Health, Labour and Welfare (2022). *Annual Report on the Medical Care System for the Elderly*. MHLW, Tokyo.

Ministry of Health, Labour and Welfare (2025). Estimates of national medical care expenditure, FY2023 (Reiwa 5-nendo Kokumin Irȳohi no Gaikyō). Technical report, Ministry of Health, Labour and Welfare.

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

Nishi, A., McWilliams, J. M., Noguchi, H., Hashimoto, H., Tamiya, N., and Kawachi, I. (2012). Health benefits of reduced patient cost sharing in Japan. *Bulletin of the World Health Organization*, 90:426–435a.

O'Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review*, 89(1):103–124.

OECD (2023). OECD economic surveys korea. https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/07/oecd-economic-surveys-korea-2024_9343c046/c243e16a-en.pdf.

Poterba, J. M. (1988). Are consumers forward looking? evidence from fiscal experiments. *American Economic Review*, 78(2):413–418.

Rice, T. and Matsuoka, K. Y. (2004). The impact of cost-sharing on appropriate utilization and health status: A review of the literature on seniors. *Medical Care Research and Review*, 61(4):415–452.

Salive, M. E. (2013). Multimorbidity in older adults. *Epidemiologic Reviews*, 35(1):75–83.

Shigeoka, H. (2014). The effect of patient cost sharing on utilization, health, and risk protection. *American Economic Review*, 104(7):2152–2184.

Statistics Bureau, Ministry of Internal Affairs and Communications (2025). Current population estimates as of October 1, 2024. Technical report, Statistics Bureau, Ministry of Internal Affairs and Communications.

Statistics Bureau of Japan (2023). Report on internal migration in japan, 2022. Technical report, Statistics Bureau, Ministry of Internal Affairs and Communications.

Table 1: Descriptive Statistics One Month Prior to Policy Notification ($k = -3$, July 2022)

	(1) Full Sample	(2) Analytical Sample	(3) Control	(4) Treatment	(5) Difference
A. Demographic characteristics					
Male	0.40	0.57	0.51	0.58	0.07***
	(0.49)	(0.50)	(0.50)	(0.49)	
Age	82.41	81.65	81.47	81.69	0.22***
	(5.91)	(4.91)	(4.87)	(4.92)	
Single-person households	0.58	0.34	0.24	0.37	0.13***
	(0.49)	(0.47)	(0.43)	(0.48)	
Household income (million JPY)	2.47	2.97	2.76	3.03	0.27***
	(4.85)	(0.59)	(0.52)	(0.60)	
B. Healthcare utilization rates					
Healthcare utilization rate	0.84	0.83	0.82	0.84	0.01***
	(0.36)	(0.37)	(0.38)	(0.37)	
Inpatient	0.05	0.04	0.04	0.04	0.00**
	(0.23)	(0.19)	(0.19)	(0.19)	
Outpatient	0.78	0.78	0.77	0.78	0.01***
	(0.41)	(0.41)	(0.42)	(0.41)	
Dental	0.21	0.22	0.20	0.23	0.03***
	(0.41)	(0.42)	(0.40)	(0.42)	
Pharmaceutical	0.65	0.64	0.63	0.65	0.01***
	(0.48)	(0.48)	(0.48)	(0.48)	
C. Healthcare costs (1,000 JPY)					
Healthcare costs	74.45	60.59	58.42	61.12	2.70***
	(216.62)	(191.60)	(187.05)	(192.69)	
Inpatient	36.13	25.42	24.81	25.57	0.76**
	(202.80)	(177.94)	(174.58)	(178.75)	
Outpatient	22.50	20.20	19.30	20.41	1.11***
	(63.06)	(56.56)	(52.78)	(57.44)	
Dental	3.10	3.30	3.04	3.36	0.32***
	(12.10)	(12.12)	(11.16)	(12.35)	
Pharmaceutical	11.90	11.18	10.84	11.27	0.43***
	(34.47)	(32.81)	(31.36)	(33.15)	
Kōryō share	0.03	0.02	0.02	0.02	0.00
	(0.17)	(0.14)	(0.14)	(0.14)	
Number of individuals	18,653,796	1,549,144	304,210	1,244,934	1,549,144

Notes: This table presents summary statistics at $k = -3$ (July 2022), one month before policy notification. Column (1) reports the full LSEH population with ordinary income certification. Columns (2)–(5) report the analytical sample, restricted to individuals within the 0.5 million JPY bandwidth around the applicable household-income threshold (2 million JPY for single-person households; 3.2 million JPY for multi-person households), with income normalized relative to the applicable threshold. Column (2) reports the analytical sample overall; column (3) reports the control group (normalized income below zero; 10% coinsurance rate); column (4) reports the treatment group (normalized income above zero; 20% coinsurance rate); column (5) reports the difference (treatment minus control) with significance tests. The analytical sample skews toward male and multi-person households relative to the full population, while baseline healthcare utilization rates are broadly comparable across samples, supporting the generalizability of the estimated effects within the analytical sample (see Section 5.2 for further discussion). Healthcare expenditures are in thousands of JPY. Standard deviations in parentheses. Statistical significance of differences: *** $p < 0.01$, ** $p < 0.05$.

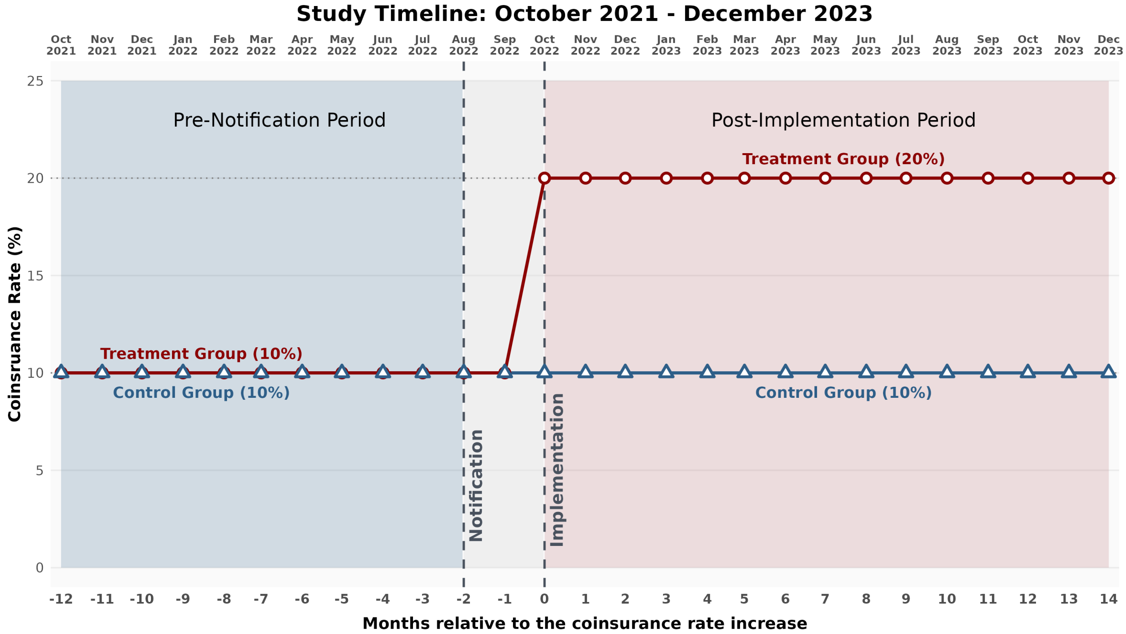

Figure 1: Study Timeline: Coinsurance Rates for Treatment and Control Groups

Notes: This figure illustrates the study timeline and coinsurance rate changes. The treatment group experienced an increase from 10% to 20% after $k = 0$ (October 2022, implementation), while the control group maintained 10% throughout. The timeline divides into three phases: Pre-Notification ($k \leq -3$, October 2021–July 2022), Notification and Anticipation ($k \in \{-2, -1\}$, August–September 2022), and Post-Implementation ($k \geq 0$, October 2022 onward). $k = -3$ denotes July 2022, the omitted reference period; $k = -2$ (August 2022) and $k = -1$ (September 2022) are the two notification months; $k = 0$ (October 2022) marks implementation.

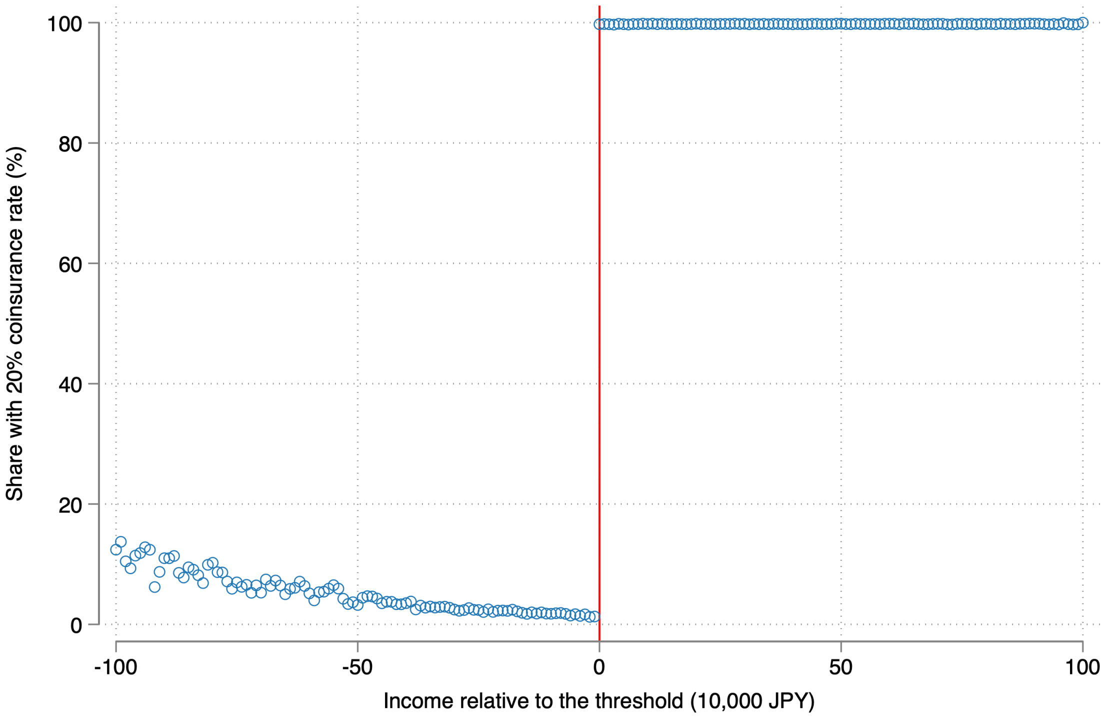

Figure 2: Treatment Assignment: Coinsurance Rate by Normalized Income

Notes: This figure illustrates the sharp discontinuity in coinsurance rate assignment at the income threshold as of $k = 0$ (October 2022). The x-axis shows annual income relative to the applicable threshold (2 million JPY for single-person; 3.2 million JPY for multi-person households), in million JPY. The y-axis represents the proportion assigned the 20% coinsurance rate. The jump from near-zero to essentially 1.0 at the threshold confirms near-perfect compliance and reflects the correct implementation of the income formula for both household types.

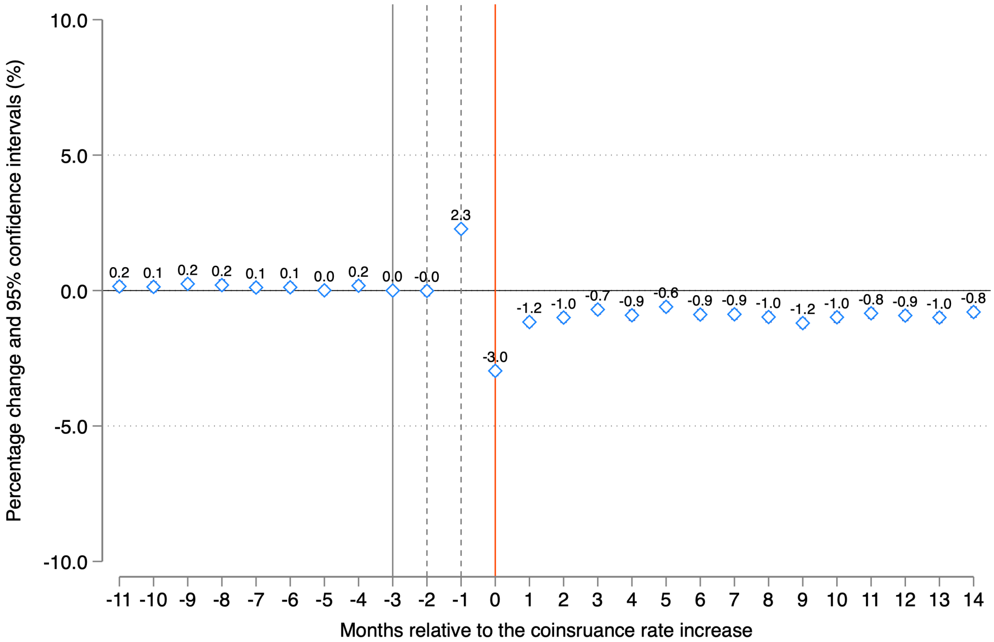

(a) Healthcare utilization rate (extensive margin)

