Market Power and the Supply of Recyclable Waste: Evidence from Indonesia's Waste Banks*

Mychaela Paetow

University of Southern California

Click here for the latest version

October 29, 2026

Abstract

Informal recycling markets are ubiquitous in developing countries. Low state capacity for waste management creates opportunities for intermediaries to buy and sell sorted recyclable waste for profit. This paper measures the extent of market power that intermediaries have over their sellers, and what this implies for the recovery of recyclable material. Leveraging 757,559 transactions from a national network of Indonesian waste banks, I estimate the elasticity of supply for plastic facing this large buyer of recyclable waste, where I employ a shift-share IV strategy that interacts the price of virgin plastic, a substitute good for plastic waste, with each bank's predetermined plastic revenue share. I find vastly different elasticities in the two segmented markets for suppliers. Waste pickers, who collect for income, have a supply elasticity of 1.97, while households, which deliver only what their own consumption generates, do not respond to the price at all. Given this segmentation, I explore how banks respond when competition for these sellers changes, using the staggered opening of rival banks. When a rival bank opens, incumbents lose about a third of their plastic, all of it from pickers. They respond by adding finer, higher-priced grades of plastic that pickers disproportionately collect, while cutting prices on the coarse, mixed grades into which most household material falls. Competition thus raises pay for sellers who can leave and widens the markdown on those who cannot. These results matter for the design of recycling subsidies. I show that a subsidy targeted to waste pickers would cost about 30 percent less per ton of induced collection than a uniform one, and that a bank that segments its sellers through its menu delivers most of that saving on its own.

Keywords: recycling; informal markets; monopsony; pass-through; waste pickers; Indonesia

JEL Codes: Q53, O13, L13, O17, Q58

*Paetow: University of Southern California, 3620 S. Vermont Ave., Kaprielian Hall 363, Los Angeles, CA 90089 (email: mpaetow@usc.edu). I am deeply grateful to Paulina Oliva for her guidance and support throughout this project. I also thank Vittorio Bassi, Matthew Kahn, Jeffrey Weaver, and Simon Quach for their advice, Matt Gordon and Anna Papp for useful comments, and participants in the USC Applied Microeconomics Reading Group and the USC Development Seminar. I am very thankful to Putra Fajar Alam and Brandon Tjiandra of SMASH.id and Ferry Husen of Bank Sampah Bersinar, without whom this project would not have been possible. I gratefully acknowledge proposal development funding from the King Climate Action Initiative (K-CAI) at J-PAL. All errors are my own.

1 Introduction

In much of the developing world, recyclable material is recovered not by the state but by a private supply chain.[1] Within this chain, households and waste pickers sell sorted material by weight to small intermediaries, who then aggregate it and resell it to traders and industrial recyclers. What this chain does not process is typically burned, dumped, or washed into waterways, which comes with significant costs to health and to the environment (Kaza et al., 2018; Cook et al., 2026). In Indonesia, informal recycling recovers about 15 percent of all waste generated, and the vast majority of the country's recycling occurs through these private channels (World Bank, 2018; Damanhuri and Padmi, 2012; Zahrah, Yu and Liu, 2024). Despite their scale, little is known about how these markets operate, and in particular how the prices that intermediaries post affect the amount of material recovered.

This paper asks whether these intermediaries exercise market power over their sellers, and what that implies for the recovery of recyclable material. A growing literature documents that employers face upward-sloping labor supply curves and pay workers less than their marginal product (Manning, 2003; Card, 2022), and that the traders who buy from farmers in developing countries do the same (Bergquist and Dinerstein, 2020; Rubens, 2023). Recent work extends this to settings with environmental externalities, studying buyer power in agricultural supply chains where expanded production causes deforestation (Domínguez-Iino, 2026; Barrozo, 2024). In these settings, a markdown that reduces upstream supply also reduces the externality. In the market for recyclable waste, the relationship runs in the opposite direction. Each kilogram a seller delivers has a private resale value and a public value, and a markdown that reduces supply reduces both.

A standard monopsonist marks down price according to the supply elasticity it faces, and one facing sellers with observably different elasticities of supply typically sets a separate markdown for each. The intermediaries I study cannot do this. Instead, they post a single price per grade that applies to every seller, despite facing two very distinct groups of sellers that are likely to have very different supply elasticities. Their sellers are of two broad types: households, who deliver the recyclables their own consumption generates, and waste pickers, who collect material for income. I use this setting to understand how a monopsonist prices when it faces a segmented population of sellers but cannot price discriminate across segments directly.

To do so, I draw on the transaction records of a national network of Indonesian *waste*

[1] Roughly 1 percent of the urban population in developing countries, at least 15 million people, earn a living by recovering recyclables (Medina, 2008), and the informal sector collects an estimated 58 percent of all plastic waste collected for recycling worldwide (Lau et al., 2020).

banks (*bank sampah*). Indonesia is a setting in which the stakes of recovery are unusually high: it generated roughly 34 million tons of recorded waste in 2024, about 40 percent of which was mismanaged, and it ranks among the largest national sources of plastic entering the ocean (Kementerian Lingkungan Hidup, 2024; Jambeck et al., 2015; Meijer et al., 2021). It is also one in which the private chain can be observed. The data cover 757,559 deposits by 38,597 sellers at 553 banks between 2013 and 2026. Each record identifies the seller, the material and grade deposited, the weight, and the price paid. This allows me to estimate the supply facing an individual buyer separately for each type of seller. I combine these records with the network's roster of verified bank opening dates, which lets me date the arrival of a rival near each incumbent. The paper proceeds in four steps.

I first document how waste banks set prices. A bank buys plastic by grade. A grade is a line on the bank's menu, defined by a resin, such as PET or polypropylene, and by quality attributes within that resin, such as color and cleanliness. For example, clear PET bottles and mixed PET are two grades of one resin. A bank posts one price per grade and credits every seller at that price, regardless of who she is or how much she delivers. The bank can change the price a seller effectively receives in two ways: (1) by repricing a grade it already carries, or (2) by adding a cleaner or more finely specified grade, which pays a premium over the coarser grade it is split from. Both margins are active. Over the course of a year, 86 percent of banks change their plastic menu, adding, dropping, or repricing a grade, and 53 percent reprice at least one grade they already carry. Month to month, a grade enters the menu more often than an existing one is repriced. The two margins differ in which sellers they reach. Repricing an existing grade changes the price paid on every kilogram of that grade, whereas a new grade pays a premium only to sellers who sort material to it. Picker sellers deliver 2.8 distinct plastic grades per visit against 1.5 for household sellers, a difference that persists when I hold the bank's menu fixed. Adding grades therefore allows a bank to raise the price paid to pickers without raising it for households, a form of second-degree price discrimination (Mussa and Rosen, 1978).

Having established how banks set prices, I estimate the elasticity of plastic supply to the price a bank posts, separately for the two types of sellers. Since a bank can pay its sellers more either by repricing a grade or by adding one, I measure the price a seller faces with a quality-adjusted price index that values both margins. My empirical strategy relies on the relationship between virgin and recycled plastic. Recycled plastic substitutes for virgin plastic, so a rise in the world price of virgin plastic also raises the value of the plastic waste a waste bank collects. Because this shock is common to all banks, identification comes from cross-sectional variation in exposure to it. I instrument a bank's quality-adjusted plastic price with its predetermined plastic revenue share interacted with the world price of virgin

plastic, so the comparison is between banks with high and low plastic exposure that face the same upstream movement in the same month. I estimate an aggregate supply elasticity of 1.18 (s.e. 0.74), with standard errors in parentheses hereafter. This average conceals a large difference between seller types. The supply elasticity of picker sellers is 1.97 (0.78). A higher price attracts more pickers and more volume per picker, with extensive- and intensive-margin elasticities of 1.09 (0.41) and 0.89 (0.65), respectively. The supply elasticity of household sellers, by contrast, is 0.64 (0.59) and statistically indistinguishable from zero. Notably, the way banks reprice reveals awareness of this difference. I find that the shock is not passed uniformly through grades. Within a bank and month, the posted price of a continuing grade rises by less the larger the share of its kilograms that households deliver, -0.73 (0.37) per unit of share. The markdown a bank sets falls hardest on the households whose supply does not respond to it.

I then study how incumbent banks respond when a rival opens nearby, which changes sellers' outside options while leaving downstream values unchanged. I date openings from the network's verified roster and estimate stacked event studies against banks not yet or never treated. Six to eleven months after an opening, an incumbent's plastic intake is lower by 0.45 (0.21) log points, about 37 percent, and the loss is entirely from pickers. A fifth of picker sellers leave, while plastic delivered by households and the number of household sellers do not change. Standard theory predicts that competition for sellers should raise the price an incumbent pays. Instead, incumbents raise it for some sellers and lower it for others. They lower the posted price of the grades they already carry by a few percent, and they split the resins they buy into finer grades about twice as often as before the opening, posting the new grades at roughly twice the price of the coarse grade they refine. The new grades are pickers' grades. Picker sellers deliver 88 percent of the material on them and move half of their plastic onto them within two years, while most household material stays on the existing grades, whose prices fell. A rival draws off the elastic sellers, and the incumbent bids for them on the instrument that reaches their sorted loads, while widening the markdown on the grades whose remaining sellers will not leave over a price cut. Competition therefore raises the effective price paid to sellers whose supply responds and lowers it for those whose supply does not.

Finally, I consider implications for environmental policy. The environmental value of a kilogram delivered to a bank depends on where it would otherwise have gone, and this differs by seller type. Household deliveries do not respond to the price, and survey evidence identifies civic and environmental motives, rather than the payment, as the principal reason Indonesian households use a waste bank (Widayat et al., 2025). Their material is therefore most plausibly diverted from municipal collection. Pickers' material responds to price, but

part of the response may be redirected from other buyers rather than newly recovered. Newly recovered materials can come from sources that differ widely in the harm they would otherwise cause.[2] I cannot measure these alternative fates directly, but I can bound the value of each. Drawing on estimates of disposal costs and of the health damages of open burning, I value a ton that would otherwise have reached a sanitary landfill at US$20–25 and a ton that would otherwise have been burned in the open at US$300–1,500. For a ton that would otherwise have leaked into the sea, the causal estimates of Du, Zhang and Zou (2024) on marine microplastics and birth weight give a benchmark of US$5–90 per ton on the infant-health channel alone and of the same order as open burning once the associated mortality is valued at the value of a statistical life (VSL). A uniform per-kilogram subsidy paid through a bank that posts one price to all of its sellers costs about US$115 per ton of induced collection. Because pickers deliver about 80 percent of the kilograms, most of this payment reaches the sellers whose supply responds, and its inframarginal component is a transfer to some of the poorest workers in the economy. The part paid on household deliveries buys no additional recovery. A payment restricted to pickers costs US$82 per ton. However, a bank that separates its sellers through its menu achieves most of this saving without any targeting by the government. Such a bank passes a uniform subsidy through at a higher rate on the grades pickers deliver than on the grades households deliver, so the same payment costs US$96 per ton. The market power that allows a bank to mark down its inelastic sellers also directs the subsidy toward its elastic ones. Whether any of these payments is worth its cost depends on where the marginal kilogram would otherwise have gone. Under the most conservative estimates of damages from open burning, the targeted payment pays for itself even if nearly three quarters of the picker response is redirected from other buyers rather than newly recovered.

Related literature. The paper contributes to four literatures. The first studies monopsony, whose central object is the elasticity of labor supply facing an individual firm (Manning, 2003; Card, 2022). Estimating it requires variation in one firm's wage that other firms do not share, so much of this literature infers it indirectly (Sokolova and Sorensen, 2021). A growing body of work shows that the elasticity differs across groups of workers within the same firm (Card, Cardoso and Kline, 2016; Sharma, 2023; Caldwell and Oehlsen, 2026), and that constraints on paying workers differently change how firms set wages (Cullen and Pakzad-Hurson, 2023). I observe every price a buyer posts and every transaction at it, so I

[2]Pickers buy sorted material from households door to door, sort the mixed stream on streets and at transfer stations, and pick the face of open disposal sites. Material recovered at a disposal site was already bound for it, whereas material recovered from streets, drains, and neighborhoods without collection would otherwise be more likely to be burned or to leak into waterways.

estimate the elasticity facing a single buyer directly, separately for two types of sellers who face an identical schedule. I also document a margin of adjustment this literature has not, to my knowledge, studied. When a rival bank opens, the incumbent responds not by raising its price but by refining the menu of grades it buys, creating premium grades that only the elastic sellers claim.

The second is the literature on intermediaries in developing-country supply chains. Traders at the first tier of these chains pay producers less than competitive prices (Bergquist and Dinerstein, 2020; Rubens, 2023; Zavala, 2022), entry among them raises producer prices (Bergquist and Dinerstein, 2020; Casaburi and Reed, 2022; Chatterjee, 2023), and world prices pass through to producers incompletely (Atkin and Donaldson, 2015; Mitra et al., 2018). Recent work links this market power to environmental outcomes in agricultural supply chains where expanded production causes deforestation (Domínguez-Iino, 2026; Barrozo, 2024). Kochhar (2024) shows that it also hinders farmers' adaptation to climate shocks. I show that an intermediary that posts a single price to sellers with different supply elasticities marks down the less elastic sellers the most, and that the entry of a rival widens their markdown.

The third is the literature on uniform pricing and on how prices adjust through changes in the set of products offered. DellaVigna and Gentzkow (2019) show that U.S. retail chains charge nearly uniform prices across stores whose demand elasticities differ substantially, and large employers similarly post one wage for a job across local labor markets (Hazell et al., 2022; Quach, 2025). In price measurement, indices that follow only continuing products miss price changes that occur at product replacement (Nakamura and Steinsson, 2012; Gagnon, Mandel and Vigfusson, 2014), and the entry of new goods complicates the measurement of quality-adjusted prices (Bils, 2009; Broda and Weinstein, 2010). I document the same forces on the buyer side of a market. I show that a buyer facing sellers who differ in their responsiveness to price can discriminate between them by refining the menu of grades it buys rather than the price of any one grade.

The fourth is the economics of waste and recycling. The classical literature shows that when disposal cannot be observed, a deposit-refund system dominates a disposal tax because it rewards the observable act of returning material (Fullerton and Kinnaman, 1995, 1996; Palmer and Walls, 1997; Fullerton and Wolverton, 2000; Kinnaman and Fullerton, 2000). Where municipal collection makes disposal observable, a unit-based tax can reduce waste (Lee, 2026). In my setting the tax is uncollectible where most households burn or dump, but a refund is already paid privately as the price a neighborhood intermediary posts. The policy question is therefore whether, and how, to subsidize a market that already operates

(Berkouwer and Dean, 2022), which requires knowing how much additional material a higher price elicits, and from whom. Existing estimates come from national or industry time series for recovered paper and scrap metal in the United States and Europe and find the response low, with own-price elasticities for recovered paper between 0.2 and 0.3 (Anderson and Spiegelman, 1977; Edwards and Pearce, 1978; Mansikkasalo, Lundmark and Söderholm, 2014; Blomberg and Söderholm, 2009; Zink, Geyer and Startz, 2018). My estimate is, to my knowledge, the first supply elasticity for recovered material identified from the transaction records of individual buyers, and the first for the informal tier of the recycling chain.

The remainder of the paper proceeds as follows. Section 2 describes the Indonesian waste system and Section 3 the data. Section 4 presents the main results, and Appendix A sets out a simple model of a waste bank that organizes their interpretation. Section 5 draws out the policy implications, and Section 6 concludes.

2 Setting

This section describes the features of the Indonesian waste system that matter for my analysis, namely the limited reach of public collection, the private recovery chain that operates in its place, and the position of waste banks and of plastic within that chain.

2.1 Public Collection and Its Limits

Responsibility for waste collection and disposal rests with Indonesia's 514 district and municipal governments, which decide independently how much collection to provide and how to pay for it.[3] Financing is nominally local, through a user charge (*retribusi*) and the general local budget, but neither raises much. In the most detailed cost exercise available, cost recovery from the *retribusi* ranged from 5 to 48 percent of service cost across five jurisdictions, with a median near 17 percent. Most regions also allocate well under one percent of their budgets to waste management (Vidyaningrum, 2020; Tempo, 2025*a*). As a result, formal collection reaches only a minority of the population. The 2021 village census records that 70.5 percent of villages and urban wards dispose of most household waste by burning or burying it, against 19.4 percent whose waste is taken to a collection point (Badan Pusat Statistik, 2022). Even in areas where waste is collected and transported to a final disposal location, the externalities still remain high. Of the 550 disposal sites counted in 2025, 343

[3]Republik Indonesia, 2008; Republik Indonesia, 2012. The central government sets policy and standards; provinces coordinate and supervise; districts operate.

operated as open dumps and about 110 met sanitary standards (Tempo, 2025*b*).[4] Public collection, where it exists, is a disposal service rather than a recovery one. Waste is collected mixed, with little to no separation at the source or at transfer points, and what the trucks carry goes to the disposal site whole. Recyclable material that enters the public stream is recovered, if at all, by the pickers who work the transfer stations and the face of the dump.

2.2 Waste Picking and the Private Recovery Chain

Given the absence of public capacity for recycling,[5] recovery is primarily performed by private actors, of which pickers make up the majority. Waste pickers (*pemulung*) most commonly operate at three points in the discard stream: buying sorted material door to door as itinerant buyers (*rosokan*), working the streets and transfer stations, and sorting at the face of the disposal sites themselves. Entry barriers to picking are low. It requires little capital, no license or fixed workplace, and no schooling or formal training, which is why it often functions as an absorbing occupation for the very poor. Pickers sell by weight to a buyer who posts a per-kilogram price by material, and earnings are the product of that price and what they can physically collect, store, and carry, with few monetary input costs. The work is highly flexible and closely resembles a gig-type job. Pickers choose where and when to search and how to allocate effort and carrying capacity across materials.

Material recovered by pickers moves up a chain of aggregators, each of which buys at a per-kilogram price, sorts and consolidates, and resells to the tier above. Besides waste pickers (*pemulung*), the neighborhood junkshop (*pengepul*, *lapak*), the regional trader (*bandar*), and the industrial recycler (Iacovidou et al., 2025*b*,*a*; World Bank, 2018) are the most commonly identified actors within this chain. The number of tiers and the identity of the buyer at each vary by area: in some markets a picker sells directly to a regional trader, in others a junkshop sells to a larger junkshop before the material reaches a trader, and a bank's onward buyer may be any of these. The waste banks I study sit at the first tier of this chain, beside the neighborhood junkshop. Both buy sorted material from pickers and households at a per-kilogram price, and both sell it onward to the tiers above.

[4] A sanitary landfill places waste in lined cells, compacts it and covers it with soil daily, and collects and treats leachate and landfill gas; an open dump does none of these.

[5] The state's own recycling infrastructure is small. A recent inventory counts 923 material recovery facilities nationally, community-scale TPS 3R sorting sites and larger TPST processing facilities, of which only 508 are active (Iacovidou et al., 2025*b*). Facilities are built to national program targets with central capital grants and handed to village associations without an operating budget, and many cease operating within a few years of handover (Vidyaningrum, 2020). Formal systems capture less than 5 percent of waste against about 15 percent recovered informally, and active household source separation is 1.6 percent, less than 0.5 percent of it involving plastics (World Bank, 2018, pp. 4 and 20).

2.3 Waste Banks

A waste bank (*bank sampah*) is a first-tier aggregator that functions much like a junkshop or *lapak*. It differs from a junkshop in three respects. First, waste banks are typically established with initial implementation support from local government or a community organization, which a junkshop does not receive. Second, a waste bank posts a per-kilogram price for each grade on a menu and does not condition it on the identity of the depositor or the size of the deposit, whereas a *lapak* need not display a menu and is free to vary the price it offers with the load or the seller. Third, material is weighed, categorized, and credited at the posted price to a savings balance that can be withdrawn periodically rather than paid as cash on the spot, which is what gives the waste bank its name. Some waste banks are also run by neighborhood associations, religious organizations, or schools, though these are the minority and are largely not the type I study.

The sector has grown rapidly since the first banks were registered in the late 2000s. The most recent count available puts it at 4,341 operating banks at the end of 2022 (Budiyarto, Clarke and Ross, 2025). The PODES 2024 village census attempted a nationwide count of waste banks for the first time, though this data is not yet publicly available. Based on the current available data, waste banks handle only a small fraction of the national waste stream and of national recycling, on the order of one and two percent (Iacovidou et al., 2025*b*). These shares matter for interpreting the results. The sector is not large enough for its direct diversion to move national aggregates. I study it because it is the only part of the private recovery chain for which prices and quantities are recorded. Because a bank closely resembles the junkshop beside it, its behavior is also plausibly informative about how the *pengepul* and other private buyers in the chain operate.

The banks I observe are a sizable share of the overall population of waste banks. The 553 banks that record transactions on the partner platform are roughly one in eight of the 4,341 banks counted as operating at the end of 2022, and the 875 banks on its roster roughly one in five; they span 30 of Indonesia's 38 provinces, though they are concentrated in West Java, where the partner is based and where 40 percent of the banks and half of the kilograms are located (Section 3.6). There is no independent census of bank size or clientele against which to compare them, and banks that record on a digital platform are plausibly better organized than the median bank. The network is nonetheless national in reach, includes banks of every scale from a few depositors to hubs that buy from other banks, and the menu structure I document is not specific to any one region or operator.

2.4 Plastic

I focus on plastic for four reasons. First, it is the second-largest component of the waste stream after food waste, at about a fifth of generated waste and rising, and together with paper, metal, and glass it makes up the fraction of the stream with positive re-sale value. Second, it is the material whose mismanagement is most costly. Unlike many other materials, plastic does not degrade – what is dumped fragments into microplastics that leach into soils, rivers, and the sea, with measurable consequences for human health (Du, Zhang and Zou, 2024; Beaumont et al., 2019). Burning is worse still. Open burning of plastic releases fine particulate matter, black carbon, and dioxins and furans, and household waste burning is among the largest sources of these pollutants in Indonesia (Wiedinmyer, Yokelson and Gullett, 2014; Wang et al., 2023; Pathak et al., 2023). Third, among the materials bought and sold in these private markets, plastic has one of the highest re-sale values: the median posted price for PET is Rp 2,000 per kilogram, against Rp 1,000 for paper and Rp 200 for glass (Table 1). Fourth, it is the one recovered material with a direct link to a globally traded commodity. Recycled plastic competes with virgin plastic, a petroleum derivative, as a production input, so its value at an Indonesian collection point is set several links up a chain that terminates in the world market for virgin plastic. The other materials waste banks buy have no comparable link, which allows them to serve as comparison materials in Section 4.

3 Data

For my analysis, I combine four sources of data: waste bank transaction records, the network's roster of bank openings, upstream price data for virgin plastic, and a two-year ledger of one bank's onward sales. This section describes each briefly and defines the variables used throughout.

3.1 Transaction Records

My core data are the transaction records of SMASH.id, a software platform on which waste banks across Indonesia record their deposits. I obtained them through a research partnership with the platform's operator, which works with Indonesia's Ministry of Environment. The application logs every transaction between a depositor and a bank at the moment it occurs. It records the ID of the depositor and the receiving bank, the depositor's type, a timestamp, a material category as labeled by the bank, the weight in kilograms, and the credit posted

to the depositor's account. Every bank carries a geocoded location. The price a depositor receives is the credit per kilogram, and I define the modal credit per kilogram within a bank–line–month as the posted price for that line and month,

$$p_{bjt} = \text{mode}_{i \in (b,j,t)} \left\{ \frac{\text{credit}_i}{\text{kg}_i} \right\}. \tag{1}$$

In practice, 98 percent of transactions in a line–month are exactly at that mode, so the mode recovers the posted schedule almost exactly, and I treat it as the price every depositor at that bank faces for that grade in that month.

Over the course of a year, 86 percent of banks change their plastic menu by adding, dropping, or repricing a grade, and 53 percent reprice at least one grade they already carry, so month to month, a grade enters the menu more often than an existing one is repriced (Figure 2). After removing records with unparseable dates or non-positive weight or credit,[6] the sample contains 757,559 transactions between 553 waste banks and 38,597 depositors, spanning 30 of Indonesia's 38 provinces, February 2013 through August 2026.[7] About 200 banks begin recording in 2017 and 2018, the network then grows slowly until 2024, and 212 banks join between 2024 and July 2026 (Figure 1, Panel A). Descriptive statistics on menus and prices begin in January 2018, when enough banks record for cross-bank comparisons to be meaningful. The supply design measures each bank's exposure over 2018–2021 and estimates on January 2022 through July 2026, so that exposure is predetermined relative to the price variation it is interacted with (Section 4.1.1). The entry design uses every dated opening at which the incumbent qualifies under the rule of Section 4.2.