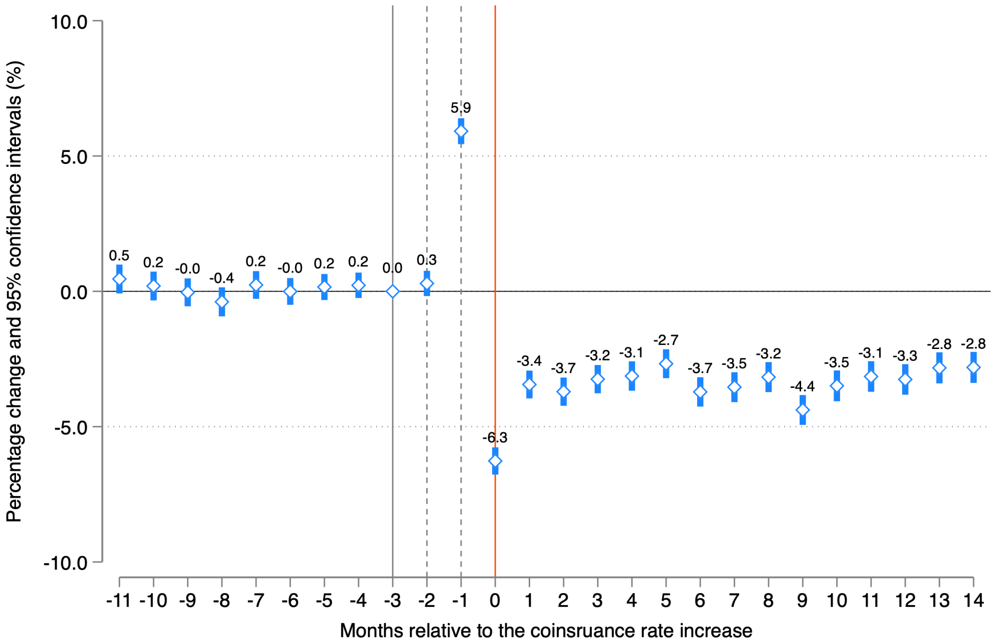

(b) Healthcare costs, log (intensive margin)

Figure 3: Effects on Healthcare Utilization and Costs

Notes: Event study estimates of the effect of increased coinsurance rates (from 10% to 20%) on overall healthcare consumption among individuals aged 75+ with normalized income within 0.5 million JPY of the applicable threshold. Panel (a) shows utilization rates (extensive margin); Panel (b) shows log expenditures (intensive margin). Diamonds are point estimates; vertical bars show 95% confidence intervals. The solid black vertical line marks $k = -3$ (July 2022, the omitted reference period). Two dashed vertical lines indicate $k = -2$ (August 2022) and $k = -1$ (September 2022), spanning the anticipation period. The solid red vertical line marks $k = 0$ (October 2022, implementation). Estimates are expressed as percentage changes relative to the treatment group's pre-treatment mean for binary outcomes and as approximate percentage changes for log-transformed expenditure outcomes. All models include individual, month, and municipality-by-month fixed effects. Standard errors are clustered at the individual level.

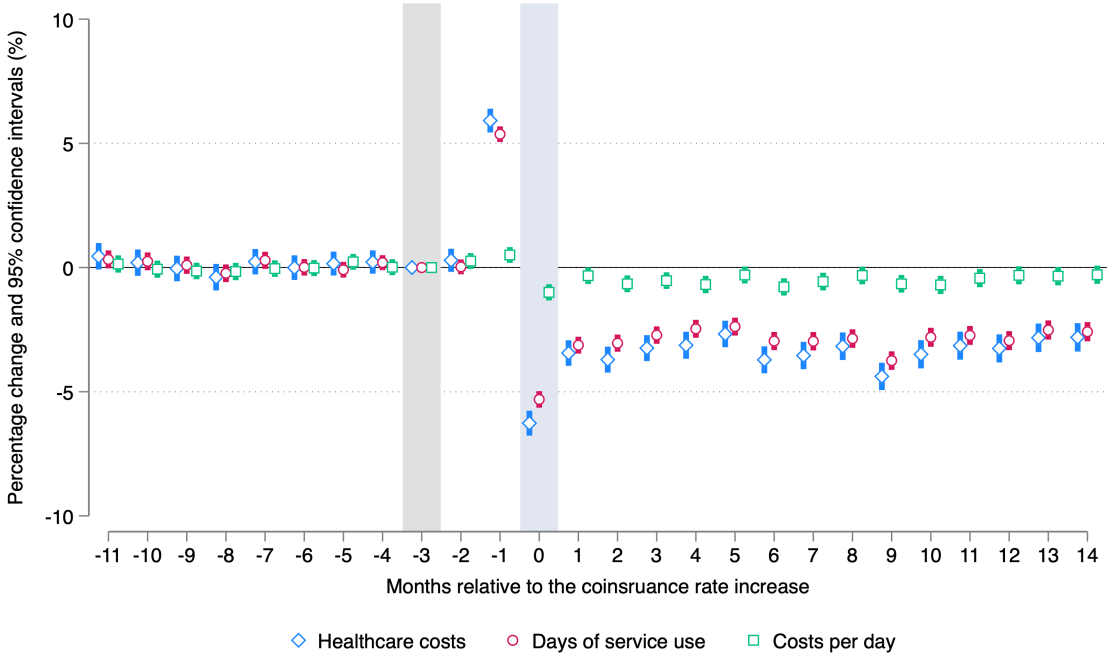

Figure 4: Decomposition of the Intensive Margin: Total Costs, Days of Service, and Costs per Day

Notes: This figure decomposes the intensive margin into three components: total medical expenditure (blue diamonds), days of healthcare service use (red circles), and expenditure per day (green squares). All series are expressed as percentage changes relative to pre-treatment means. The gray shaded region marks $k = -3$ (July 2022, the omitted reference period); the light blue shaded region marks $k = 0$ (October 2022, implementation); the unshaded region between them ($k \in \{-2, -1\}$) constitutes the anticipation period. In the anticipation period, total expenditure and days of service both surge at $k = -1$, while expenditure per day shows little change, indicating that the anticipatory response operates through visit frequency rather than treatment intensity. Post-implementation, expenditure and days of service track each other closely, both declining by approximately 3% in steady state, while expenditure per day shows only a 0.5–1% decline—confirming that the dominant adjustment mechanism throughout is visit frequency reduction rather than within-visit treatment intensity reduction. All models include individual, month, and municipality-by-month fixed effects. Standard errors are clustered at the individual level.

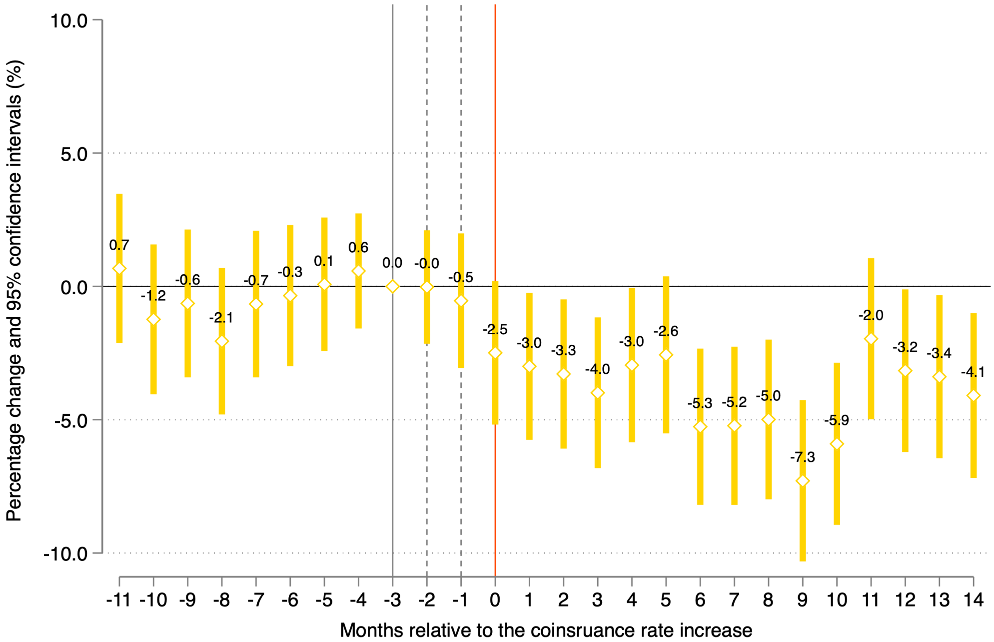

(a) Hospital admission rate (extensive margin)

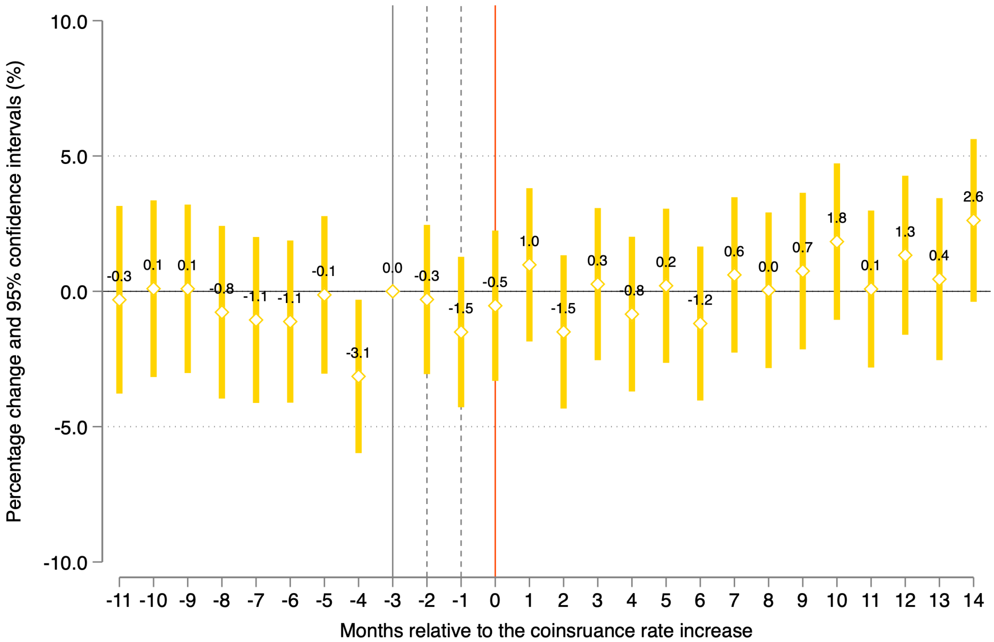

(b) Inpatient costs, log (intensive margin)

Figure 5: Heterogeneous Responses for Inpatient Care

Notes: Event study estimates for inpatient care. Panel (a) shows hospital admission rates (extensive margin); Panel (b) shows log inpatient expenditures conditional on admission (intensive margin). Diamonds are point estimates; vertical bars show 95% confidence intervals. The solid black vertical line marks $k = -3$ (July 2022, the omitted reference period). Two dashed vertical lines indicate $k = -2$ (August 2022) and $k = -1$ (September 2022), spanning the anticipation period. The solid red vertical line marks $k = 0$ (October 2022, implementation). The divergence—declining admission rates alongside stable per-admission costs—is consistent with deferral of elective or semi-elective admissions rather than foregone acute care. All models include individual, month, and municipality-by-month fixed effects. Standard errors are clustered at the individual level.

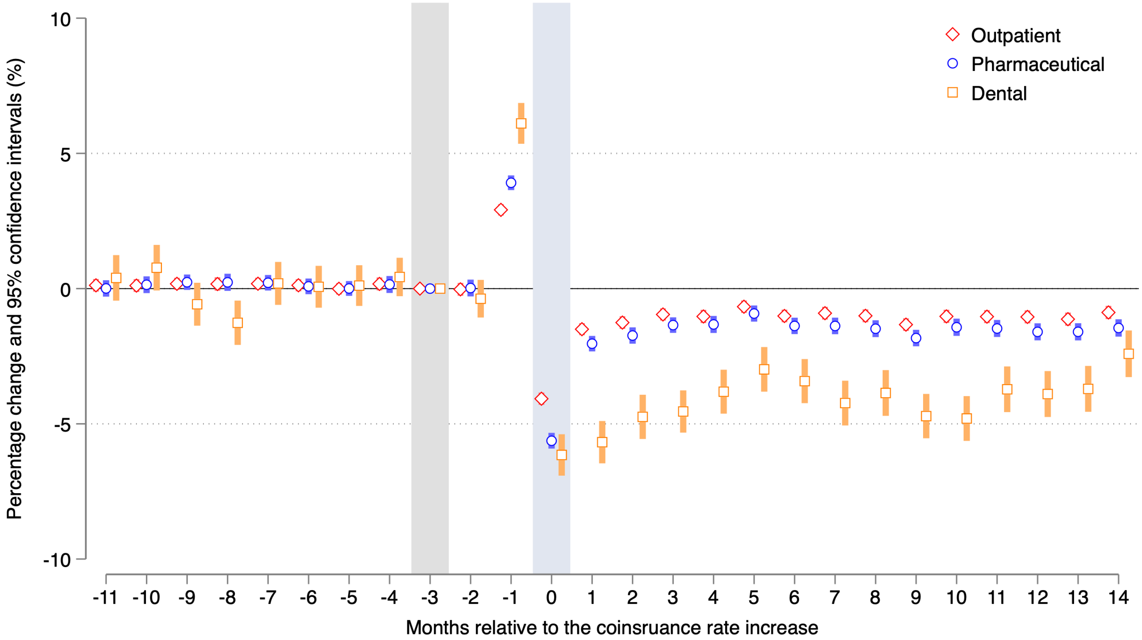

(a) Utilization rates (extensive margin)