The material category is a label that each bank defines for itself, so the raw records contain thousands of distinct labels. I organize them at three levels, and I use the same three terms throughout the paper. A *material* is a broad class of waste, such as plastic, paper, metal, or glass. Within plastic, a *resin* is a polymer type, such as PET, HDPE, or polypropylene. Within a resin, a *grade* is a line on a bank's menu, distinguished from the other lines of the same resin by quality attributes such as cleanliness, color, and form. Clean PET, clear PET bottles, and mixed PET are three grades of one resin, and a bank that carries all three posts three prices. Because every grade is a line on some bank's menu, I use

[6] I drop transactions whose timestamp fails to parse or falls before 2013 or after the extract date, most of which parse to a default date around 1970, and transactions with non-positive recorded weight or credit.

[7] Some banks are hubs (*bank sampah induk*) that buy from smaller unit banks as well as from individual depositors; the roster identifies 27 hubs and 34 unit banks among 875 banks. A unit bank's delivery to its hub appears in the records as a deposit by an account belonging to another bank. These accounts, and those of institutions such as schools and offices, are excluded from every quantity outcome, but retained in the price index, since a bank's posted schedule faces every account. Unit banks are also excluded as entrants in the entry design. Appendix Table A9 adds these accounts back.

"grade" and "line" interchangeably, and a grade is always specific to the bank that posts it. For descriptive statistics and for the comparison materials of Section 4, I collapse the labels into eight groups: PET, other plastic resins, paper, metal, glass, used cooking oil, undifferentiated waste, and a small remainder.[8] For the analysis of plastic prices, I do not collapse. The unit is the grade, so a bank that distinguishes clean from mixed PET is treated as posting two prices rather than one, and a line called "PET" at one bank need not have the same quality as a line with that name elsewhere. Between January 2018 and July 2026, the plastic data contain 3,677 grades, 25,924 grade–months, and 395 banks, and a bank posts a median of five plastic grades in a month (Table 2). The grades of a resin differ in quality rather than in kind. Within PET, a clean or clear grade posts at roughly twice the price of a colored or mixed one (Figure 3, Panel C), so a bank that splits a coarse grade raises what a depositor who sorts to the finer grade receives without changing any number already on the board. Paper, metal, and glass menus are instead short lists of distinct products, such as cardboard and newsprint or tin and copper, with a median of two or three lines per bank. The attribute model below therefore prices plastic grades against close substitutes on the same menu, and I use the other materials only as placebos (Appendix Table A5).

3.2 The Quality-Adjusted Price

My headline price is a quality-adjusted, chained hedonic-imputation index. I parse each bank-specific line name into 148 attributes,[9] covering resin, cleanliness, color, form, platform codes, and their principal interactions, and estimate

$$\log p_{bjt} = \alpha_{bt} + x_j'\gamma + e_{bjt}, \tag{2}$$

weighting by transaction credit and absorbing bank–month effects, on every plastic line–month in the records. The attributes explain 81 percent of the variation in log prices across the plastic lines a bank posts in a month, pooling every resin and quality on the menu, and the attribute prices are reported in Appendix Table A1. Within PET and polypropylene, dirty, mixed, colored, and black grades are discounted, and clean and clear grades carry premia. Estimated on PET lines alone, the attributes explain 86 percent of the variation in log prices across a bank's PET lines in a month. The fit is poor for the residual "other plastic" group and for expanded polystyrene, whose labels carry little attribute information,

[8] Nearly all of the undifferentiated-waste ("Residu") volume is a bulk municipal hand-off that one bank logged at a nominal Rp 1 per kilogram over 2023–24. It appears as its own row in Table 1, is excluded from the per-bank volume figures, and enters no regression.

[9] The records contain 1,988 distinct plastic labels built from 98 recurring descriptors. Paper and metal each have roughly 600 labels built from 35 to 38 descriptors and glass has 258 labels built from 15.

and I place no composition claim on either.

The chained index P_{bt}^{H} is a credit-weighted average of month-to-month log price changes across a bank's plastic lines.[10] Constructing it requires a rule for which lines are on the menu. I observe a grade only in the months in which someone delivers it, so a grade absent from the records in a given month may have been dropped from the bank's schedule or may simply not have been supplied. Deposits are intermittent, and the two cases are indistinguishable in any single month. I separate them by the length of the silence. A line that transacts again after a silence of six months or less is treated as a continuing line, and its contribution to the index is the change in its posted price since it last transacted. A line transacting for the first time, or returning after a silence of more than six months, is treated as new to the menu, and its contribution is the gap between its price and what the bank's continuing grades would imply for a grade of its attributes under equation (2). The rule is motivated by the records themselves. Scoring every grade absent in the previous recorded month as new, six in ten entering lines are grades the bank had already transacted on, and eight in ten of those return at exactly their previous price, so the rule keeps those returns as repricing events, of zero size when the price is unchanged, and reserves the new-grade term for grades the bank is writing onto the menu.[11] The results are not sensitive to where the line is drawn. Shortening the rule to three months, lengthening it to nine, or scoring every return as new does not materially impact the main results (Appendix Tables A7 and A21). The index separates exactly into a continuing-grade component and a new-grade component, and the new-grade component contains both grades transacting for the first time and grades returning after a silence.

I use this index as my headline price because banks change what they pay both by repricing the grades they already carry and by revising the menu, and the index is the closest measure I have of the realized change in the price that households and pickers face. It measures the bank's schedule, the price it pays for a grade of given attributes relative to its other grades. A number on the board that rises moves it. A line that enters moves it only if it is posted above what the bank's existing grades would imply for its attributes. A bank that posts coarse PET at Rp 2,000 and adds a finer grade whose attributes command a premium of 0.6 log points, the median for the refinements of Table 12, leaves the index unchanged if

[10] The weights are the average of a line's share of the bank's plastic credit in months $t-1$ and t, the Törnqvist formula standard in the price-index literature (Diewert, 1976). A line that leaves the menu contributes nothing to the change, and a new line's imputed gap is capped at one log point in absolute value.

[11] Appendix Figure A1 shows how much of each bank's cumulative index change arrives through lines new to the menu. On the quality-adjusted index, a schedule revision moves the effective price by a median of 0.015 log points, one revision in six by more than 10 percent, and the change is fully persistent a year later (Figure 2, Panel B).

it posts the new grade at about Rp 3,600, and raises it only above that. The index does not measure what a particular seller takes home. A seller who sorts her material to the finer grade is paid more for the same material, and the index credits the whole difference to the grade rather than to the price, because it is built to hold quality fixed. The alternatives each miss one of the two margins. A unit value would attribute to prices the changes in composition that a menu revision produces, and a matched-model index that followed only continuing grades would miss the price at which the bank writes new lines. A refinement is identified from the bank's own menu, so a grade the bank adds beside one it already carries is treated as a finer grade of that resin even where the parsed attributes cannot say in what respect it is finer. Throughout, $\log P_{bt}^{H}$ is the price in both designs and K_{bt} denotes plastic kilograms bought from individual depositors.

3.3 The Network Roster

SMASH.id also provided the network's roster, which records for each registered bank its coordinates and its verified opening date. It covers 875 banks, of which 753 have both an opening date and coordinates, including 284 that register but never record a transaction. I use the roster to date the entry of rival banks near an incumbent (Section 4.2). The roster contains more banks than the transaction records, and many of the entrants in the entry design never record a transaction. A bank competes for an incumbent's sellers whether or not it records on the platform, so the roster is the right source for dating entry, but it means that I can only fully observe the incumbent's side of the market. I observe what the incumbent loses and how it reprices but not what the entrant posts or whom it attracts.

The roster records openings but not closures, and the transaction data cannot supply them. An empty bank–month means either that the bank has stopped recording or that it has closed, and the two are indistinguishable. Gaps in recording are common. Only 71 banks recorded a transaction in January 2026 against 553 that have ever recorded, and many banks go quiet for months and return. I therefore impute no closures and treat an empty month as non-recording rather than as a zero. Every outcome is measured on months in which the bank records. The supply design uses banks with long and consistent histories, and the entry design holds treated and comparison banks to the same recording rule before and after an opening, so that neither group is selected on recording gaps (Appendix Table A11).

3.4 Upstream Price Series

I measure the virgin plastic price with two monthly series. The first is the unit value of Indonesia's imports of PET polymer in primary form, the closest measure of the replacement cost facing the Indonesian plastics chain and the headline shifter.[12] The second is the U.S. producer-price index for plastics materials and resins, which I use as an independent check.[13] Both are converted to rupiah at the monthly average exchange rate and deflated by the Indonesian consumer price index,[14] so that for the resin index

$$v_t = \log(\text{VirginIndex}_t \times e_t/\text{CPI}_t)\,, \tag{3}$$

and analogously v_t^{PET} for the PET import price. Both series are determined in world markets. The identifying variation in Section 4.1.1 is their interaction with a bank's predetermined plastic exposure, with calendar-month effects absorbing the series themselves, so the design requires only that nothing else co-moving with the upstream price affect plastic-exposed banks differentially. Because virgin plastic is a petroleum derivative, I also collect the monthly Brent crude oil price[15] and use it to ask whether waste-bank prices follow the component of upstream plastic values specific to the resin market or the component shared with oil (Appendix Table A4).

3.5 Onward Sales

For one large operator, I observe a ledger of realized onward sales: the date, buyer, product, kilograms, and price of every shipment from January 2024 to December 2025, 5,489 transactions and 1,526 tons of material. The ledger's product codes are the bank's own purchase grades, so a shipment can be matched to the bank's purchases of the same grade in the same month for every material except PET, which the bank buys as loose bottles and sells as pressed bales and which I match at the resin level. Buyers are anonymized within each annual ledger. Section 4.3 uses the ledger to measure the bank's gross buy–sell spread by material and the pass-through of the price it receives into the price it pays.

[12] HS 3907.61; the value in U.S. dollars divided by the weight in kilograms, summed over all partner countries, from the monthly foreign-trade statistics of Badan Pusat Statistik.

[13] Bureau of Labor Statistics series WPU0662, retrieved from FRED.

[14] FRED series IDNCPIALLMINMEI, from the OECD Main Economic Indicators.

[15] FRED series POILBREUSDM.

3.6 Summary Statistics

Table 1 summarizes the sample. Transaction-level variables are heavily right-skewed. The mean transaction weighs 12.6 kilograms against a median of 2.0, and the mean credit of Rp 20,166 is more than eight times the median of Rp 2,400, which reflects the heterogeneity in depositors that organizes the paper. Paper is the largest material at 37 percent of kilograms and plastic the second at 35 percent, of which PET is 15 points. Of the 553 banks, 40 percent are in Jawa Barat, where the platform's operator is based, and the province accounts for half of the kilograms collected (Figure 5, Appendix Figure A2). Plastic's share of collected kilograms rose from 26 percent in 2018 to 36 percent in the final year of the sample (Figure 6).

The records distinguish two kinds of individual depositors, households and waste pickers.[16] Table 3 describes them. The two groups differ sharply in scale. The median picker visit weighs 41 kilograms against 3 for a household visit, 17 kilograms of it plastic against 1.6. Pickers are 13 percent of depositor–bank records but deliver 64 percent of all individual-depositor kilograms and 84 percent of the plastic (Panel C). The average bank–month with any picker depositor contains 12 pickers delivering 2,063 kilograms of plastic, against 45 household depositors delivering 295 (Table 5). The groups differ in what they bring as well. Plastic is 60 percent of a picker's kilograms and 26 percent of a household's, whose deliveries are mostly paper. Pickers also bring more varied plastic, 2.8 distinct grades per visit against 1.5, and bring two or more grades in 63 percent of visits against 27 percent. These differences survive bank fixed effects, which hold the menu fixed. At the same bank, a picker visit spans 0.95 (0.08) more grades and 0.43 (0.03) more resin groups (Panel B).

Table 4 describes what each group delivers in the sample of the supply design. The groups differ less in which resins they bring than in how they bring them. Pickers deliver sorted PET at scale. Bottles free of caps and labels, sorted by clarity or colour and paid about Rp 5,900 per kilogram, are 71 percent of a picker's plastic, and pickers supply 98 percent of what arrives on those grades. Households deliver unsorted material in small loads. A third of their plastic is PET as it comes, paid about Rp 1,400, and the rest is spread in roughly equal shares across HDPE, polypropylene, film, and generic plastic, at a median of under three kilograms a visit. The contrast is sharpest where the choice is on the board. At banks that post both a sorted and an unsorted PET grade in the same month, households put 23 percent of their PET on the sorted grade and pickers 87 percent. The pattern is

[16] The platform records each depositor as a household, an informal-sector worker, a unit bank, or an institution. I refer to informal-sector workers as pickers. At the largest banks the picker group includes junkshop (*lapak*) and sorter (*pemilah*) accounts, which belong to the same collection chain. Accounts that belong to neither group, the corporate clients of a hub's collection programme, community waste facilities (TPS 3R), cash accounts, and one-off events, are removed from the supply-design sample alongside the unit-bank and institutional accounts. They are concentrated at the two largest banks (Section 4.1.1).

network-wide. Across all banks, sorted PET is a third of a picker's plastic and an eighth of a household's, and where a bank posts both tiers pickers put three quarters of their PET on the sorted grade against a third for households (Panel B).

3.7 Price Dispersion Across Banks

The price a seller receives depends on which buyer she faces. Median posted prices span an order of magnitude across materials and a factor of two and a half across plastic resins, and within PET a clear or clean grade posts at roughly twice the price of a colored or mixed one (Figure 3). Within a resin, prices also differ substantially across banks. For each month with at least three reporting banks I compute each bank's kilogram-weighted PET price and take the standard deviation of its log across banks. This statistic averages 0.51 log points over 103 months from January 2018 to July 2026, with no trend, so a bank one standard deviation above the monthly mean pays roughly two thirds more per kilogram than a bank at the mean for what the platform records as the same resin in the same month. The dispersion is not an artifact of banks buying different grades. Holding the full set of grade attributes fixed with the hedonic model of Section 3.2, the quality-adjusted price level is as dispersed as the unit value itself, with a cross-bank standard deviation of 0.47 log points, and grade composition accounts for only 29 percent of the cross-bank variance in the unit value (Appendix Figure A4). Nor is it mainly regional. Among the 55 banks in Greater Bandung, which sell into the same downstream market and face the same transport costs, the standard deviation of the bank-average quality-adjusted residual is 0.38, against 0.42 across the whole network, and Appendix Figure A5 shows high- and low-paying banks side by side within one metropolitan area. Two buyers a few kilometers apart pay prices for the same grade that differ by more than half, month after month, over eight years.

Prices this dispersed are difficult to reconcile with an integrated market in which sellers arbitrage across buyers. They are consistent with each bank facing its own upward-sloping residual supply curve, and they are informative about who the sellers are. A household that deposits for civic reasons does not shop across banks, so a bank whose depositors are households can pay less without losing them, whereas a picker who sells for income compares buyers. Persistent differences in levels are also what a buyer that pays a fixed fraction of the value of its material produces. Each bank moves one for one with its own downstream value and keeps its place in the distribution, which is the pattern Section 4.3 documents at the one bank where that value is observed. Dispersion alone does not establish market power, since sorting costs and unobserved quality also vary locally. Section 4 provides more direct evidence.

4 Main Results

In this section, I present main results for two separate but related designs. First, I estimate the supply elasticity of plastic waste, in aggregate and for the two types of sellers (Section 4.1). Second, I examine the role of competition. Specifically, I use the opening of a rival bank to study how incumbents respond when competition for their sellers changes (Section 4.2). Section 4.3 concludes with a brief analysis of one bank's onward sales.

4.1 The Supply of Plastic to a Waste Bank

4.1.1 Design

I estimate the supply elasticity facing a single waste bank by instrumenting the quality-adjusted price with the interaction of its predetermined plastic exposure and the upstream value of plastic. Recovered plastic competes with virgin plastic as a production input, so a rise in the upstream value of virgin plastic raises the value of what a bank collects, and it plausibly does so by more at a bank whose business is concentrated in plastic than at one that buys mainly paper or glass. Baseline exposure is the plastic share of a bank's 2018–2021 transaction credit,

$$\varpi_b = \frac{\text{plastic transaction credit}_{b,\text{2018–2021}}}{\text{total transaction credit}_{b,\text{2018–2021}}}, \tag{4}$$

and the instrument is this share interacted with an upstream price shifter, $Z_{bt} = \varpi_b s_t$, where s_t is a monthly series common to all banks. The headline shifter is the real-rupiah unit value of Indonesian PET imports, v_t^{PET}. The real-rupiah virgin-resin producer-price index, v_t, serves as a robustness check (Appendix Figure A3 plots both series and the distribution of exposure).

The object of interest is the supply elasticity ε^S in

$$\log K_{bt} = \alpha_b + \psi_t + \varepsilon^S \log P_{bt}^H + u_{bt}, \tag{5}$$

where K_{bt} is plastic kilograms bought by bank b from individual depositors in month t, P_{bt}^H is the quality-adjusted plastic price index of Section 3.2, α_b is a bank fixed effect, ψ_t is a calendar-month fixed effect, and u_{bt} is an error term. I estimate it by two-stage least squares, with first stage

$$\log P_{bt}^H = \alpha_b + \psi_t + \pi Z_{bt} + \nu_{bt}, \tag{6}$$

where π measures how much the bank's quality-adjusted price moves with its exposure-

weighted upstream value. I report alongside it the reduced form

$$\log K_{bt} = \alpha_b + \psi_t + \varphi Z_{bt} + e_{bt}, \tag{7}$$

so that the elasticity can be read as the ratio φ/π of the quantity response to the price response. Bank fixed effects absorb permanent differences across banks and calendar-month effects absorb the common shifter, so identification comes from the differential response of exposed and less-exposed banks to the same movement in upstream prices. The estimation window runs from January 2022 to July 2026. After singleton removal, the sample holds 578 bank–months from 21 banks. The design requires banks with long and consistent transaction histories, since exposure is measured over one window and the response over a later one. The 21 banks that meet this requirement are not the largest in the network but the most persistent, and they carry 66 percent of its plastic kilograms (Appendix Table A2). With 21 bank clusters, conventional cluster-robust standard errors over-reject. Standard errors in parentheses are two-way clustered by bank and month, and significance stars are taken from a wild-cluster bootstrap p-value with the null imposed. For the two-stage least-squares elasticities, the tables also report 95 percent Anderson–Rubin confidence sets, which remain valid when the instrument is weak (Cameron, Gelbach and Miller, 2008; MacKinnon, Nielsen and Webb, 2023; Anderson and Rubin, 1949; Andrews, Stock and Sun, 2019).[17]

The identifying assumption is that the instrument is orthogonal to the error in equation (5) conditional on the fixed effects, $\mathbb{E}[u_{bt}\, Z_{bt} \mid \alpha_b, \psi_t] = 0$. Conditional on bank and month fixed effects, exposure-weighted movements in the upstream price of virgin plastic must affect a bank's plastic intake only through the price it posts. Banks with higher baseline plastic exposure must not experience differential supply shocks that happen to coincide with movements in the upstream price. Section 4.1.4 discusses the main threats to this assumption and reports tests of it.

4.1.2 The First Stage

Table 5 reports the first stage. A one-unit increase in the instrument raises a bank's quality-adjusted plastic price by 3.62 log points (1.08), with a first-stage F of 11.2 and a wild-bootstrap p-value of 0.011. The coefficient is the difference in the response to the virgin price between banks that differ in exposure by one unit. Calendar-month effects absorb the

[17]The bootstrap uses Webb weights and 9,999 replications clustered by bank. Anderson–Rubin sets are reported under both the clustered covariance and the wild bootstrap. In the text I quote point estimates with clustered standard errors and note where the bootstrap or the Anderson–Rubin set changes the conclusion a conventional interval would support.

response common to all banks, so the level of pass-through from the virgin price to posted prices is not identified, only the gradient in exposure. Entered in terciles, the response is a step rather than a gradient. Banks in the middle and top thirds of exposure respond by about one log point more per unit of the shifter than banks in the bottom third, and the two upper thirds are indistinguishable from each other.[18] The response to the virgin price bundles two links, from the world price to the price a bank's own buyer pays and from that price to the schedule the bank posts. The first link is where banks differ, and the onward-sales ledger of Section 4.3 shows the second at the one bank where it is observed. The first stage is therefore informative about which banks' schedules follow the plastic market and where on the menu the response arrives, not about the size of any bank's markdown. Panel B shows where the response comes from. Repricing of grades already on the menu accounts for 1.61 (0.88) of the 3.62, and grades entering the menu for the remaining 2.01 (0.79). About 55 percent of the differential response therefore arrives on lines the bank writes rather than lines it reprices, and a matched-model index that followed only continuing grades would recover 1.80 (0.97) of the 3.62, understating the response by half (Figure 4, Panel A).[19] The entering component splits into grades transacting for the first time, 1.14 (0.65), and grades returning after a silence of more than six months, 0.87 (0.63), so it is not an artifact of the reappearance rule of Section 3.2.

Table 6 separates the price at which a bank revises its menu from how often it does so. The results indicate that the shock moves the former and not the latter. Panel B reports how often a bank revises its menu in a month. Thirty percent of bank–months contain a repriced continuing grade and 40 percent a grade entering the menu, and neither probability moves with the instrument. The same holds for within-resin refinements and for the short returns that the reappearance rule treats as resupply. Menu revision is therefore no more frequent when the upstream value of plastic is high than when it is low. Panel A reports what does move. Within a grade already on the menu, the posted price rises by 1.43 (0.61) per unit of the instrument. Grades that enter the menu are written at a price further above the bank's continuing schedule than their attributes would imply, which is the entering-grade component of the first stage, 1.14 (0.65) for grades transacting for the first time. A rise in the upstream value of plastic leaves the turnover of the menu alone and raises the numbers

[18] With exposure entered as tercile indicators interacted with the shifter, in place of the continuous interaction, the first-stage coefficients are 0.97 (0.26) for the middle third relative to the bottom third and 1.29 (0.24) for the top third, with seven banks in each tercile and cutpoints at 0.25 and 0.38. A linear-in-exposure first stage would put the two gaps at 0.41 and 1.02.

[19] The unweighted mean of posted prices recovers 2.13 (0.86) of the 3.62 and the raw unit value 2.84 (0.93). By credit, 57 percent of the entering component is finer grades of a resin the bank had stopped buying and brings back, 39 percent is finer grades added beside the grades it carries, and 4 percent is new resins or coarser additions, and this composition does not move with the instrument (Appendix Table A19, Panel A).

on it, both on the grades the bank already carries and on the grades it would have added anyway.

Table 6, Panel C, reports heterogeneity in the price response across the lines a bank carries by who supplies them. I find that the response is concentrated on picker lines. Sorted PET, which pays four times the unsorted price and is 98 percent picker-supplied, makes up 71 percent of pickers' plastic kilograms (Table 4). Households deliver mostly unsorted material, with mixed PET at 35 percent of their plastic and the rest spread across lines defined by object rather than by sorted resin. With line and bank-by-month fixed effects, so that the level of the shock is absorbed, a continuing line's posted price rises by less the larger the household share of its kilograms, −0.73 (0.37) per unit of share, and by −1.60 (0.63) when the share is measured in 2022 and the outcome over 2023–26 so that the share is predetermined (Table 6, Panel C). Allowing PET lines their own response steepens the gradient to −1.17 (0.54), and the PET term itself is negative, −0.80 (0.46). On PET, the bank does not pass the shock through by repricing the lines it already carries. The posted prices of continuing PET grades, coarse and fine, do not rise with the instrument, while those of the other resins do (Appendix Table A16, Panels A and B). PET value instead reaches sellers through the finer PET grades the bank adds or re-lists, which post 10.1 (6.7) log points further above the coarsest PET grade, and 5.7 (3.2) above the bank's continuing schedule, than they otherwise would. These within-resin comparisons use only bank–months in which the bank transacts at least two PET grades, 88 bank–months at 10 banks for the entering grades, and I read them as a consistent pattern rather than as precise estimates. A picker thus faces a different menu after the shock on two counts. The numbers on the lines she already supplies rise by more than those on household lines, and finer PET grades appear above them at a premium. This is where a quantity response should appear, and I turn to that next.

4.1.3 The Elasticity of Supply

Table 5, Panel A, reports the supply elasticity, first for all individual depositors and then separately for households and pickers. The reduced form on log plastic kilograms is 4.26 (2.26), and two-stage least squares gives a supply elasticity of 1.18 (0.74). A ten percent increase in the quality-adjusted price a bank posts brings in roughly twelve percent more plastic from individual depositors. Ordinary least squares gives 0.43 (0.26), the attenuation one expects when banks raise prices where material is scarce. The resin-PPI shifter gives a first stage of 6.48 (2.30) and an elasticity of 1.26 (1.03), directionally the same but less stable across windows, and I treat it as corroborating rather than as an equally preferred instrument (Appendix Table A7). The aggregate elasticity is less precisely estimated than

the price response that generates it. The clustered Anderson–Rubin set reaches zero at its lower end and the bootstrap set contains it.[20] The first stage explains the imprecision. The shock reprices the lines pickers supply and leaves the household lines alone, so the aggregate averages a response with a non-response and understates the elasticity of the sellers whose prices moved. I therefore estimate the two separately.