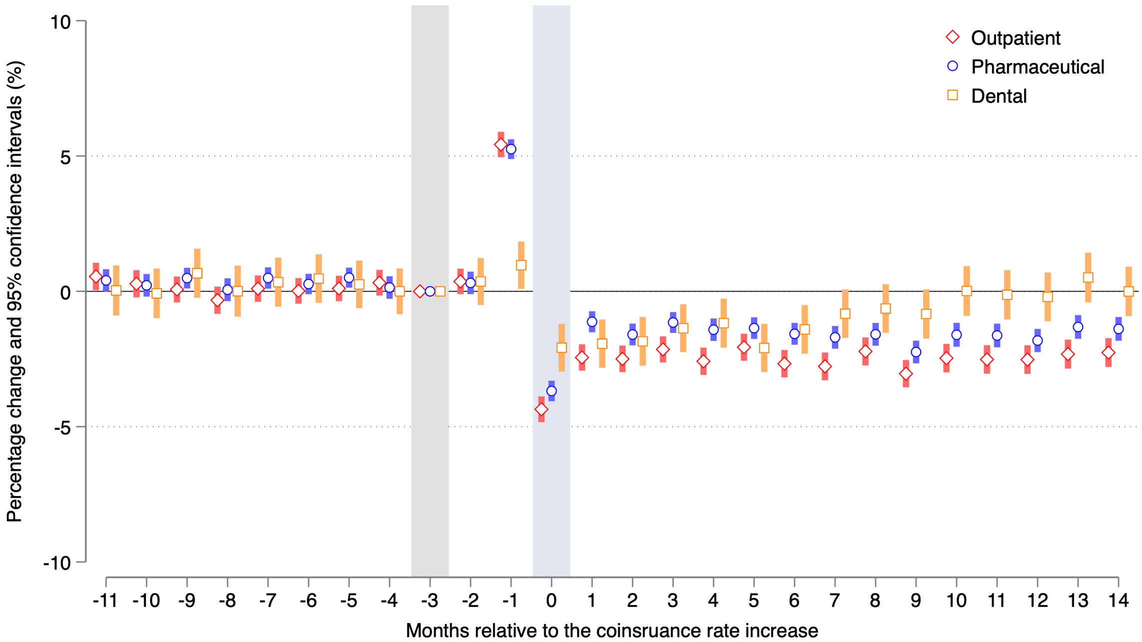

(b) Costs, log (intensive margin)

Figure 6: Heterogeneous Responses for Outpatient, Pharmaceutical, and Dental Care

Notes: Event study estimates for outpatient (red diamonds), pharmaceutical (blue circles), and dental (orange squares) services. Panel (a) shows utilization rates (extensive margin); Panel (b) shows log expenditures (intensive margin). The gray shaded region marks $k = -3$ (July 2022, the omitted reference period); the light blue shaded region marks $k = 0$ (October 2022, implementation); the unshaded region between them ($k \in \{-2, -1\}$) constitutes the anticipation period. All models include individual, month, and municipality-by-month fixed effects. Standard errors are clustered at the individual level.

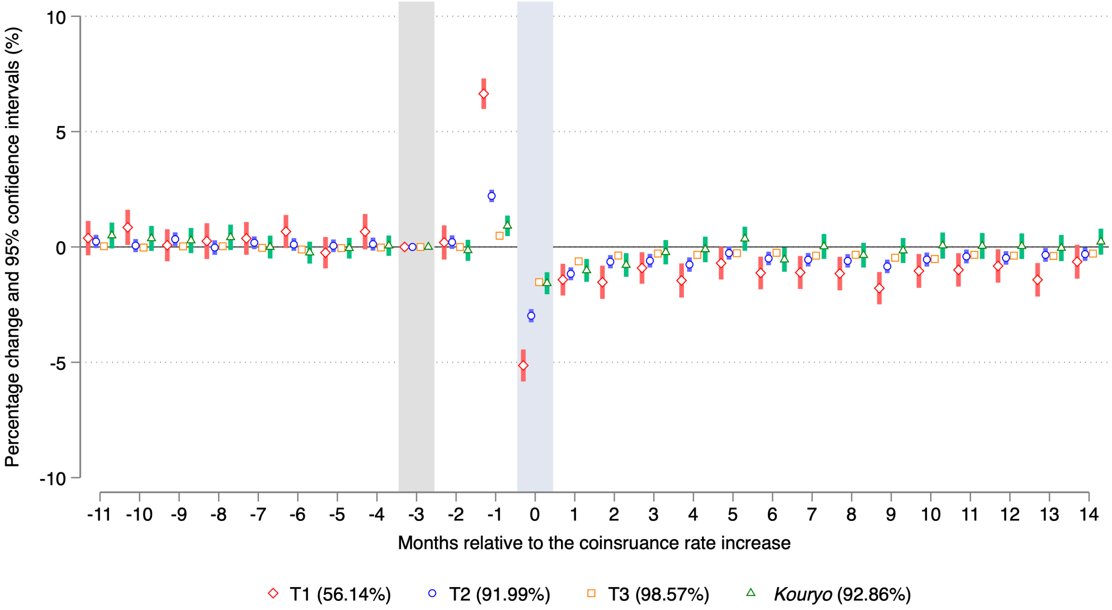

(a) Healthcare utilization rate (extensive margin) by pre-treatment utilization group

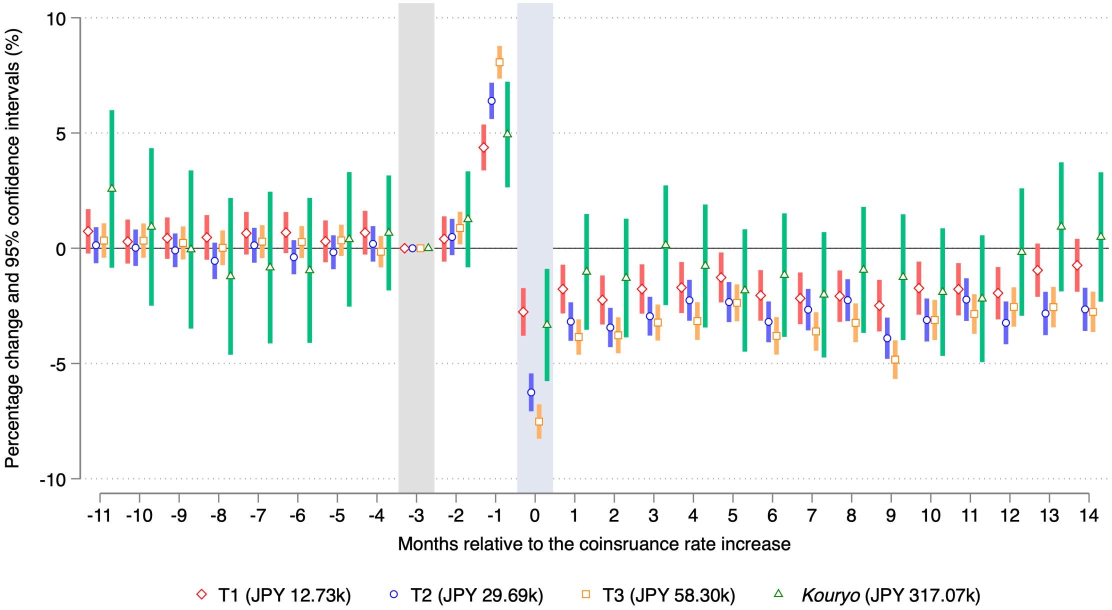

(b) Healthcare costs, log (intensive margin) by pre-treatment utilization group

Figure 7: Heterogeneous Responses by Medical Needs

Notes: Event study estimates by pre-treatment utilization group. Red diamonds: T1 (low utilization, mean expenditure JPY 12,730, baseline utilization rate 56.14%); blue circles: T2 (middle, JPY 29,690, 91.99%); orange squares: T3 (high, JPY 58,300, 98.57%); green triangles: *Kōryō* (High-Cost Medical Expense Benefit eligible, JPY 317,070, 92.86%). The gray shaded region marks $k = -3$ (July 2022, the omitted reference period); the light blue shaded region marks $k = 0$ (October 2022, implementation); the unshaded region between them ($k \in \{-2, -1\}$) constitutes the anticipation period. Subsamples are constructed as follows: (1) observations with monthly medical expenditures $\geq$ JPY 576,000 prior to $k = -3$ (July 2022) are classified as *Kōryō*-eligible; (2) patients eligible at least once form the *Kōryō* group; (3) the remaining patients are divided into three tertiles (T1, T2, T3) based on average outpatient visit days prior to $k = -3$. All models include individual, month, and municipality-by-month fixed effects. Standard errors are clustered at the individual level.

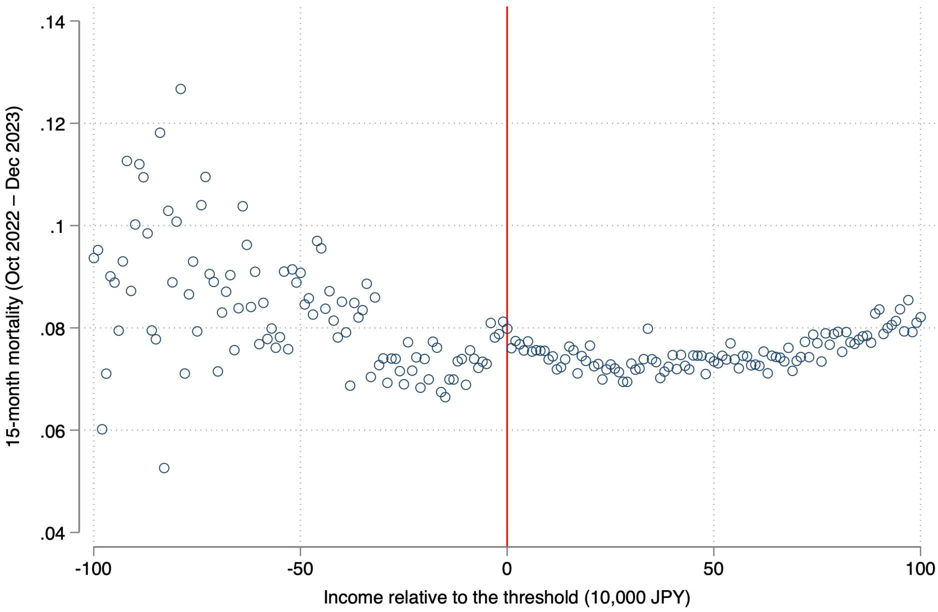

(a) Normalized income and 15-month mortality (full sample)

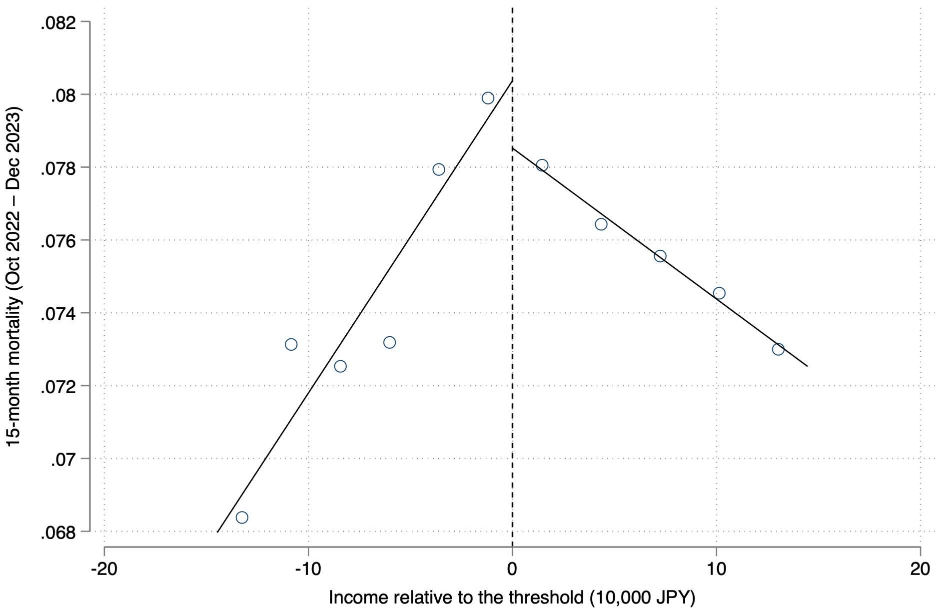

(b) RDD estimate of 15-month cumulative mortality

Figure 8: Effects on 15-Month Cumulative Mortality

Notes: Panel (a) plots the relationship between normalized income and 15-month mortality (probability of sample attrition from $k = 0$ to $k = 14$, October 2022 to December 2023) across the full income distribution; this descriptive plot is not restricted to the RDD bandwidth. Panel (b) presents the regression discontinuity estimate within the optimal bandwidth selected by `rdrobust`. Dots represent binned averages; solid lines are local linear fits estimated separately on each side of the cutoff (dashed vertical line). Mortality is proxied by sample attrition, which arises primarily from death given that inter-prefectural migration among individuals aged 75+ is rare (approximately 0.3% annually; Statistics Bureau of Japan 2023) and transition to public assistance is unlikely near the income threshold. See Appendix E for validation details.