Picker depositors have a supply elasticity of 1.97 (0.78), significant at the five percent level under the bootstrap and with an Anderson–Rubin set that excludes zero under either form of inference. Household depositors have an elasticity of 0.64 (0.59) whose confidence set comfortably contains zero, so no household response is statistically detectable.[21] Both types face the same bank-level price index and may deliver to any grade on the menu, but only pickers appear to respond to it. The difference between the two groups is better identified than either level, because it comes from within-bank, within-month variation across sellers facing the identical schedule. Stacking the two panels with bank-by-type and month-by-type effects, and instrumenting the price and its interaction with the picker indicator, gives a gap of 1.34 (0.63), $p = 0.046$, on 940 bank–months and 21 bank clusters. Dropping one bank at a time, the gap ranges from 1.14 to 1.90, never changes sign, and remains significant at the five percent level in 10 of the 21 samples.

Margins. Both types can respond on two margins. More sellers can participate, and each participating seller can bring more. Because $\log K_{bt} = \log N_{bt} + \log(K_{bt}/N_{bt})$, where N_{bt} is the number of active plastic depositors, the elasticities of the two components sum exactly to the total,

$$K_{bt} = N_{bt} \times \frac{K_{bt}}{N_{bt}}, \tag{8}$$

and the same decomposition can be estimated by type,

$$Y^m_{gbt} = \alpha^m_{gb} + \psi^m_t + \kappa^m_g \widehat{\log P^H_{bt}} + u^m_{gbt}. \tag{9}$$

Table 7 reports the results. In the pooled sample, participation carries 0.95 (0.59) of the 1.18 and kilograms per depositor 0.22 (0.54). The participation elasticities, 1.09 (0.41) for picker and 0.84 (0.56) for household depositors, cannot be distinguished from each other, with a gap of 0.25 (0.60). The groups separate on the intensive margin, with elasticities of 0.89 (0.65) for pickers and −0.20 (0.55) for households, a gap of 1.09 (0.63) with $p = 0.10$.

[20] The clustered Anderson–Rubin set is $[0.00, 3.75]$ and the Webb wild set $[-0.45, 5.25]$. The bootstrap p-value on the reduced form is 0.17 (Table 5).

[21] The Webb wild Anderson–Rubin sets are $[0.45, 9.65]$ for picker and $[-1.20, 3.30]$ for household depositors, with bootstrap reduced-form p-values of 0.033 and 0.386.

A higher price draws in more pickers and raises collection per picker. Whether it draws in additional households I cannot tell, but a household that already sells to the bank does not bring more, consistent with household supply being bounded by own consumption. No individual margin is significant at the five percent level in any column, and the within-seller version of the intensive margin, log kilograms of a continuing picker on the instrument with seller and month effects, is positive in every specification I have run and precise in none, at 1.27 (1.38) in Appendix Table A17. I read the decomposition as evidence on the direction of the difference between the groups rather than on the size of either margin.

Where the response lands. The additional plastic comes almost entirely from pickers, and Table 8 asks on which lines it arrives. I classify every line each month from its own price history and regress picker kilograms on lines of each status, scaled by the bank's mean monthly picker kilograms, on the instrument. Because the outcomes sum to total picker kilograms, each coefficient is a share of the response. Lines whose posted price did not change carry 28 percent of picker kilograms and 80 percent of the response, and lines repriced up carry 40 percent of the kilograms and 58 percent of the response. Lines repriced down lose kilograms. Lines transacting for the first time carry 1 percent of picker kilograms and none of the response, whereas lines returning after a silence of more than six months carry 1 percent of the kilograms and 28 percent of the response. These returning lines are the re-listed finer grades through which PET value reaches pickers. The components are individually imprecise, but about three-quarters of the response arrives on lines the picker already supplies. The finer grades a bank adds draw pickers' material up the grade ladder over the months after they enter rather than at once, a lag consistent with sellers learning that a new line is on the board and changing how they sort, and one the entry design shows directly (Section 4.2). Appendix Table A15 reports the components and a within-menu version. With line and bank-by-month effects, picker kilograms tilt toward lines transacting for the first time relative to lines at an unchanged price, by 0.83 (0.29) in inverse hyperbolic sine. The two readings are consistent. The within-menu estimate is a proportional change on lines that carry one percent of picker kilograms and are often at zero, so a large proportional tilt is a negligible share of the kilograms, which is what the additive decomposition reports. In kilograms, the response arrives on the lines a picker already supplies; in proportional terms, the new lines grow fastest, from a base near zero.

Interpretation. My results suggest that the marginal supplier to a waste bank is a picker, for whom recovery is an income-generating activity with a scalable effort margin and an outside option, not a household disposing of its own waste. The gap is above one and the

household elasticity is indistinguishable from zero, which is what one expects if the price is a small part of what a household that delivers for civic reasons gets from doing so. The margins point the same way. A picker who already sells to the bank brings more when the price rises, because she can search longer, carry more, or divert a load she would have sold elsewhere. A household that already sells to the bank brings what her consumption generates and cannot bring more at any price. A continuing household's plastic kilograms do not rise with the instrument, 1.08 (0.84), even though the schedule she faces does, 2.91 (1.04) (Appendix Table A17). The picker response is a quantity response rather than a change in sorting. A picker faces a higher number on the sorted grades she already delivers and a new premium on finer grades she does not. In the short-run, she takes the first. The attribute value of her load does not rise, -1.21 (1.07), nor does her share on grades finer than the resin's coarsest, 0.46 (2.14). Sorted PET is already 71 percent of what she delivers, and when the price rises she brings more of it. The new grades fill more slowly. Their premium is paid as sellers learn they are on the board and sort to them, to the resin's incumbents and to sellers new to the bank, so the contemporaneous elasticity measures the response on the lines a picker already supplies and does not capture the response on the new grades, which arrives over the following months.

To put the picker elasticity in perspective, a bank that sets its price against a supply elasticity of 2 marks down by one half of the price it pays, or one third of the price it receives (Section 5.1 states the pricing rule, and Appendix A derives it). That is roughly the size of the gross spread recorded in the onward-sales ledger of Section 4.3, so the markdown the elasticity implies is about what one bank actually earns between buying and selling, which gives some confidence that the elasticity is of the right order. A bank cannot post one price to pickers and another to households. When it raises its price, pickers bring more material and households bring what they would have brought anyway. The same elasticities imply that a cut loses pickers and keeps households, and the entry design, which lowers an incumbent's price relative to its sellers' outside option, shows that loss directly. In this sense, a waste bank has market power over households but not over pickers, and the menu margin of Section 4.1.2 is the same fact seen from the bank's side. It raises the numbers on the lines pickers supply and posts its PET premium on the finer grades it adds, while the unsorted PET line that households supply is the line whose number does not move.

Bank-level and market supply. It is important to note that the elasticity estimated here is the supply facing a single bank, not the supply of recyclable plastic to the market as a whole. The distinction matters because the bank-level elasticity is the object that determines a bank's markdown, whereas a system-wide subsidy moves the market elasticity, which is

smaller. In my estimation, month fixed effects absorb the common movement in upstream prices, and because almost every local market contains a single exposed bank, the instrument moves that bank's price relative to the other buyers its sellers can reach. Off-platform buyers such as junkshops are not in the data. If their prices also rise with the upstream value of plastic, a picker's outside option rises with the instrument, and the estimate understates the response to a change in the bank's price alone. Part of the response may also be material redirected from other buyers rather than new recovery. The bank-level elasticity is the right object for a change in one bank's price, but for a system-wide subsidy it is an upper bound on the collection response. Section 5 treats the share of the response that is new collection as a parameter.

4.1.4 Threats to Identification

This section reports tests of the identifying assumption.

Persistence of the upstream series. The PET import price is highly autocorrelated, at 0.945 over the 55-month window, so the design rests on two to three independent episodes rather than on month-to-month variation (Appendix Figure A6 plots the exposure gradient of the price month by month against the shifter). A persistent series raises a specific concern. If plastic-intensive banks were on a different trend from other banks for reasons unrelated to the value of plastic, the instrument, which interacts exposure with a slowly moving series, would be correlated with that trend. The first stage would then reflect differential growth across banks rather than the response to upstream prices. One test is to allow each bank a linear trend proportional to its exposure. Doing so lowers the first stage from 3.62 (1.08) to 1.67 (1.41). This specification is demanding, because the deflator is itself nearly a trend, and much of the identifying variation is absorbed along with the confound. A more direct test retains only the high-frequency variation in the shifter. I replace each shifter with its deviation from a centered moving average. By construction, this deviation is orthogonal to any trend that is linear in exposure, so the test isolates the episodes in which upstream prices moved against their recent average. The first stage is 3.74 (1.61) with a 24-month window and 3.00 (1.28) with a 36-month window (Appendix Table A8). The price response is therefore not an artifact of the common drift in upstream prices. The aggregate quantity elasticity, however, remains imprecise under either version.

Other shocks to plastic-intensive banks. The identifying assumption requires that exposure-weighted movements in the upstream price affect a bank's plastic intake only

through its posted price. This assumption fails if banks with high plastic exposure experience other shocks that coincide with movements in the PET price. Four tests address this concern (Appendix Tables A3 and A4). First, leads of the instrument do not load on the price, so exposed banks were not already adjusting before the upstream price moved. Second, exposure is unrelated to observable bank characteristics, which limits the scope for exposure to proxy for some other dimension along which banks differ. Third, paper exposure interacted with the PET price produces no plastic-price response. The result is therefore specific to plastic and not a property of banks with a concentrated business of any kind. Fourth, the response is to the PET-specific component of the import price rather than to crude oil, which affects the wider economy through fuel and transport costs. Projecting crude out of the PET price leaves the first stage at 4.05 (1.38).

Placebo materials. If the instrument picked up a shock to exposed banks' business as a whole rather than the value of plastic, the prices these banks post for other materials would move with it. Applying the same index construction and first stage to paper and glass gives first stages of 1.45 (1.00) and 0.75 (1.32), and the continuing-grade components of both are indistinguishable from zero (Appendix Table A5). The metal first stage is negative, -4.43 (1.58), but it comes entirely from the price at which dormant metal lines are re-listed after a long silence. Metal menus consist of two or three distinct products that banks buy intermittently, so the reappearance rule scores a returning metal line against a different product rather than against close substitutes. An index that treats only first-ever grades as new gives -1.00 (0.87) for metal, against 2.75 (1.03) for plastic. The instrument therefore moves the price of plastic and not the prices of the other materials a bank buys, consistent with the identifying assumption that it operates through the value of plastic alone.

Measurement of menu revision. The entering-grade component carries about half of the first stage, and it depends on a rule that scores a grade returning after more than six months of silence as new. The concern here is one of measurement. If a returning grade reflects intermittent supply rather than a decision by the bank to re-list it, then the rule would classify ordinary fluctuations in what sellers happen to bring as changes to the menu. The price response would then be overstated, since part of the price variation attributed to the bank's menu would instead reflect the composition of deliveries. Two facts indicate that the rule measures menu revision rather than intermittent resupply. Grades transacting for the first time, which cannot be resupply, carry 1.14 (0.65) of the 2.01. And grades that return within six months transact at exactly their previous price 69 to 84 percent of the time, whereas grades returning after a longer silence are repriced to the bank's current schedule

(Appendix Table A21). Shortening the rule to three months, lengthening it to nine, or scoring every return as new leaves the first stage between 3.64 and 4.67 and the elasticity between 0.91 and 1.17 (Appendix Table A7).[22]

4.2 The Arrival of a Rival

4.2.1 Design

The supply design traces out a bank's residual supply curve by moving the price it posts. Competition moves the curve itself. When a rival opens nearby, each seller gains an alternative buyer. The quantity an incumbent can retain at any given price falls, and the incumbent must decide whether and how to respond. I date openings from the verified network roster rather than from an entrant's first recorded transaction. For each incumbent, treatment is the first roster-dated opening within 15 kilometers that occurs while the incumbent is an established, recording bank. The treated incumbents are young when the rival opens. Dated from its roster opening, the median incumbent is 10 months old (mean 26, standard deviation 39, since a few banks date from the 2000s) and has a median of 4 recorded months before the opening (mean 6, standard deviation 6). Because a month without a record is not evidence that a bank has stopped operating, I do not impute closures and measure every outcome on months in which the bank records (Section 3.3). The requirement to record in months 6 to 11 keeps the comparison from being dominated by accounts that stop recording shortly after they start. It applies symmetrically, so that treated and comparison banks record at the same rates over the first year and a half after an opening (Appendix Table A11). The recording requirement is the key specification choice in this design, since an incumbent that stops recording after a rival opens leaves the estimation sample, and Appendix Table A9 reports how the estimates change when it is relaxed or tightened. Comparison banks are never treated or not yet treated at the horizon in question, with a not-yet-treated bank's months kept only while it is untreated.

The event universe contains 97 treated incumbents and 51 distinct entrants at a median distance of 8.0 kilometers. Twenty-nine incumbents face a rival within 5 kilometers and 68 a rival between 5 and 15. Appendix Table A10 compares treated incumbents with their comparison banks before the opening. The two groups are similar in scale and in recent growth. Plastic and total kilograms, depositor counts, and plastic growth over the preceding

[22]The estimates are also robust to other choices in the construction of the index and to the exclusion of any single bank. Across seven variations in the attribute vocabulary, weighting, new-grade cap, and code attributes, the first stage ranges from 2.82 to 4.19 and the elasticity from 1.02 to 1.51, and leaving one bank out at a time gives first stages from 3.17 to 4.93 and elasticities from 0.51 to 1.53 (Appendix Table A7).

half-year differ by a tenth of a standard deviation or less, and the difference in growth is −0.03 (0.24) log points. Treated incumbents are newer on the platform, with 6 recorded months before the opening against 14, and had been open for 26 months on average against 39 (normalized difference −0.35, $p = 0.21$). A smaller share of their plastic comes from picker sellers (0.22 against 0.46), with correspondingly fewer plastic grades on the menu, they post a higher plastic price (normalized difference 0.31), and more of them are on Java (0.72 against 0.56). Bank-age fixed effects and the alternative comparison groups of Appendix Table A9 address the difference in age. The symmetric recording rule holds recording history fixed, and the differences in clientele, price level, and location are level differences absorbed by the bank and line fixed effects.

Identification relies on a parallel trends assumption. Conditional on bank, month, and age fixed effects, the outcomes of treated incumbents would have evolved in parallel with those of their comparison banks in the absence of the opening. This assumption does not require that the location of entry be as good as random. Entrants may select into markets with more sellers or with more successful incumbents, and such differences in levels are absorbed by the bank fixed effects. The assumption is violated if entry is correlated with the incumbent's pre-existing trajectory, for example if rivals enter markets where incumbents are already losing sellers. In that case, the estimated effect of entry would be biased upward. The pre-period provides a direct check on this assumption. Treated and comparison banks exhibit similar growth in plastic over the six months preceding the opening, with a difference of −0.03 (0.24) log points (Appendix Table A10). The event-study coefficients for the pre-period, reported in Section 4.2.2, also provide no evidence of differential trends prior to entry. Treated incumbents are younger on average than their comparison banks, though this difference is absorbed by the age fixed effects, and Appendix Table A9 shows that the estimates are robust to alternative comparison groups.

Quantities are monthly totals over individual depositors and are estimated in levels by Poisson pseudo-maximum likelihood on a stacked panel. Each treated incumbent s defines a stack containing its own bank–months and those of its comparison banks over event months −12 to 23. I estimate the quantity response in bins of event time, −12 to −7, 0 to 5, 6 to 11, 12 to 17, and 18 to 23 months, with months −6 to −1 omitted as the reference period, and cluster standard errors by the entrant, because one opening can treat several incumbents. Stacking each opening with its own comparison banks is the standard remedy for the failure of two-way fixed-effects event studies under staggered treatment, in which units treated earlier serve as controls for units treated later and effects that grow with time since treatment receive negative weight (Goodman-Bacon, 2021; Cengiz et al., 2019; de Chaisemartin and D'Haultfœuille, 2020; Callaway and Sant'Anna, 2021). Every comparison here is between a

treated incumbent and banks that are untreated at the horizon in question, so the coefficients are averages of clean two-by-two contrasts. The coefficients β_k are log-point changes in expected quantity relative to the incumbent's own comparison banks and to its own reference period. I estimate in levels by Poisson because the outcomes are counts and kilograms with many small values, and a log or inverse-hyperbolic-sine transformation would let a bank that grows from a few kilograms to a few dozen dominate the average (Chen and Roth, 2024).

4.2.2 Who a Rival Takes

For each outcome, I estimate the Poisson regression

$$\mathbb{E}\big[Q_{bts} \,\big|\, b, t, s\big] = \exp\Big(\alpha_{bs} + \lambda_{ts} + \gamma_{a(b,t)} + \sum_k \beta_k \, D^k_{bts}\Big), \tag{10}$$

where $D^k_{bts} = \mathbb{1}[b = s] \cdot \mathbb{1}[t - E_s \in \mathcal{B}_k]$ equals one when bank b is the treated incumbent of stack s and month t falls in event-time bin $\mathcal{B}_k$, Q_{bts} is the monthly quantity at bank b in month t within stack s, α_{bs} and λ_{ts} are bank-by-stack and month-by-stack fixed effects, $\gamma_{a(b,t)}$ is a fixed effect for the bank's age in months since its roster opening, E_s is the month of the opening that defines stack s, and β_k is the effect in bin $\mathcal{B}_k$. The exponential conditional mean is the Poisson specification, so the β_k are log-point changes while the outcome stays in levels.

Table 9 and Figure 7 report the response pooled over all 97 treated incumbents, with the pre-opening level of each outcome beside it. An incumbent loses pickers and their material, and keeps its households. In the year after the opening, plastic kilograms fall by 0.45 (0.21) log points, about 37 percent relative to comparison banks, and among incumbents still recording two years after the opening the coefficient is −0.35 (0.16).[23] The coefficient at −12 to −7 months is −0.02 (0.14), so there is no evidence of a pre-trend in plastic.[24] The loss is absent at 0–5 months, 0.05 (0.10), and appears at 6–11. The roster dates an entrant's registered opening, which can precede the month in which it is buying at scale. No other

[23]Later bins are estimated only from incumbents that are still recording, 57 at 12–17 months and 30 at 18–23, so the trajectory compares different sets of banks, and I report the second-year cells to show its shape rather than as precise estimates. The year-one loss comes from incumbents that record in months 6–11 but not afterward. Requiring a record in months 12–23 as well leaves 71 incumbents and a plastic estimate at 6–11 months of 0.11 (0.20) (Appendix Table A9). Because the data do not record closures (Section 3.3), I cannot determine whether these banks exited or stopped recording. Household-intensive incumbents that record into a second year also begin acquiring picker sellers and posting grades for them over that year, so the second-year cells by incumbent type describe which banks survived and changed rather than the response of a fixed type.

[24]Total kilograms at −12 to −7 months are 0.25 (0.09) among the 33 incumbents old enough to be observed then. The pre-trend is in paper and glass at the older incumbents, not in plastic, and the age fixed effects do not remove it.

outcome responds. Paper, glass and metal kilograms are unchanged, 0.02 (0.06), 0.02 (0.32) and −0.26 (0.30) at 6–11 months, and so is the total number of depositors, −0.04 (0.07).

Splitting plastic by the type of seller who delivered it shows that the loss comes entirely from pickers. Plastic delivered by picker sellers falls by 0.76 (0.14) log points at 6–11 months, a decline of more than half, whereas plastic delivered by household sellers does not change, 0.02 (0.13). The two series sum to the total, so the whole of the plastic loss is pickers' material. The pickers' loss is robust to the choice of reference period, the bank-age fixed effects, alternative comparison groups, and the exclusion of the largest opening cohort (Appendix Table A9).[25] Depositor counts show the same pattern. Picker accounts active in a month fall by 0.19 (0.09) log points, about a fifth, while household accounts are unchanged, 0.04 (0.07). The pre-opening levels explain how a bank can lose half its pickers' plastic without a visible change in its depositor count. The average treated incumbent buys 488 kilograms of plastic a month from 26 depositors. Only 2.4 of those depositors are pickers, but they deliver 409 of the 488 kilograms. A fifth of these pickers is under 2 percent of the bank's depositors, but a 53 percent fall in the plastic they deliver, on 84 percent of the bank's plastic, is roughly 45 percent of its plastic, in line with the pooled loss of 37 percent. The sellers whose deliveries respond to a bank's own price in Section 4.1.3 are also the sellers who respond to a rival's.

Table 10 examines the same pattern from the incumbent's side, classifying incumbents by the share of their plastic that picker sellers delivered before the opening and splitting at the treated median of 0.12 (Appendix Table A12 reports every outcome for each type). Picker-intensive incumbents lose 0.62 (0.17) log points of plastic, −0.85 (0.11) of the plastic their pickers delivered, and a quarter of their picker accounts, −0.24 (0.09). Household-intensive incumbents show no loss, 0.13 (0.17), and their household accounts are unchanged, −0.06 (0.06). Estimated in one regression with the bins interacted with type, the difference between the two types is 0.77 (0.22) log points in the plastic response and 0.41 (0.20) in the picker account response. Both types are compared with the same pool of untreated banks. Restricting each incumbent's comparisons to banks of its own type reproduces the picker-intensive estimates and leaves the household-intensive estimates too imprecise to be informative, because household-supplied banks record more intermittently and few of them record long enough to serve as comparisons (Appendix Table A9).

Appendix Table A13 reports the loss by distance to the rival. The estimates do not vary with distance. Pickers' plastic falls where the rival is between 5 and 15 kilometers away,

[25] Across the alternative reference periods, age fixed effects, qualification rule, and cohort exclusion, the pickers' loss at 6–11 months ranges from −0.66 to −1.44 and the pooled plastic loss from −0.34 to −0.79. Under the age-matched and reweighted comparison groups the pickers' loss is −2.25 (0.42) and −1.30 (0.36), and under all three alternative comparison groups the plastic loss is larger, between −0.93 and −1.36. Adding unit-bank and institutional accounts to the quantities gives a plastic loss of −0.40 (0.19).

−0.68 (0.17), and where it is within 5, −0.94 (0.25) (Appendix Table A13). I read this as evidence that pickers are mobile within this radius. Some travel farther for a better price, which is how a rival several kilometers away draws them from an incumbent, and for others the rival may simply be closer to where they collect or live, so that the move reflects lower travel cost rather than a higher price. The data do not separate the two, and the loss within 15 kilometers is of the same size either way. Because seller identifiers are bank-specific, the design also cannot distinguish a seller who moves to the entrant from one who leaves the platform. It measures what an incumbent loses rather than what the market recovers, which is the redirection share that Section 5 treats as a parameter.

4.2.3 How Incumbents Respond

The posted price of a grade on the menu, the modal credit per kilogram of a bank–grade line ℓ, is estimated on the same stacks by least squares,

$$\log p_{\ell ts} = \alpha_{\ell s} + \lambda_{ts} + \gamma_{a(\ell,t)} + \sum_{k} \theta_k \, D^{k}_{b(\ell)ts} + \varepsilon_{\ell ts}, \tag{11}$$

where $\alpha_{\ell s}$ is a line-by-stack fixed effect, θ_k is the price effect in bin $\mathcal{B}_k$, and the remaining terms are as in equation (10). Every grade enters, including grades added after the opening, and the line fixed effect measures each price against the grade's own mean, so the price estimates do not depend on the reference period.

The textbook prediction for a buyer facing one kind of seller is that competition raises the price it pays. The incumbents in this setting, however, respond differently. Table 11 and Figure 4, Panel B, report their response, pooled over all treated incumbents. The posted price of a plastic grade already on the menu drifts down after the opening, by 3 log points in the first year (standard errors 0.01 and 0.02 in the two bins) and by 4 to 6 points in the second, when the estimates are less precise. The coefficient before the opening is 0.02 (0.03). The posted prices of paper and glass grades do not move, and metal rises only in the second year, on few lines. The quality-adjusted plastic index does not move in any bin, at −0.02 (0.02), −0.03 (0.05), 0.05 (0.05) and 0.06 (0.09), and neither does its continuing-grade or its new-grade component. The two measures weight grades differently, and the gap between them is informative. The line-level estimate counts every grade equally, whereas the index weights grades by credit, and the grades on which pickers deliver carry far more kilograms than those on which households do, so a cut on a low-volume household grade moves the index little.