Anticipation and Adaptation: Cost-Sharing Responses and Mortality Among the Oldest-Old

Online Appendix

Rong Fu, Masato Oikawa, Akira Kawamura, Haruko Noguchi

A Coinsurance Rate Determination Process

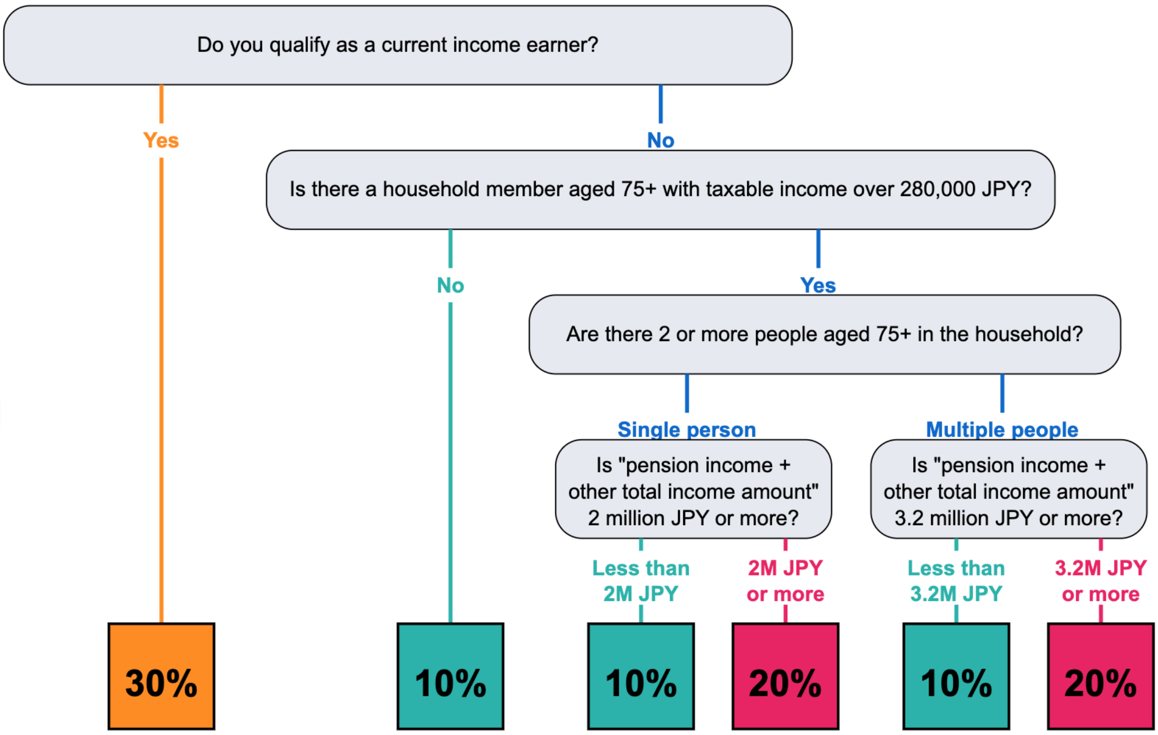

Figure A.1: Coinsurance Rate Determination Process for Population Aged 75+

Notes: This flowchart illustrates the decision process used to determine coinsurance rates under Japan's Medical Care System for the Elderly in the Latter Stage of Life. The process begins by identifying current income earners (who pay 30%), then assesses household members aged 75+ with taxable income exceeding 280,000 JPY. For qualifying elderly individuals, rates depend on household composition (single vs. multiple) and total income thresholds (2 million JPY for single households, 3.2 million JPY for multiple-person households). Individuals exceeding these thresholds pay a 20% coinsurance rate, while those below pay 10%. The income calculation uses: $\text{AOI} = \text{SUMI} + \max(\text{ANSI}, 0) - \text{MIRPP}$, where SUMI aggregates business, agricultural, real estate, interest, dividend, miscellaneous, capital gains, and forest income; and ANSI combines net salary income with an adjusted income deduction amount.

B Identification Validity

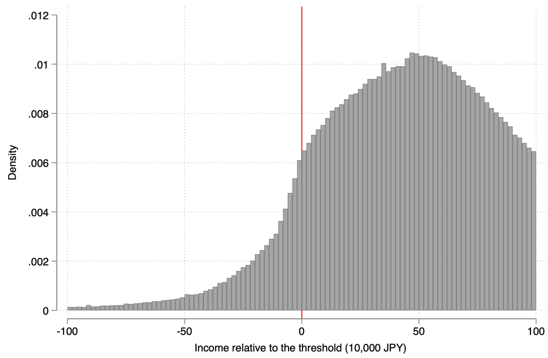

(a) Entire sample

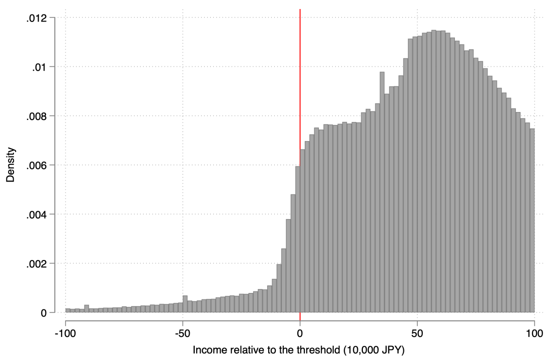

(b) Single-person Households

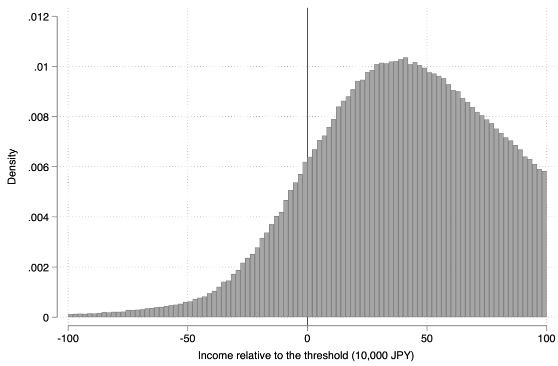

(c) Multi-person Households

Figure B.1: Distribution of the Running Variable around the Income Threshold

Notes: This figure plots the distribution of normalized income around the cutoff (defined as zero). The three panels show the full sample, single-person households, and multi-person households, respectively. There is no evidence of bunching immediately below the threshold, consistent with the absence of income manipulation. The smooth density through the cutoff supports the validity of the regression discontinuity design.

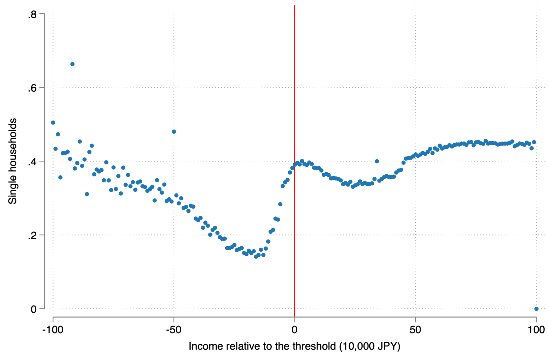

(a) Single-person Household Ratio

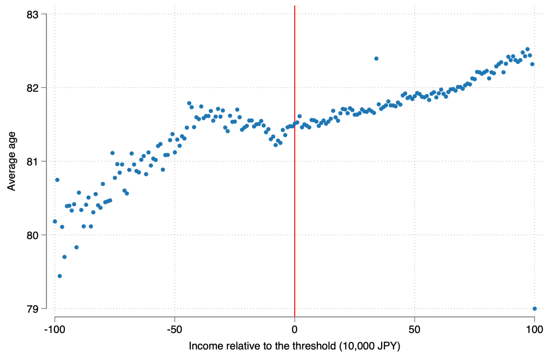

(b) Average Age

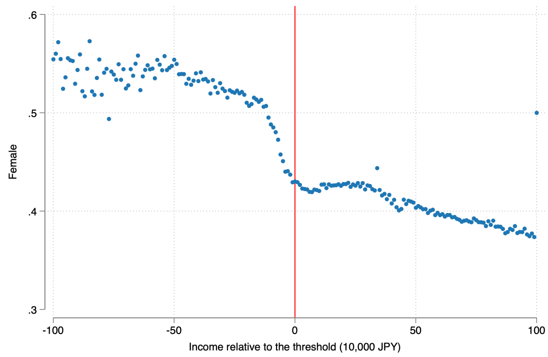

(c) Female Ratio

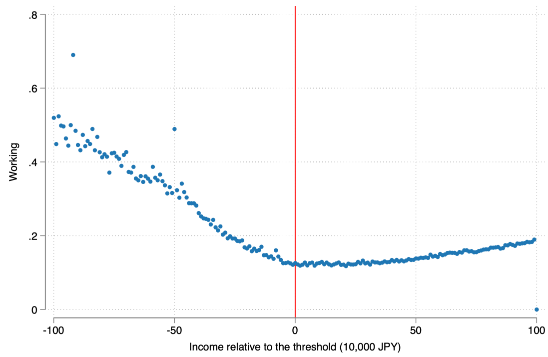

(d) Employment Ratio

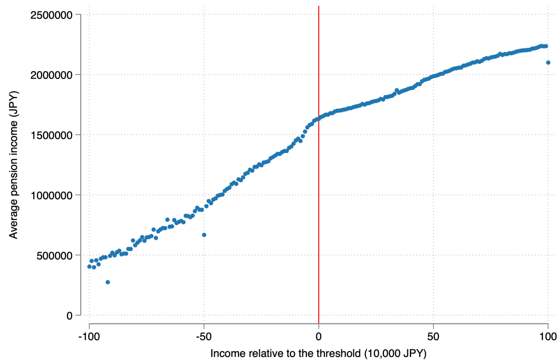

(e) Average Pension Amount (JPY)

Figure B.2: Continuity of Observable Characteristics around the Cutoff

Notes: This figure plots the mean values of five observable characteristics against normalized income, with zero corresponding to the eligibility cutoff. Each dot represents a binned average. All five covariates—single-person household ratio, average age, female ratio, employment ratio, and average pension amount—display smooth distributions without detectable discontinuities at the cutoff, supporting quasi-random assignment near the threshold.

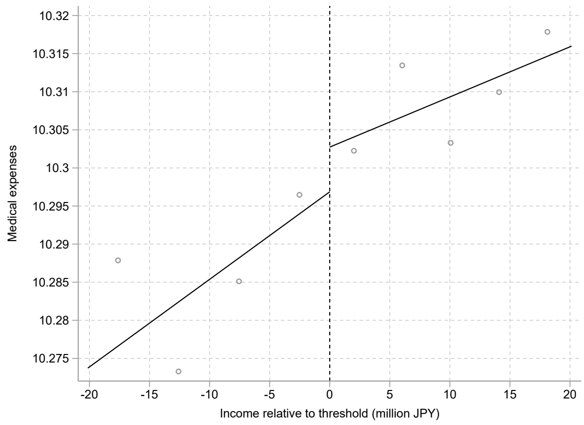

(a) $k = -3$ (July 2022)

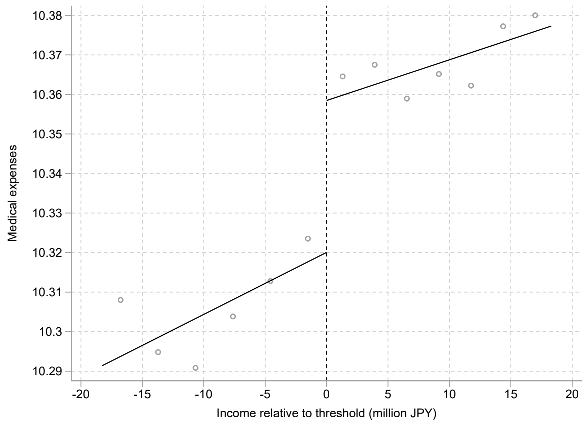

(b) $k = -1$ (September 2022)

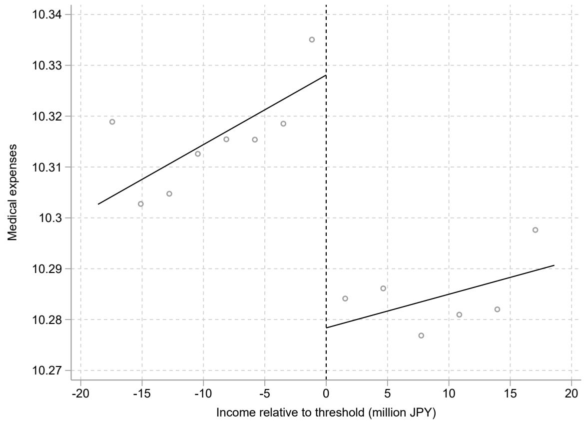

(c) $k = 0$ (October 2022)

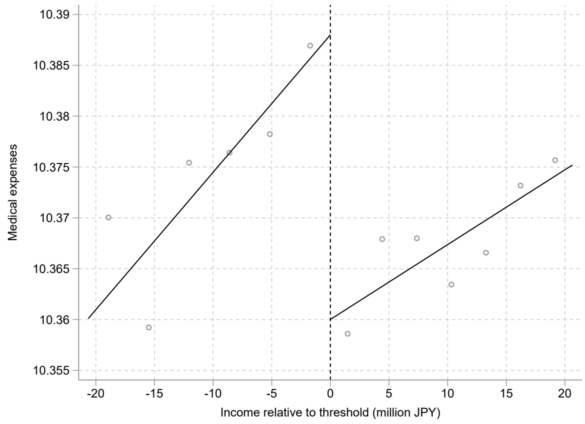

(d) $k = 5$ (March 2023)

Figure B.3: Cross-Sectional RDD Estimates by Month: Total Healthcare Costs

Notes: Cross-sectional RDD estimates for total healthcare costs (log) in four key months ($k = -3, -1, 0, +5$), corresponding to the four columns reported in Table B.1. Each panel plots binned averages (dots) and local linear fits (solid lines) estimated separately on each side of the income threshold (dashed vertical line). The bandwidth is selected by `rdrobust` with a local linear specification. Estimated discontinuities: $k = -3$ (July 2022, pre-notification baseline): $+0.4\%$ (statistically insignificant, consistent with no pre-existing difference); $k = -1$ (September 2022, anticipation): $+3.6\%$ (last-minute surge); $k = 0$ (October 2022, implementation): -5.0%; $k = 5$ (March 2023, steady state): -3.0%. These estimates are free of the DD parallel-trends assumption and independently corroborate the event-study findings.