The cut in posted prices is only one part of the incumbent's response. At the same time,

as they lower the price on existing grades, incumbents also add new, finer grades to the menu that pay substantially more. The probability that an incumbent adds a grade finer than the coarsest grade of a resin it already buys, a within-resin refinement, rises by 7 (2) percentage points at 0–5 months, 11 (3) at 6–11, 15 (4) at 12–17 and 19 (7) at 18–23, on a base of 11 percent of recorded bank–months, with a pre-period coefficient of −0.03 (0.04). Incumbents refine about twice as often as before in the first year and nearly three times as often in the second. The two placebo additions do not rise. Additions that are coarser than or equal to the bank's existing grade of the resin move by 4 (3) points at most, and additions in a resin the bank did not carry are not significant at the five percent level in any bin (they are too few, 141 events, for the stacked count regression to be stable, and I do not tabulate them). The response is specific to refining a partition the bank already has.

Table 12 describes what a refinement is. Of 850 first transactions on a finer same-resin grade, 73 percent are posted above the coarse grade's current price, with a median gap of 0.69 log points, about twice the coarse price, and the gap is the same for refinements that follow a rival's opening. After a refinement, the coarse grade loses 10 to 12 points of its share of the resin's kilograms, which move to the new grade, and the new grade's kilograms come from pickers (88 percent) and from sellers whose first deposit at the bank falls in or after the month of the refinement (69 percent). Table 13 shows who uses the grades added after a rival opens. Of the 90 treated incumbents with post-opening plastic, 63 post at least one plastic grade that was not on their menu before the opening. At those incumbents, the new grades carry half of post-opening plastic kilograms and are priced 0.21 (0.08) log points above the same bank's continuing grades, a raw comparison that makes no adjustment for attributes. Picker sellers deliver 88 percent of the kilograms on new grades, against 72 percent on the same banks' continuing grades. Within two years, 55 percent of all picker-delivered plastic moves onto the new grades, against 31 percent of household-delivered plastic. A picker delivers about 200 kilograms a month on new grades, a household about 4. The grades are open to everyone, and 47 percent of household sellers deliver on one at least once, but a household that does so brings a few kilograms while pickers move their loads. Most household material remains on the existing grades, whose posted prices fell.

Taken together, these results imply that the incumbent does not pay more for a given quality. It posts a line that only a seller who sorts can claim, at roughly twice the coarse price, and the quality-adjusted index, which nets the premium out, records no change in the price of given quality. A seller's earnings, however, can and do change. A picker who was delivering clear bottles into the mixed grade can now sort them to the clear grade and be paid about twice as much for the same material. Conditional on sorting, the refinement raises the price available to her by as much as a doubling of the posted price would have.

That pickers take up the offer is not in doubt. They are the sellers whose supply responds to price (Section 4.1), half of a bank's differential price response to a higher upstream value arrives on the grades it re-lists or adds (Table 5, Panel B), and the new grades here are populated by pickers, most of them new to the bank.

The price response does not depend on how close the rival is. The cut in existing-grade prices is −0.06 (0.02) at 12–17 months for a rival within 5 kilometers and −0.04 (0.03) for one between 5 and 15, and by band it is −0.07 (0.02) at 2–5 kilometers, −0.03 (0.01) at 5–10 and −0.06 (0.07) at 10–15 (Table 14). The cut is larger when further rivals follow. For an incumbent whose first rival is followed by a second opening within 15 kilometers, measured against once-treated incumbents with no second rival, the posted price of existing grades falls by 0.04 (0.01) log points immediately and by 0.07 (0.03) at 12–17 months, with a pre-period coefficient of 0.00 (0.04) (Appendix Table A14). Refinement is present at every distance. The probability of a within-resin refinement rises by 15.5 (7.6) points at 6–11 months and 25.0 (8.9) at 12–17 for a rival within 5 kilometers, and by 9.7 (3.7) and 12.0 (4.2) for one between 5 and 15 (Table 11, Panels B and C).

Appendix Table A12 reports the menu response by incumbent type. The response is concentrated among the incumbents whose sellers sort. Picker-intensive incumbents refine their menus by 8 (4), 17 (5), 19 (6) and 16 (12) percentage points in the four bins after the opening, and household-intensive incumbents by 5 (3), 6 (4), 11 (5) and 14 (6) (Appendix Table A12). The design does not separate the two types on the posted price. The difference between them is 0.02 (0.03) at 6–11 months and 0.05 (0.05) at 12–17, which reflects imprecision rather than a precise zero, because household-intensive incumbents post few grades and reprice them rarely.

Interpretation. My results show that incumbents bid for pickers through the menu rather than through the posted price. When a rival draws away their mobile sellers, they add finer grades that only sorted loads can claim and trim the posted price on the grades they already carry. Raising the posted price on the coarse grades would be the more costly response. A posted price is paid on every kilogram of the grade, and many of those kilograms are delivered by households who would deliver them anyway at the old price. Bidding for the departing pickers through the price on the coarse grades would pay the increase mostly on material that is not at risk. A finer grade is a price increase that applies only to sorted loads. The premium is available to everyone and claimed by pickers, and the incumbent does not need to know who is who. In this way, the bank raises the effective price a picker receives without raising the price on a household's kilograms. The cut in the posted price of the existing grades follows from the same logic. The sellers who remain on those grades

after the opening are disproportionately households, whose supply barely responds to the price, so the incumbent's residual supply on those grades is less elastic than before, and the markdown that maximizes its surplus is therefore wider. A uniform increase on the coarse, mixed grades is also costly for a second reason. Picker supply can expand quickly, and a bank that raised the price of mixed plastic could take in more unsorted material than it can sort itself or place with its onward buyers at the current price, whereas a premium on a finer grade raises intake only of material that is already sorted to a buyer's specification. Competition may also push a bank to seek out new off-takers, who often have more specific material requirements, so that the grades a bank adds after a rival opens reflect the buyers it has found as much as the sellers it is trying to keep.

A natural question is why an incumbent does not price this way before a rival arrives. One explanation is the fixed cost of a finer grade. Posting one requires a downstream buyer who pays a premium for it, enough sorted volume to make up a shipment, and the handling and storage of an additional line, which together make up the fixed cost of carrying a grade. A bank adds the grade only when the gain from paying pickers and households differently exceeds this cost. That gain is increasing in the gap between the two types' elasticities, and before a rival opens the gap is small, because a picker with no nearby alternative is nearly as captive as a household. Entry raises pickers' mobility. It raises the share of a picker's response that is loads moved between buyers, widens the gap between the two groups' elasticities, and pushes the return to separating the types above the fixed cost at incumbents where it had been below. A second explanation is incomplete information. It may be difficult to discern the responsiveness of sellers before a rival reveals it, and a bank that learns its pickers can leave adjusts its pricing accordingly. A third is that competition pushes incumbents to search for new off-takers, who may require particular or more finely sorted grades, and the incumbent revises its schedule to meet that demand. The onward-sales ledger of Section 4.3 shows that a bank sells each material to several buyers at prices that differ by grade, so such buyers exist. I cannot observe an incumbent's search for them, but a rival that raises the return to separating the two types also raises the return to finding the buyer who makes the separation worth posting, and the two margins move together.

4.3 Onward Sales

The supply elasticity measures a bank's market power from the sellers' side, through how much material a higher price brings in. The onward-sales ledger of one large operator allows a look from the bank's side. The transaction data record only what a bank pays, p. For this operator I also observe what it receives when it sells onward, P, for every shipment over

two years (Section 3.5). A bank with market power pays its sellers less than their material is worth to it, and the gap between the two prices is the room it has to do so. Because the bank also bears a handling cost c out of P, the gross spread $P - p$ bounds the markdown from above, and its ratio to the price paid, $(P - p)/p$, bounds the inverse elasticity $1/\varepsilon$. The residual supply the bank faces is at least as elastic as the spread implies. The ledger also shows how the bank's buying price moves with the price it receives, the last link of the chain that Section 4.1 observes only from its upstream end, and how many buyers the bank sells to, which bears on the assumption in Appendix A that the downstream price is taken as given. Table 15 and Figure 8 report the results. The operator is a hub bank that buys about half its kilograms from feeder banks, so what follows describes one large node of the chain rather than a typical bank.

The spread is large. Across the 1,526 tons the bank sold in 2024–25, it received Rp 3,580 per kilogram and, valued at its own material-level buying prices, had paid Rp 2,396, a gross spread of 0.33 of the sale price, or 0.49 of the price paid (Table 15, Panel A). The spread is narrowest for PET, 0.27 of the sale price and 0.37 of the price paid, and for used cooking oil, 0.31 and 0.45, the two materials with the highest unit values. It is 0.39 and 0.64 for paper, 0.38 and 0.62 for metal, 0.42 and 0.73 for other plastics, and widest for glass, 0.64 and 1.74, which the bank buys at Rp 212 per kilogram and sells at Rp 581. The PET spread is earned largely on the grade partition. The bank buys mixed bottles at Rp 6,526 per kilogram and sells the clear bales it sorts from them at Rp 8,806, a third more, so the partition that banks use to adjust the price they pay is also, at this bank, the main source of its buy–sell margin. If the whole spread were markdown, the ratios would bound the elasticity of this bank's residual supply from below at about 2.0 overall, 2.7 for PET, and 1.4 to 1.6 for other plastics and paper. The picker elasticity of 2.0 in Table 5 implies a markdown ratio of 0.50, which is the size of the overall spread, larger than the PET spread, and smaller than the paper and other-plastic spreads. Handling costs are part of every spread, and half of this bank's supply comes from other banks, whose supply should be more elastic than an individual picker's. The bounds are therefore consistent with the elasticities of Section 4.1.3 but do not tighten them.

The bank's buying price moves one for one with the price it receives. The cleanest comparison is between the bank's largest sale product, pressed clear PET bales, and its largest purchase grade, loose cap- and label-free PET bottles of mixed colour (grade P12 Mix), from which the bales are made. Over the 24 months the two unit values move together (Figure 8, Panel B). A regression of the log price paid on the log price received gives a coefficient of 1.09 (0.06) in levels, 0.98 (0.06) in first differences, and 1.07 (0.12) cumulated over the current and two lagged months (Table 15, Panel B). Between September and December 2025 the

bale price fell by 0.36 log points, from Rp 8,919 to Rp 6,200 per kilogram, and the price paid for loose bottles fell by 0.38 within the same three months. The gross spread on the pair averages 0.26 of the sale price, ranges from 0.20 to 0.33, and is the same before and after the fall, 0.26 against 0.25 (Figure 8, Panel C). This is what a locally constant elasticity predicts, since then $p = v\varepsilon/(1+\varepsilon)$, so that the log price paid moves one for one with the log value and the spread is a constant share of the price. Pooling materials gives a similar answer, 0.61 (0.27) in levels and 0.97 (0.16) cumulated over three months. At the level of individual grades the estimates are smaller, 0.24 (0.10) in levels and 0.55 (0.18) cumulated, because grade-month unit values are noisier and the posted prices of the bank's smaller grades are revised less often than its PET price, the rigidity documented in Section 3.1.

Pass-through of one in proportional terms is what a buyer with market power is expected to show, not evidence against it. A bank that pays a fixed fraction of the value of its material, as in equation (26), raises what it pays by 10 percent when that value rises by 10 percent whatever fraction it keeps, so the markdown appears in the level of the price and not in how the price moves. A buyer with no market power that pays the value of its material less a fixed handling cost would instead pass a rupiah change through in full and, because it pays a smaller number than it receives, a proportional change through by more than one, 1.35 at this pair. In levels, the price paid moves by 0.83 (0.05) rupiah per rupiah of the price received, and the spread moves with the sale price, by 0.17 (0.05), where a fixed handling cost predicts zero.[26]

The bank sells to few buyers, and to the same ones month after month. It sold to 80 distinct buyers in 2024 and 92 in 2025, but the set of buyers for any one material is small. PET bales went to 24 and 26 buyers, paper to 23 and 24, and cooking oil and glass to three and five (Table 15, Panel C). Within each material the bank returns repeatedly to one buyer. The largest buyer of PET bales took 39 and 46 percent of the value in the two years and bought in every month, as did the largest buyers of paper and of other plastics. Repeated sales to a small set of buyers, rather than re-matching at each shipment, are the pattern documented in agricultural supply chains in developing countries, where buyer and seller trade under relational contracts that substitute for enforcement (Macchiavello and Morjaria, 2015, 2021), and they are consistent with the abstraction that the bank takes the downstream price as given rather than bidding it out shipment by shipment.