Table B.1: Cross-Sectional RDD Estimates by Month with Statistical Inference: Total Healthcare Costs

	(1) Jul. 2022	(2) Sep. 2022	(3) Oct. 2022	(4) Mar. 2023
RD Estimates	0.004	0.036***	-0.050***	-0.030***
	(0.007)	(0.008)	(0.008)	(0.008)
Observations left	173,723	166,874	170,199	178,046
Observations right	358,905	334,212	326,769	370,826
Bandwidth left	20.134	18.294	18.596	20.648
Bandwidth right	20.134	18.294	18.596	20.648

Notes: This table reports conventional local linear regression discontinuity estimates of the effect of the increase in the coinsurance rate on monthly medical expenditures for four key months ($k = -3, -1, 0, +5$), corresponding to the four panels in Figure B.3. The running variable is income normalized relative to the applicable household-type-specific income threshold: 2 million JPY for single-person households and 3.2 million JPY for multi-person households. Individuals below the cutoff remained subject to the 10% coinsurance rate, while those above the cutoff became subject to the 20% coinsurance rate. The models are estimated separately on each side of the cutoff, with the optimal bandwidth selected by `rdrobust`. Robust bias-corrected standard errors are reported in parentheses. Medical expenditures are measured in Japanese yen and include total insurer-recorded healthcare spending during the corresponding month. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C Robustness Checks

Table C.1: Regression Discontinuity Estimates for 15-Month Mortality

	(1)
RD Estimate	-0.002
	(0.002)
Observations left	176,788
Observations right	306,615
Bandwidth left	14.468
Bandwidth right	14.468

Notes: This table reports conventional local linear regression discontinuity estimates of the effect of the increase in the coinsurance rate on 15-month mortality, proxied by sample attrition from $k = 0$ to $k = 14$ (October 2022 to December 2023). The running variable is income normalized relative to the applicable household-type-specific income threshold: 2 million JPY for single-person households and 3.2 million JPY for multi-person households. Individuals below the cutoff remained subject to the 10% coinsurance rate, while those above the cutoff became subject to the 20% coinsurance rate. The models are estimated separately on each side of the cutoff, with the optimal bandwidth selected by `rdrobust`. Robust bias-corrected standard errors are reported in parentheses. Mortality is proxied by sample attrition, which arises primarily from death given that inter-prefectural migration among individuals aged 75+ is rare, approximately 0.3% annually Statistics Bureau of Japan 2023, and transition to public assistance is unlikely near the income threshold. See Appendix E for validation. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

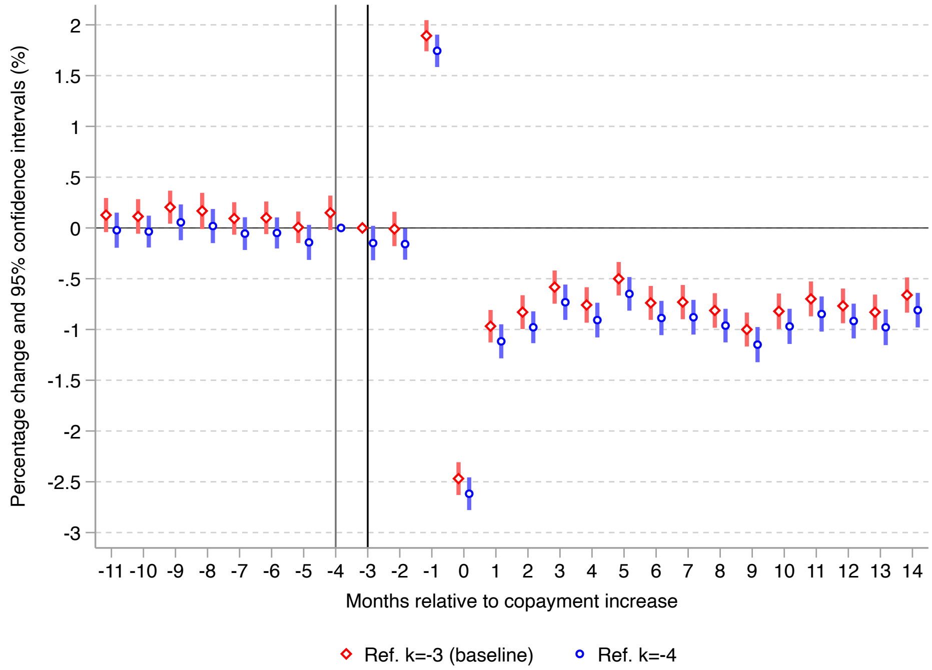

(a) Healthcare utilization rate (extensive margin)

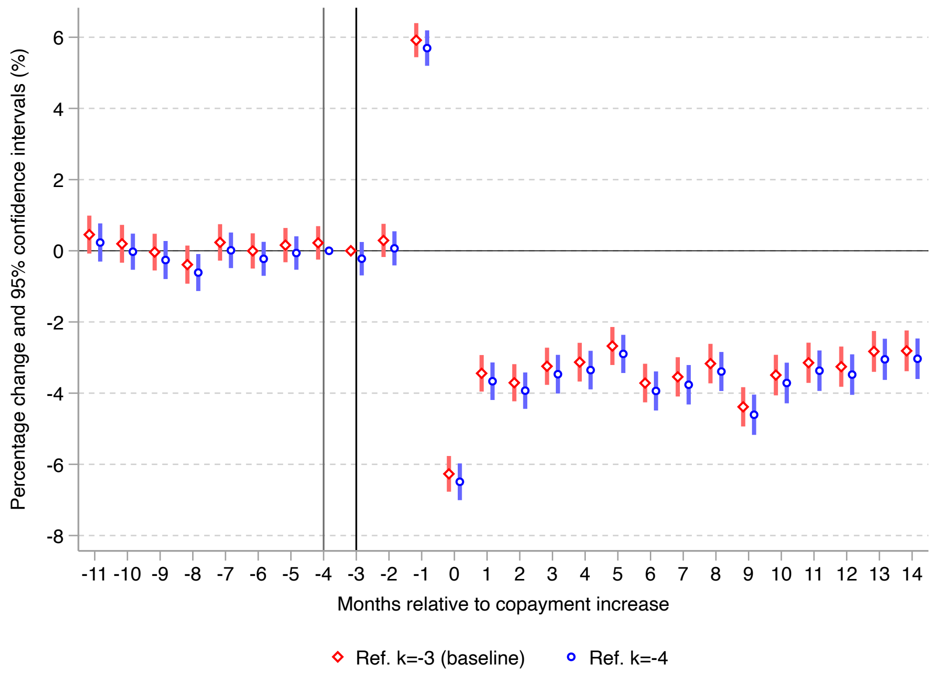

(b) Healthcare costs, log (intensive margin)

Figure C.1: Robustness: Alternative Reference Period

Notes: This figure compares the baseline event-study estimates, which use $k = -3$ (July 2022) as the omitted reference period, with estimates that instead use $k = -4$ (June 2022) as the omitted reference period. Panel (a) reports estimates for healthcare utilization rates (extensive margin), and Panel (b) reports estimates for log healthcare costs (intensive margin). The specification is otherwise identical to the main specification in equation (1), including individual, month, and municipality-by-month fixed effects. Estimates are expressed as percentage changes relative to the treatment group's pre-treatment mean for binary outcomes and as approximate percentage changes for log-transformed expenditure outcomes. Standard errors are clustered at the individual level. The similarity between the two series indicates that the main findings are not sensitive to the choice of omitted pre-notification reference month.

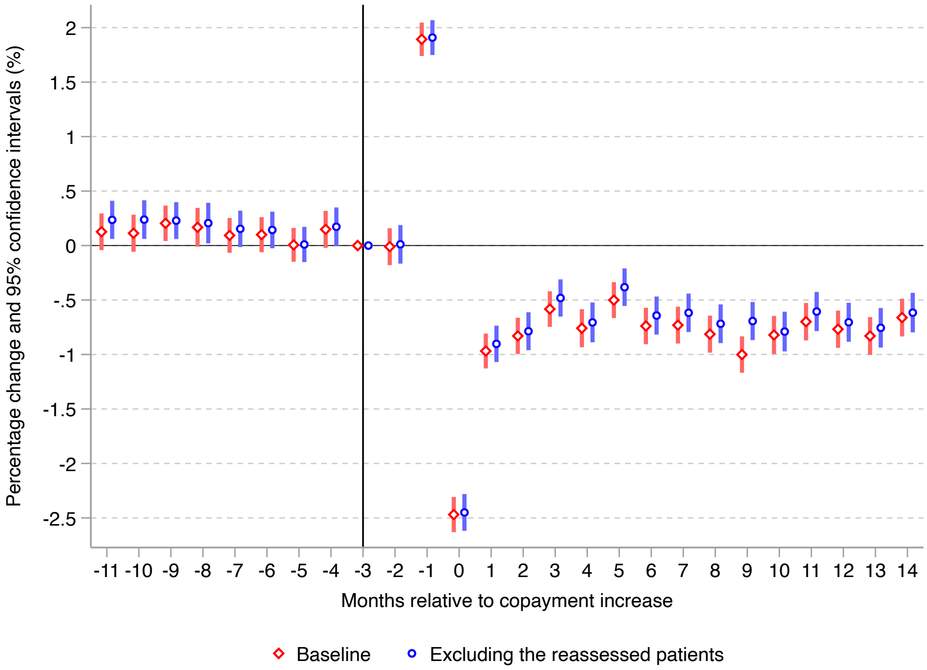

(a) Healthcare utilization rate (extensive margin)

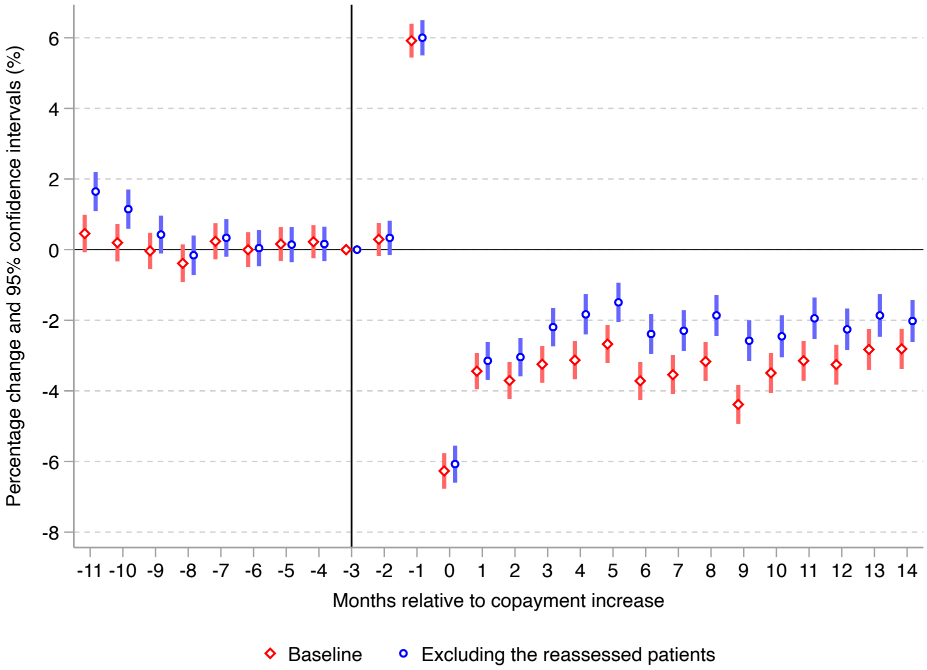

(b) Healthcare costs, log (intensive margin)

Figure C.2: Robustness: Excluding Individuals with Post-Implementation Changes in Coinsurance Status

Notes: This figure compares the baseline event-study estimates with estimates that exclude individuals whose coinsurance status ever changed after the October 2022 implementation. Among individuals observed in July 2023, 16.85% experienced a change in coinsurance status at some point after implementation, with most first changes occurring in August 2023. Panel (a) reports estimates for healthcare utilization rates (extensive margin), and Panel (b) reports estimates for log healthcare costs (intensive margin). The specification is otherwise identical to the main specification in equation (1), including individual, month, and municipality-by-month fixed effects. Estimates are expressed as percentage changes relative to the treatment group's pre-treatment mean for binary outcomes and as approximate percentage changes for log-transformed expenditure outcomes. Standard errors are clustered at the individual level. The results remain similar after excluding these individuals, indicating that the main estimates are not driven by subsequent changes in coinsurance status.

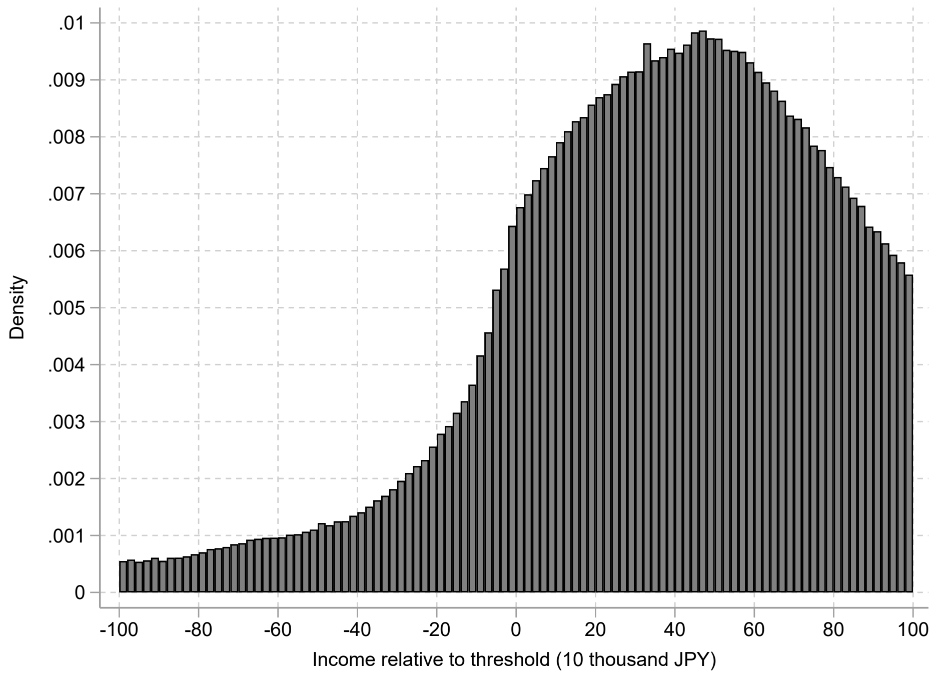

Figure C.3: Post-Period Income Distribution around the Threshold

Notes: This figure plots the distribution of normalized income in August 2023 around the applicable household-income threshold. Income is measured relative to the household-specific threshold and reported in units of 10,000 JPY. The sample is restricted to individuals with baseline normalized income within 2 million JPY of the threshold as of October 2022. If individuals adjusted income to avoid the higher coinsurance rate, bunching would be expected just below the threshold. The absence of visible bunching below the cutoff provides no evidence of selective post-implementation income manipulation around the threshold.

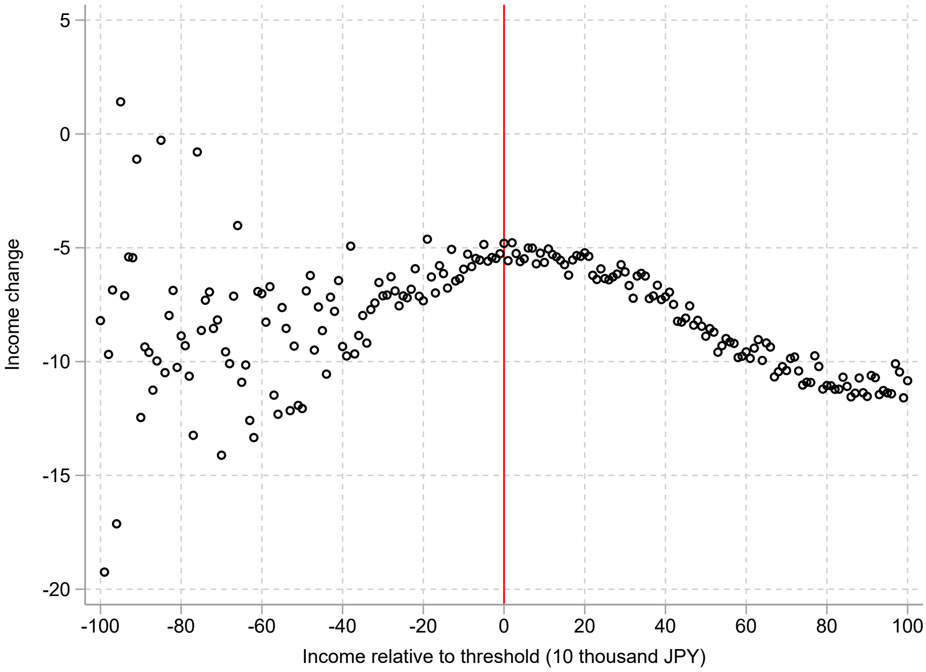

(a) Income change by baseline normalized income

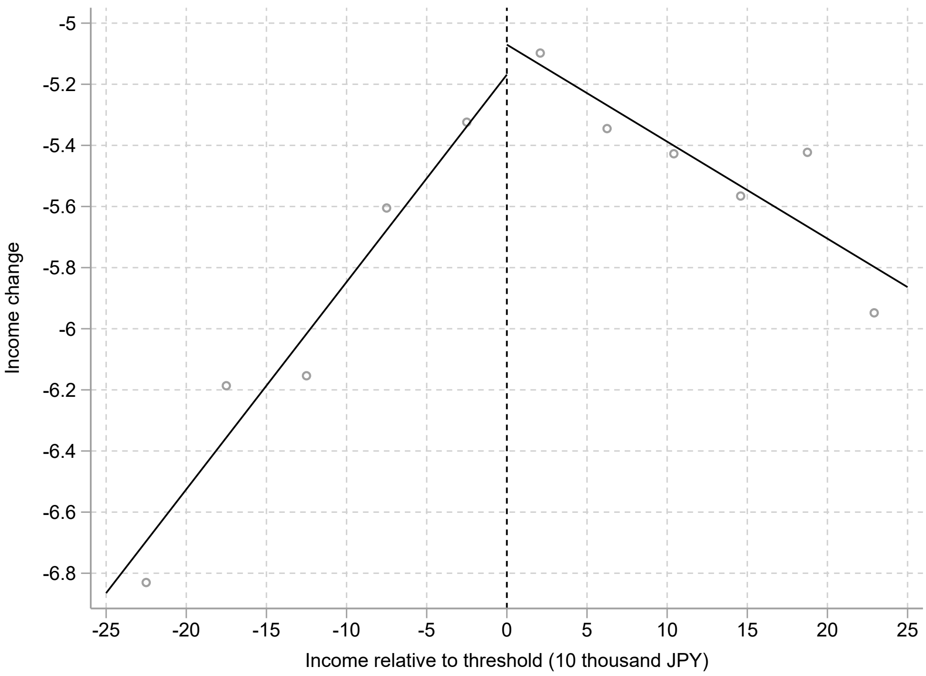

(b) RDD estimate for income change

Figure C.4: Post-Implementation Income Changes around the Threshold

Notes: This figure examines whether individuals near the income threshold selectively changed their income after the October 2022 implementation. The sample is restricted to individuals with baseline normalized income within 2 million JPY of the applicable household-income threshold as of October 2022. Panel (a) plots the change in normalized income between October 2022 and August 2023 against baseline normalized income. Panel (b) presents a regression discontinuity plot of income change at the baseline threshold, using baseline normalized income as the running variable. Income is measured relative to the applicable household-specific threshold and reported in units of 10,000 JPY. If individuals just above the cutoff selectively reduced their income after implementation to avoid the higher coinsurance rate, income changes would display a negative discontinuity at the cutoff. The absence of such a pattern provides no evidence of selective post-implementation income reduction among individuals just above the threshold.

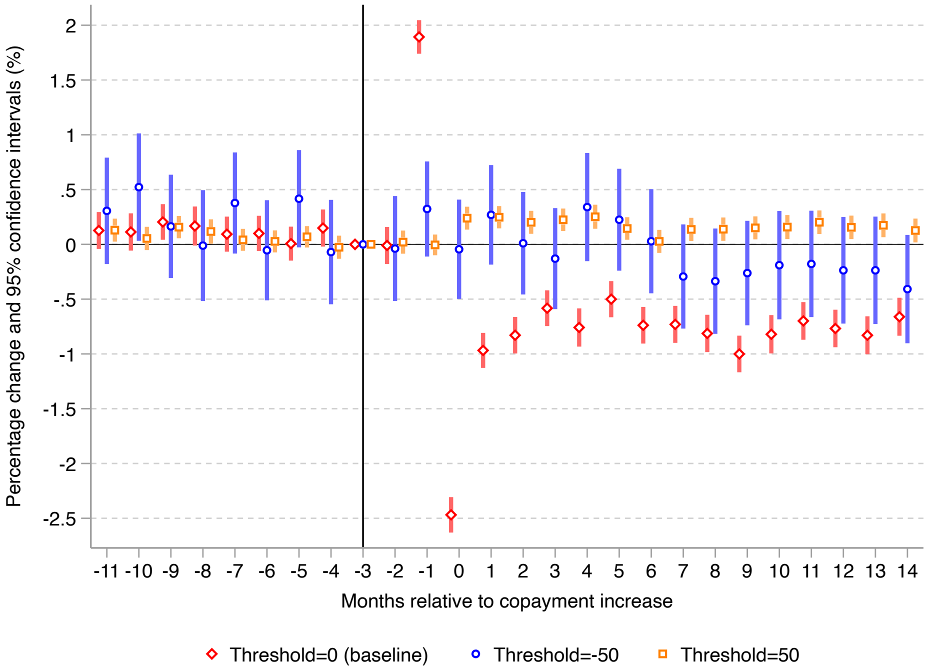

(a) Healthcare utilization rate (extensive margin)

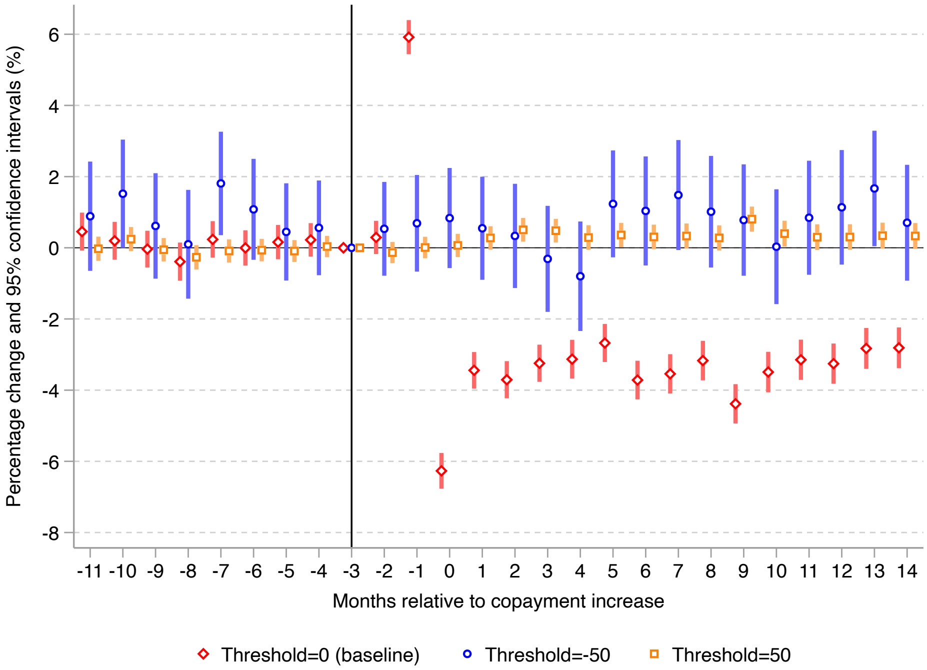

(b) Healthcare costs, log (intensive margin)

Figure C.5: Robustness: Placebo Income Cutoffs

Notes: This figure reports placebo cutoff checks using artificial income thresholds at which no policy-induced change in coinsurance rates occurred. The baseline estimates use the actual threshold normalized to zero and the main analytical sample around this cutoff. The placebo estimates use artificial thresholds located 0.5 million JPY below and above the actual cutoff. For the lower placebo cutoff, the sample is restricted to individuals with normalized income between −0.99 and −0.01 million JPY, and treatment is redefined as being above the placebo cutoff at −0.5 million JPY. For the upper placebo cutoff, the sample is restricted to individuals with normalized income between 0.01 and 0.99 million JPY, and treatment is redefined as being above the placebo cutoff at 0.5 million JPY. These restrictions ensure that each placebo comparison is conducted within regions where the actual coinsurance rate does not change at the placebo threshold. Panel (a) reports estimates for healthcare utilization rates (extensive margin), and Panel (b) reports estimates for log healthcare costs (intensive margin). Each placebo specification preserves the same event-study structure as in equation (1). All models include individual, month, and municipality-by-month fixed effects. Standard errors are clustered at the individual level. The absence of systematic effects at the placebo cutoffs supports the interpretation that the baseline estimates are driven by the actual policy-induced discontinuity in coinsurance rates rather than by smooth income-related trends.

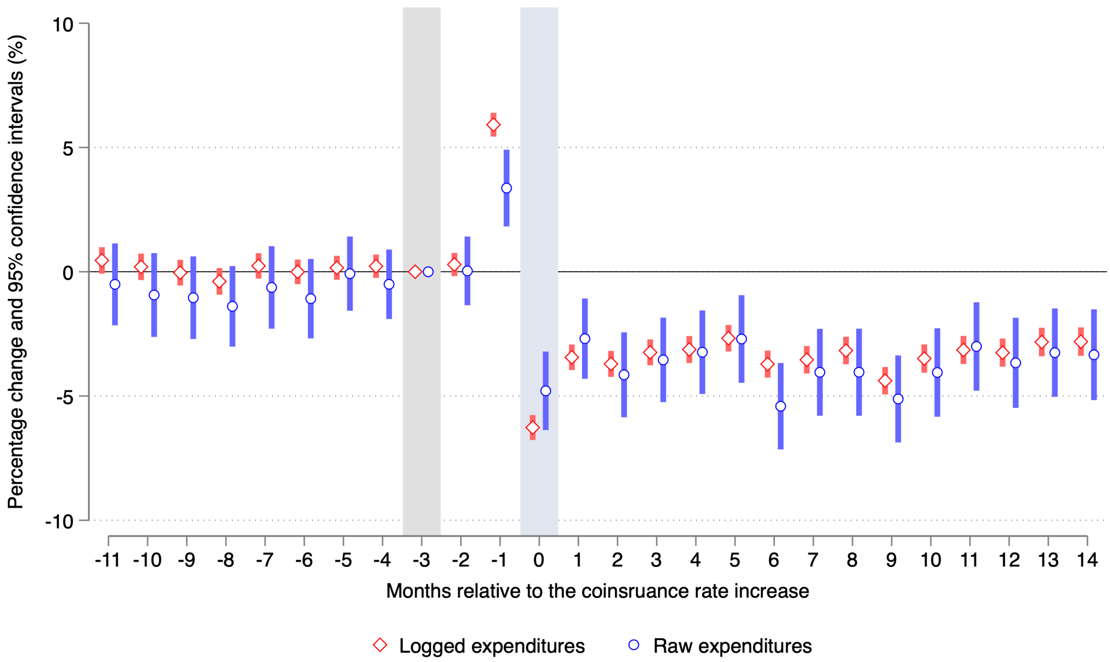

Figure C.6: Robustness: Log Specification vs. Raw Level Values

Notes: Comparison of event study estimates from the log specification (red diamonds) and raw-level specification expressed as percentage changes relative to the pre-treatment mean (blue circles). The gray shaded region marks $k = -3$ (July 2022, the omitted reference period); the light blue shaded region marks $k = 0$ (October 2022, implementation); the unshaded region between them ($k \in \{-2, -1\}$) constitutes the anticipation period. Both specifications produce consistent patterns. The log specification yields tighter confidence intervals; the raw-level specification reveals a slightly negative pre-trend that the log transformation eliminates, supporting the log specification as the preferred approach.

D Out-of-Pocket Cost Analysis

This appendix documents two complementary dimensions of the reform's financial impact on patients. First, it quantifies how the temporary outpatient relief measure attenuated the effective price increase faced by patients. Second, it examines the temporal pattern of out-of-pocket burden relative to patients' baseline income.

D.1 Actual Price Increase after the Outpatient Relief Measure

The nominal coinsurance rate increase from 10% to 20% implies a 100% increase in the patient price for a given level of healthcare utilization. For outpatient care, however, the reform included a temporary relief measure that capped the monthly increase in out-of-pocket payments at 3,000 JPY. This safeguard implies that the effective increase in outpatient cost-sharing was smaller than the nominal 10-percentage-point increase for patients with sufficiently high monthly outpatient spending.

To quantify the magnitude of this attenuation, we use pre-policy outpatient utilization among individuals in the treatment group. Abstracting from the high-cost medical expense benefit, we define the nominal monthly increase in outpatient out-of-pocket payments as

$$\Delta_{it}^{\text{nom}} = 0.1 \times M_{it}^{\text{out}},$$

where (M_{it}^{out}) denotes total outpatient medical expenditure for individual (i) in month (t). The actual monthly increase after the relief measure is

$$\Delta_{it}^{\text{act}} = \min \Delta_{it}^{\text{nom}}, 3000.$$

Because the cap operates at the person-month level, we apply this calculation before aggregating across months or individuals.

Table D.1 summarizes the implied attenuation of the outpatient price increase. Among treated outpatient-months in the pre-policy period, the 3,000 JPY cap would have been binding in 20.6% of months. At the individual level, 53.0% of treated outpatient users would have experienced a binding cap at least once. The effective monthly price increase, measured as the ratio of average actual to nominal increases, was 61.0% of the nominal increase. The distribution, however, is highly skewed: the median individual-level actual-to-nominal ratio was 98.2%. This contrast indicates that the cap had little effect for the median patient but substantially attenuated the price increase for high-utilization patients in the upper tail of outpatient spending.

Table D.1: Outpatient Relief Cap and Effective Price Increase

Measure	Estimate	Interpretation
Binding cap among treated outpatient-months	20.6%	Share of treated months with outpatient use
Treated outpatient users ever facing binding cap	53.0%	Share of treated individuals with outpatient use
Mean actual-to-nominal ratio	61.0%	Ratio of average monthly actual to nominal increase
Median individual actual-to-nominal ratio	98.2%	Median of individual-level actual-to-nominal ratios

Notes: This table summarizes the implied effect of the outpatient transitional relief measure among treated individuals in the pre-policy period. The sample is restricted to individuals within the ($\pm$)500,000 JPY income bandwidth around the applicable household-income threshold and to months before August 2022. The nominal monthly increase is calculated as 10% of total outpatient medical expenditure. The actual monthly increase applies the statutory 3,000 JPY cap at the person-month level. The calculations abstract from the high-cost medical expense benefit. The mean actual-to-nominal ratio is computed as the ratio of average monthly actual to nominal increases. The median actual-to-nominal ratio is computed from individual-level ratios.

These results imply that elasticities based on the nominal doubling of the coinsurance rate are conservative in absolute value for outpatient care, because the effective outpatient price increase was smaller than the full nominal increase. At the same time, the attenuation does not overturn the basic interpretation of the estimates. Even after the relief measure, the effective outpatient price increase remained substantial on an aggregate amount basis. Thus, the reduced-form estimates should be interpreted as responses to the policy-induced effective cost-sharing increase, inclusive of the outpatient relief measure, rather than as responses to an unconstrained full doubling of outpatient out-of-pocket prices.

D.2 Out-of-Pocket Burden Relative to Income

Figure D.1 presents event-study estimates for out-of-pocket spending as a share of baseline monthly income, using the same difference-in-differences specification as the main analysis. This exercise complements the mechanical calculations above by showing how patients' realized out-of-pocket burden evolved around implementation.

Three features of the figure are notable. First, prior to notification ($k \leq -3$), the differential in out-of-pocket burden between treatment and control groups is close to zero, consistent with the parallel-trends assumption underlying the main analysis and suggesting no pre-existing differential trend in financial exposure.

Second, at implementation ($k = 0$), out-of-pocket burden rises sharply by approximately 57–60% relative to baseline monthly income. This increase is smaller than a literal doubling of out-of-pocket payments, consistent with the outpatient relief cap and other institutional safeguards, but it remains large. The result confirms that the reform generated an immediate and economically meaningful

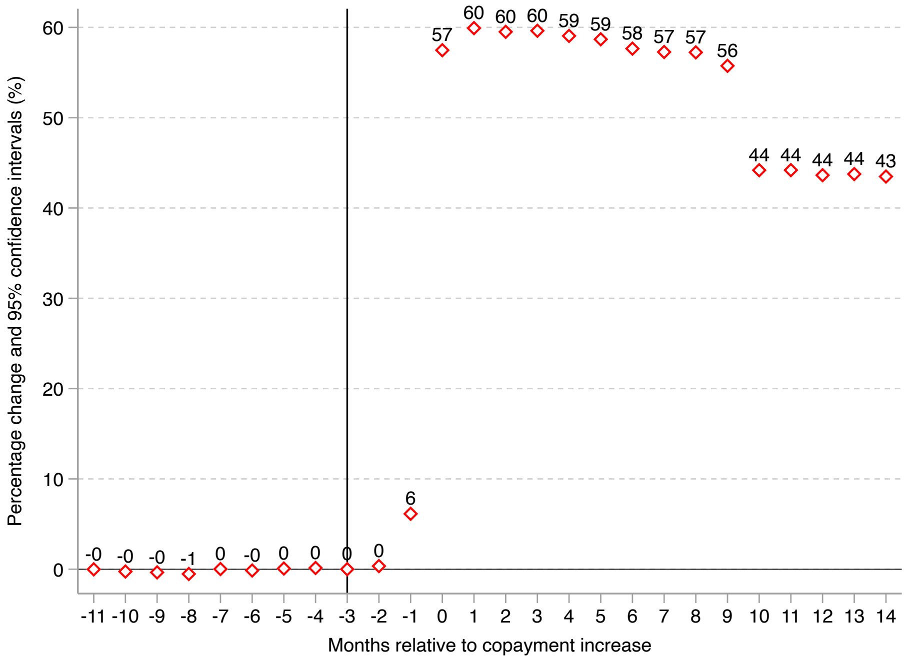

Figure D.1: Effect on Out-of-Pocket Burden Relative to Baseline Income

Notes: This figure presents event-study estimates for out-of-pocket medical spending relative to baseline monthly income. The outcome is (log(out-of-pocket spending/baseline monthly income)), where baseline monthly income equals annual household income at ($k = -3$) divided by 12. Observations with zero out-of-pocket spending are excluded. The solid black vertical line marks ($k = -3$) (July 2022, the omitted reference period), and ($k = 0$) denotes October 2022 (implementation). Estimates are approximate percentage changes. All models include individual, month, and municipality-by-month fixed effects. Standard errors are clustered at the individual level.

increase in patients' financial exposure.

Third, the burden attenuates modestly over the post-implementation period, stabilizing at approximately 44–56% above baseline in later months. This partial attenuation is mechanically consistent with the utilization reductions documented in the main analysis: as patients reduce healthcare use in response to higher cost sharing, realized out-of-pocket spending falls relative to what would have occurred under unchanged utilization. The fact that the burden remains substantially elevated throughout the post-implementation period supports the interpretation of steady-state utilization reductions as responses to a persistent increase in effective patient prices rather than as a purely transitory adjustment.

E Validation of Sample Attrition as a Mortality Proxy

This appendix provides detailed evidence supporting our use of sample attrition as a proxy for mortality (Section 3.4). We present three complementary strands of evidence: (i) a comparison of attrition counts against official national death counts, (ii) healthcare utilization trajectories in the months preceding attrition, and (iii) disease composition among those who subsequently attrit.

E.1 Comparison with Official Death Records

Table E.1 compares, for calendar year 2022, the number of individuals who permanently exit our dataset against official national death counts by age group and sex published by the Ministry of Health, Labour and Welfare.[20] The LSEH system covers approximately 98.6% of Japan's population aged 75 and above, so in the absence of non-mortality attrition one would expect our counts to equal roughly 99% of the official figures.

Table E.1: Comparison of Dataset Attrition with Official Death Counts, 2022

	Official Deaths (A)			Our Attrition (B)			
Age group	Total	Male	Female	Total	Male	Female	Ratio B/A
75–79	170,192	109,955	60,237	163,278	96,855	66,423	0.959
80–84	247,487	144,728	102,759	198,206	113,000	85,206	0.801
85–89	320,647	159,604	161,043	248,737	122,464	126,273	0.776
90–94	295,883	110,745	185,138	224,817	83,694	141,123	0.760
95–99	155,165	38,692	116,473	115,609	28,799	86,810	0.745
100+	40,183	5,608	34,575	29,113	4,157	24,956	0.725
Total 75+	1,229,557	569,332	660,225	979,760	448,969	530,791	0.797

Notes: "Official Deaths" are from Japan's *Vital Statistics 2022* (Vital Statistics of Japan 2022), restricted to deaths of individuals aged 75 and above. "Our Attrition" counts individuals who permanently exit the MCD-Tx dataset during 2022. As described in the text, new entrants at age 75 are excluded from the attrition count for the 75–79 group because their first appearance in the data may reflect relocation rather than entry into LSEH. The ratio B/A ranges from 0.725 to 0.959 across age groups, consistent with LSEH covering approximately 98.6% of the 75+ population and a small share of exits reflecting non-mortality attrition.

Several patterns emerge. First, across all age groups, our attrition counts are a substantial fraction of the official death figures—ranging from 73% to 96%—consistent with attrition being driven primarily by mortality rather than administrative exit. Second, the ratio of our count to the official count is highest for the 75–79 group (overall ratio 0.96), and then stabilizes in the 0.72–0.80 range for older cohorts. The somewhat lower ratios at older ages likely reflect two factors: (a) individuals aged 80 and above who transitioned from regular health insurance into LSEH at age 75 but subsequently shifted to public assistance (*seikatsu hogo*), a non-mortality exit route that is rare but not zero at older

[20] Official counts are drawn from the *Vital Statistics of Japan 2022* (Vital Statistics of Japan 2022). The sample is not restricted to the analytical bandwidth. Note that the 75–79 age group is not directly comparable: individuals first appearing in the dataset at age 75 cannot be distinguished from those who relocated into a new prefecture at that age, so we exclude new entrants at age 75 from the attrition count for this group. The remaining five age groups are unaffected by this issue.

ages; and (b) inter-prefectural moves accompanied by a change in individual identifiers, which are more common in the 75–79 group precisely because that group contains new LSEH entrants for whom no prior within-dataset trajectory is observed. Third, within the five age groups unaffected by the new-entrant issue (80 and above), male and female ratios track closely within each age group, indicating no systematic sex-specific bias in the attrition measure. In the 75–79 group, by contrast, the sex-specific ratios diverge—the female ratio slightly exceeds one—reflecting the same new-entrant and relocation dynamics that make this cohort's attrition count not directly comparable to official deaths (see note to Table E.1), rather than any sex-specific measurement problem. Taken together, the high and consistent ratios confirm that sample attrition in our dataset is a valid, though not perfect, proxy for mortality.

E.2 Healthcare Utilization Trajectories Prior to Attrition

Tables E.2 and E.3 report the mean utilization of individuals in the months leading up to their final appearance in the dataset, indexed by months relative to the exit month ($attr = 0$). The sample is not restricted to the analytical bandwidth.

Inpatient care. Table E.2 reports the probability of any inpatient admission (extensive margin) and the mean number of inpatient days (intensive margin) for months $attr = -11$ through $attr = 0$.

Table E.2: Inpatient Utilization Trajectories Prior to Sample Exit

Months before exit ($attr$)	Inpatient probability (extensive margin)	Mean inpatient days (intensive margin)
-11	0.125	2.70
-10	0.130	2.85
-9	0.133	2.95
-8	0.138	3.10
-7	0.143	3.25
-6	0.152	3.50
-5	0.163	4.25
-4	0.183	4.70
-3	0.208	5.40
-2	0.262	6.25
-1	0.333	7.90
0	0.422	9.80
Change (0 vs. -11)	+29.7 pp	+7.10 days

Notes: Means are computed across all individuals who subsequently exit the dataset (not restricted to the analytical bandwidth), indexed by months relative to their final observed month ($attr = 0$). Both the probability of any inpatient admission and mean inpatient days increase monotonically and accelerate in the final three months, consistent with end-of-life hospitalization. "pp" denotes percentage points.

Both series increase monotonically and accelerate markedly in the final months. The probability of hospitalization rises from 12.5% eleven months before attrition to 42.2% in the final month—a

3.4-fold increase. The steepest acceleration occurs in the last four months: admission probability roughly doubles from 20.8% at *attr* $= -3$ to 42.2% at *attr* $= 0$. Mean inpatient days follow a similar pattern, rising from 2.70 days at *attr* $= -11$ to 9.80 days at *attr* $= 0$, with the majority of the increase concentrated in the final three months. This accelerating end-of-life escalation in inpatient care is inconsistent with administrative relocation—which would generate no systematic change in utilization before exit—and is precisely the pattern expected as individuals enter terminal inpatient care.

Outpatient care. Table E.3 presents analogous trajectories for outpatient services. The outpatient visit probability declines monotonically from 77.5% at *attr* $= -11$ to 65.0% at *attr* $= 0$, with the pace of decline accelerating in the final two months. This pattern is consistent with a progressive shift from ambulatory outpatient care toward inpatient hospitalization as health deteriorates: individuals who are admitted to hospital cease attending outpatient clinics, generating a mechanical negative correlation between the two series. The sharp drop in outpatient probability from 71.0% at *attr* $= -1$ to 65.0% at *attr* $= 0$ mirrors precisely the sharp jump in inpatient admission probability (33.3% to 42.2%) over the same interval, consistent with a transfer of care setting.

Table E.3: Outpatient Utilization Trajectories Prior to Sample Exit

Months before exit (*attr*)	Outpatient probability (extensive margin)	Mean outpatient days (intensive margin)
-11	0.775	2.167
-10	0.775	2.188
-9	0.770	2.150
-8	0.766	2.083
-7	0.763	2.110
-6	0.760	2.128
-5	0.756	2.130
-4	0.750	2.145
-3	0.745	2.148
-2	0.735	2.155
-1	0.710	2.163
0	0.650	2.210
Change (0 vs. -11)	-12.5 pp	$+0.043$ days

Notes: Means are computed across all individuals who subsequently exit the dataset, indexed by months relative to their final observed month (*attr* $= 0$). The outpatient visit probability declines monotonically, consistent with a progressive shift from ambulatory care to inpatient hospitalization. Mean outpatient days exhibit non-monotone variation (range: 2.08–2.21), reflecting competing forces: the mechanical reduction from hospitalization and increased specialist outpatient contact among those not yet admitted. The final-month drop in outpatient probability (from 0.710 to 0.650) coincides with the final-month surge in inpatient probability (from 0.333 to 0.422) shown in Table E.2. "pp" denotes percentage points.

Mean outpatient days exhibit a less smooth trajectory, reflecting the competing forces of declining ambulatory contact and increased specialist outpatient visits among those not yet admitted. The small variation in mean outpatient days (ranging from 2.08 to 2.21 across the 12-month window) relative to the large variation in mean inpatient days (2.70 to 9.80) underscores that it is the inpatient margin that dominates the end-of-life care transition.

E.3 Disease-Specific Trajectories Prior to Attrition

Table E.4 traces disease-specific utilization trajectories for the three leading causes of death among the Japanese elderly: cancer/neoplasms, circulatory diseases, and respiratory diseases. For each disease, we report mean inpatient days and outpatient visit probability by months prior to exit.

Table E.4: Disease-Specific Utilization Trajectories Prior to Sample Exit

	Cancer		Circulatory		Respiratory	
attr	Inp. days	Out. prob.	Inp. days	Out. prob.	Inp. days	Out. prob.
−11	0.175	0.050	0.650	0.410	0.200	0.037
−10	0.185	0.053	0.680	0.415	0.218	0.038
−9	0.195	0.058	0.720	0.407	0.240	0.038
−8	0.210	0.063	0.760	0.403	0.265	0.038
−7	0.230	0.068	0.820	0.408	0.295	0.039
−6	0.270	0.073	0.900	0.410	0.340	0.040
−5	0.330	0.080	1.000	0.412	0.400	0.041
−4	0.430	0.093	1.160	0.415	0.500	0.042
−3	0.600	0.105	1.350	0.420	0.680	0.043
−2	0.830	0.118	1.620	0.430	0.900	0.044
−1	1.100	0.138	1.920	0.448	1.120	0.047
0	1.430	0.158	2.140	0.450	1.350	0.051
Change	+1.255	+0.108	+1.490	+0.040	+1.150	+0.014

Notes: "Inp. days" denotes mean days of inpatient care for the indicated disease category. "Out. prob." denotes the probability of any outpatient visit with the indicated primary diagnosis. Values are read from graphical output and are approximate to two decimal places. All three disease categories show monotonically increasing inpatient intensity approaching exit, confirming end-of-life disease escalation across the leading causes of elderly mortality in Japan (Ministry of Health, Labour and Welfare, 2022).

All three disease categories exhibit accelerating inpatient intensity in the approach to exit, confirming that the aggregate pattern in Table E.2 reflects genuine end-of-life disease progression across multiple major diagnostic categories. Cancer inpatient days rise from 0.175 eleven months before attrition to 1.430 in the final month—an 8.2-fold increase. Respiratory disease days rise most steeply in absolute terms in the final three months (from 0.680 to 1.350 days), consistent with acute respiratory failure as a proximate cause of death. Circulatory disease inpatient days rise from 0.650 to 2.140 days, with the steepest acceleration in the final four months.

The outpatient trajectories differ by disease. Cancer and respiratory outpatient probabilities rise monotonically toward exit, reflecting increasing specialist contact for disease management prior to terminal admission. Circulatory disease outpatient probability is more variable, exhibiting a slight U-shape before rising sharply in the final two months, consistent with the heterogeneity of circulatory conditions (ranging from acute cardiac events to chronic heart failure management).

These disease-specific patterns are what one would expect if sample attrition reflects mortality from the leading causes of elderly death in Japan, providing direct diagnostic confirmation of the mortality interpretation.

F Hospital Admission Probabilities by Selected Diagnosis Group

To assess whether the aggregate inpatient decline documented in Section 5.3 is concentrated in diagnoses with greater scope for scheduling or admission-margin adjustments, we classify selected principal inpatient diagnoses into three groups: non-deferrable acute-sensitive diagnoses (ischemic heart disease, subarachnoid hemorrhage, intracerebral hemorrhage, and cerebral infarction), urgent but margin-sensitive diagnoses (pneumonia and fracture), and elective or deferrable proxies (cataract and arthrosis).[21]

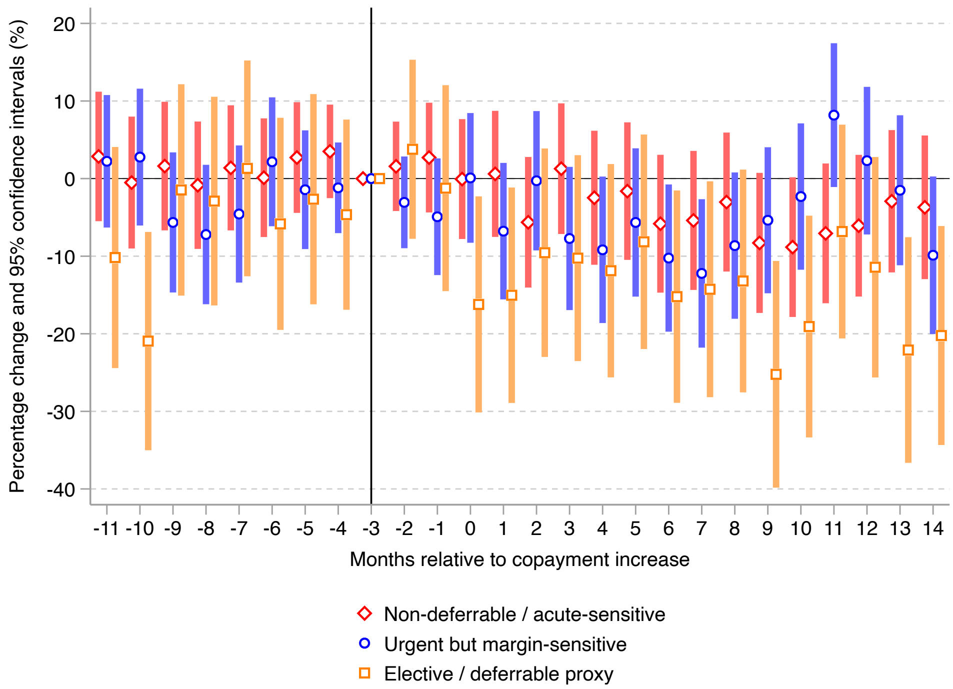

Figure F.1: Effects on Hospital Admissions by Selected Diagnosis Group

Notes: This figure presents event-study estimates for inpatient admission probabilities by selected principal diagnosis groups. The outcomes are binary indicators for whether an individual had an inpatient admission in the corresponding diagnosis group in a given month. Estimates are expressed as percentages of the category-specific pre-treatment mean admission probability. For each diagnosis group, event-study coefficients and confidence intervals are divided by the corresponding pre-treatment mean and multiplied by 100. The non-deferrable acute-sensitive group includes ischemic heart disease, subarachnoid hemorrhage, intracerebral hemorrhage, and cerebral infarction. The urgent but margin-sensitive group includes pneumonia and fracture. The elective or deferrable proxy includes cataract and arthrosis. All models use the same difference-in-differences specification as the main analysis, with individual, month, and municipality-by-month fixed effects. Standard errors are clustered at the individual level.

Figure F.1 presents event-study estimates for monthly inpatient admission probabilities in each

[21] These groups are not intended to provide an exhaustive decomposition of all inpatient admissions; rather, they capture selected diagnosis groups with qualitatively different degrees of plausible scope for scheduling or admission-margin adjustment. The non-deferrable acute-sensitive group comprises conditions requiring immediate intervention where delay would directly compromise outcomes (ischemic heart disease, stroke subtypes). The urgent but margin-sensitive group comprises conditions that are clinically urgent but where the admission decision may involve some discretion at the margin (pneumonia, fracture). The elective or deferrable proxy comprises conditions for which scheduled procedures are standard practice and timing is largely at the patient's and provider's discretion (cataract, arthrosis). These classifications are consistent with standard clinical categorizations used in the health economics literature.

Table F.1: Period-Average Effects on Inpatient Admissions by Selected Diagnosis Group

	Non-deferrable / acute-sensitive		Urgent but margin-sensitive		Elective / deferrable proxy	
	coefficients	se	coefficients	se	coefficients	se
Pre-period average, k=-11 to -1	1.5	(2.8)	-2.1	(3.0)	-4.5	(5.1)
Early pre-period, k=-11 to -4	1.4	(3.1)	-1.6	(3.3)	-5.9	(5.3)
Immediate pre-period, k=-2 to -1	2.1	(2.9)	-4.0	(3.1)	1.3	(5.7)
Post-period average, k=0 to 14	-3.9	(3.4)	-4.6	(3.5)	-14.6	(5.3)
Immediate post-period, k=0 to 2	-1.7	(3.6)	-2.3	(3.9)	-13.6	(6.1)
Middle post-period, k=3 to 8	-2.8	(3.8)	-8.9	(3.9)	-12.2	(5.7)
Late post-period, k=9 to 14	-6.1	(3.9)	-1.4	(3.9)	-17.5	(5.8)
Pre-treatment mean	0.0038		0.0039		0.0016	

Notes: This table summarizes period-average effects for selected inpatient diagnosis groups. Estimates are expressed as percentages of the category-specific pre-treatment mean admission probability. Each estimate is calculated as the average event-study coefficient over the indicated period, divided by the corresponding pre-treatment mean and multiplied by 100. Standard errors, reported in parentheses, are computed using the variance-covariance matrix of the event-study estimates and therefore account for covariance across monthly coefficients. The omitted reference month is $k = -3$. The non-deferrable acute-sensitive group includes ischemic heart disease, subarachnoid hemorrhage, intracerebral hemorrhage, and cerebral infarction. The urgent but margin-sensitive group includes pneumonia and fracture. The elective or deferrable proxy includes cataract and arthrosis.

group, scaled by the category-specific pre-treatment mean to make magnitudes comparable across groups with different baseline admission rates. Because diagnosis-specific admissions are rare events, monthly estimates carry wide confidence intervals and should be interpreted with appropriate caution. Table F.1 summarizes these estimates as period averages.[22]

The pre-period estimates are imprecise across all three groups, reflecting the rarity of these diagnosis-specific admissions. In the post-implementation period, however, a clear gradient emerges across groups. Non-deferrable acute-sensitive diagnoses show no systematic decline, with a post-period average of -3.9% that is statistically indistinguishable from zero. Urgent but margin-sensitive diagnoses show a similarly modest and noisy post-period average of -4.6%. By contrast, the elective or deferrable proxy exhibits a substantially larger and more persistent decline, averaging approximately -14 to -15% below the category-specific pre-treatment mean throughout the post-period, with the reduction deepening over time. While the wide confidence intervals for the non-deferrable group preclude strong claims about the precise magnitude of any effect on acute admissions, the contrast between groups is consistent with the interpretation advanced in the main text: the aggregate inpatient decline primarily reflects adjustments in admissions with greater scope for scheduling decisions, rather than a broad reduction in clearly unavoidable acute care.

[22] For diagnosis group d and period P, the period-average effect is computed as $100 \times |P|^{-1} \sum_{k \in P} \hat{\beta}_{dk} / \bar{Y}^{d}_{\text{pre}}$, where $\hat{\beta}_{dk}$ is the event-study coefficient for group d in relative month k, and $\bar{Y}^{d}_{\text{pre}}$ is the category-specific pre-treatment mean. Standard errors are computed using the variance-covariance matrix of the event-study estimates, accounting for covariance across monthly coefficients. The omitted reference month is $k = -3$.