[26]Incomplete pass-through in levels is the evidence of intermediary market power in Bergquist and Dinerstein (2020) and Atkin and Donaldson (2015). The log coefficient on its own does not identify market power, since pass-through depends on the curvature of supply as well as on market structure (Weyl and Fabinger, 2013). A competitive buyer whose handling cost were proportional to the price would also show a log coefficient of one, so the ledger is consistent with a constant markdown and inconsistent with competition under a per-kilogram cost rather than a proof of the former.

These results carry the limitations of a single operator's records. The spreads are gross of handling, baling, storage, and transport, which the ledger does not record, so they bound the markdown rather than measure it. In 2024 the bank bought 366 tons of PET and sold 210, so its unit values for that year partly reflect inventory carried into 2025, when purchases and sales balance within 13 percent for every material, and the pooled spreads use both years. Because buyers are anonymized within each annual ledger, the persistence of buyer relationships across years cannot be measured.

5 Policy Implications

This section asks what the results imply for public policy. I consider two instruments available to Indonesian policymakers, a per-kilogram price subsidy paid through the banks and a payment targeted at pickers. Section 5.1 values the marginal kilogram by who delivers it and derives the planner's subsidy. Section 5.2 prices induced collection under a uniform and a targeted payment and asks who receives the subsidy. Section 5.3 asks what a recovered kilogram is worth under each of its alternative fates and derives break-even conditions.

5.1 Environmental Value and The Planner's Problem

Setup. A bank posts one price p per kilogram of a grade and resells at a net downstream value v. Because the posted price does not depend on who delivers, residual supply is the sum of household and picker supply, $Q(p) = H(p) + W(p)$, with elasticity

$$\varepsilon = s_H \, \varepsilon_H + s_W \, \varepsilon_W, \tag{12}$$

where s_H and $s_W = 1 - s_H$ are the two groups' shares of the bank's kilograms. Household supply responds to the price only through the pecuniary part of a household's return, which surveys identify as secondary to its civic part (Widayat et al., 2025), so ε_H is small. Picker supply responds on two margins, the loads a picker redirects from other buyers and the additional material she collects,

$$\varepsilon_W = \lambda(1 - \sigma) + \sigma \, \varepsilon_C, \tag{13}$$

where σ is the bank's share of picker collection in its catchment, λ governs how readily loads move between buyers, and ε_C is the elasticity of total collection to the return to collecting. A bank that maximizes $(v - p)\, Q(p)$ sets $p^* = \varepsilon v/(1 + \varepsilon)$, so the markdown below downstream

value is $1/\varepsilon$ and the posted price moves one for one in proportional terms with the value of the material (Appendix A derives these expressions and the menu that lets a bank set group-specific markdowns without pricing by identity). Let $\varepsilon^S = d\ln Q/d\ln p$ denote the supply elasticity estimated in Section 4, the empirical counterpart of equation (12), and let

$$\omega_\theta = \frac{s_\theta\, \varepsilon_\theta}{\varepsilon} \tag{14}$$

be the share of a price-induced increase in deliveries attributable to type $\theta \in \{H, W\}$. Because $\omega_H \approx 0$, the marginal kilogram a price instrument buys is a picker's kilogram whatever the composition of the bank's throughput. Its environmental value depends on where it would otherwise have gone.

Counterfactual fates. Let d_k denote the damage avoided when a kilogram that would otherwise have remained at source k, or followed the fate that source implies, is recovered instead. Material that would have been landfilled has avoided damage d_L, material that would have been burned in the open, dumped, or leaked into drains and waterways has $d_U \gg d_L$ (Table 17), and material another buyer would otherwise have received has an avoided damage of zero. The two types differ in their counterfactuals because they differ in why they deliver. A household that sorts and carries its waste for civic reasons is unlikely to be one that would otherwise burn or dump it, so if it stops using the bank its material most plausibly enters the municipal stream and is landfilled. Let m_H denote the share of material a household withdraws from the bank that would be landfilled, with the remainder burned, dumped, or sold to an itinerant buyer. I maintain that m_H is close to one. This is an assumption about selection into waste-bank use, not about the average Indonesian household, more than half of which report burning their waste (Kementerian Kesehatan, 2023). Because the marginal kilogram a price instrument buys is a picker's, the results for the price instruments below depend little on it. The assumption matters for any policy that draws in additional households, such as opening a bank where none exists. If half of household material would otherwise have been burned rather than landfilled, e_H rises from roughly US$22 to roughly US$160 per ton at the lower end of the open-burning range. A picker's collection draws instead on sources $k \in \mathcal{K}$ that differ in what would have happened to the material had she not taken it. These include litter along roads and in drains and waterways, informal dumpsites, household and communal bins, transfer points, and the

landfill itself. The expected damage avoided by a marginal kilogram from each type is then

$$e_H = m_H\, d_L + (1 - m_H)\, \bar{d}_{H,\neg L}, \tag{15}$$

$$e_W = \underbrace{\frac{\sigma \varepsilon_C}{\varepsilon_W}}_{\delta_W} \underbrace{\sum_{k \in \mathcal{K}} \mu_k\, d_k}_{\bar{d}_W}, \tag{16}$$

where $\bar{d}_{H,\neg L}$ is the average avoided damage over the household's non-landfill alternatives, zero for material sold to an itinerant buyer, δ_W is the share of the picker response that is new collection rather than redirection, which follows from equation (13), and μ_k is the share of new picker collection drawn from source k. The source shares are those of the marginal kilogram. A picker who collects more at a higher return extends her effort to the sources that were not worth working at the lower one, which need not be those that supply most of her material on average.

The expected damage avoided by a marginal kilogram delivered to the bank is the response-weighted average

$$e = \omega_H\, e_H + \omega_W\, e_W, \tag{17}$$

and the corresponding diversion share, the fraction of the response not redirected from another buyer, is

$$\delta = \omega_H\, \delta_H + \omega_W\, \delta_W, \qquad \delta_H = 1 - (1 - m_H)\, r_H, \tag{18}$$

where r_H is the share of the household's non-landfill alternative accounted for by itinerant buyers. The two types are asymmetric on every dimension that matters for e. Household kilograms are diverted with near certainty but from the least damaging destination, so $e_H \approx d_L$. Picker kilograms are partly redirected, so $\delta_W < 1$, but the kilograms that are new may come from sources whose counterfactual is burning or leakage, so $\bar{d}_W$ can be far above d_L. Neither δ_W nor the source shares μ_k can be recovered from the transaction data, which record a deposit identically whatever its origin. Section 4.2 explains why the entry design bounds these objects only loosely, so I treat the diversion share as a parameter and vary it over a range below.[27] Under the selection assumption, then, $e_H \approx d_L$, while e_W can lie anywhere in $[0, d_U]$: zero if the picker response is redirection, near $\delta_W d_L$ if new collection is drawn from sources already bound for managed disposal, and near $\delta_W d_U$ if it is drawn from roadsides, drains, and waterways.

[27] An event study of satellite-detected open-dump extent around new banks, using the Java-wide classifier of Gordon, Papp and Shandal (2026), has no power at this scale. The median bank collects 0.4 tons a year and the 90th percentile 8 tons, whereas a hectare of open dump holds 5,000 to 10,000 tons, so the exercise is a null by construction and I do not report it.

The planner's problem. The social value of a marginal kilogram delivered by type θ is $v + e_\theta$, while the seller receives p. Welfare from collection is

$$\mathcal{W}(p) = \sum_{\theta \in \{H,W\}} \int_0^{Q_\theta(p)} \left[v + e_\theta - Q_\theta^{-1}(q) \right] dq,$$

and because a single price faces both types, the price that maximizes it satisfies $\sum_\theta Q'_\theta(p)\,(v + e_\theta - p) = 0$, or

$$p^\circ = v + e. \tag{19}$$

The efficient uniform price equals net downstream value plus the damage avoided by the *marginal* kilogram, weighted by the types' responses in equation (14). The laissez-faire gap decomposes into two wedges,

$$p^\circ - p^* = \underbrace{\frac{p^*}{\varepsilon}}_{\text{market power}} + \underbrace{e}_{\text{environmental}}. \tag{20}$$

A subsidy of τ rupiah per kilogram of plastic, of any grade, bought from an individual depositor and paid to the operator raises the net value of a kilogram to $v + \tau$ and the posted price to $p^*(\tau) = \rho\,(v + \tau)$ with $\rho = \varepsilon/(1 + \varepsilon)$. Setting $p^*(\tau) = p^\circ$ gives

$$\tau^* = \underbrace{\frac{v}{\varepsilon}}_{\text{corrects the markdown}} + \underbrace{\frac{e}{\rho}}_{\text{environmental value, grossed up}}, \tag{21}$$

which converges to the Pigouvian benchmark e as residual supply becomes perfectly elastic. At a bank that separates the types through the menu, pass-through is type-specific by equation (29). A subsidy raises the picker price at the rate $\rho_W = \varepsilon_W/(1 + \varepsilon_W)$ and the household price at the rate $\rho_H = \varepsilon_H/(1 + \varepsilon_H)$. The induced response is then $s_W \varepsilon_W \rho_W + s_H \varepsilon_H \rho_H$ in place of $\varepsilon\rho$, and the fiscal cost per induced kilogram is

$$\left.\frac{C}{\Delta Q}\right|_{\text{menu}} = \frac{p}{s_W \varepsilon_W \rho_W + s_H \varepsilon_H \rho_H}. \tag{22}$$

As $\varepsilon_H \to 0$ this converges to $p/(s_W \varepsilon_W \rho_W)$, which exceeds the cost of a payment restricted to pickers, $p/(\varepsilon_W \rho_W)$, only by the factor $1/s_W$, the share of the subsidy paid on household kilograms. The bank's markdown on each type is $1/\varepsilon_\theta$, the inverse-elasticity rule, so a profit-maximizing menu passes a uniform subsidy through to each type in proportion to that type's responsiveness. The bank has no interest in the planner's objective. It separates its sellers in order to retain more of the subsidy paid on kilograms that would arrive anyway, and in

doing so it directs the remainder to the kilograms that respond.

For a small subsidy, collected kilograms rise by $\Delta Q = \varepsilon^S \rho \, (\tau/p) \, Q$, of which

$$\Delta D = \delta \, \varepsilon^S \rho \, \frac{\tau}{p} \, Q \tag{23}$$

are not redirected from another buyer, and the damage avoided is $e \, \Delta Q$. Because the subsidy is paid on inframarginal kilograms as well, its fiscal cost is τQ to first order, and

$$\frac{C}{\Delta D} = \frac{p}{\delta \, \varepsilon^S \, \rho}, \qquad \frac{C}{e \, \Delta Q} = \frac{p}{e \, \varepsilon^S \, \rho}. \tag{24}$$

The first ratio is the fiscal cost per kilogram diverted. The second is the fiscal cost per rupiah of damage avoided, which exceeds one whenever the subsidy costs more than the damage it prevents. Because $\omega_H \approx 0$, the environmental return to a price instrument is governed by the picker margin whatever the share of households in a bank's throughput, while the payment on household kilograms, whose size is governed by s_H, does not change where those kilograms go.

5.2 The Cost of a Per-Kilogram Subsidy

The cost of a per-kilogram subsidy depends on three objects, the supply elasticity ε^S, pass-through ρ, and the share δ of the induced response that is new collection rather than material redirected from another buyer. I take ε^S from the estimates of Section 4, separately for all depositors and for each seller type, so that a uniform payment and one conditioned on depositor type can be priced side by side. Pass-through is not estimated and is set at the constant-elasticity benchmark $\varepsilon^S/(1 + \varepsilon^S)$ throughout. The diversion share δ is left as a parameter, since nothing in the data identifies it, and Section 5.3 asks what value it would need to take for the subsidy to pay for itself. The formulas are local, so the figures below describe a small subsidy. A large one would require integrating along the supply curve.

Cost per ton. Table 16 evaluates the first ratio in equation (24) at the mean bank–month unit value of plastic bought from individual depositors, Rp 1,787 per kilogram. A bank that posts one price to a pooled set of sellers passes the subsidy through at the rate implied by the aggregate elasticity it faces, 1.18 from Section 4, or 0.54 of each rupiah. The kilograms the higher price draws in are pickers' and households', which respond at their own elasticities of 1.97 and 0.64 and make up 82 and 18 percent of the network's plastic. At these values a uniform subsidy costs Rp 1.91 million (US$115) per ton of induced collection, and twice that

if half the response is redirected from other buyers.[28] A payment conditioned on depositor type, which the per-depositor ledger makes feasible, costs Rp 1.37 million (US$82) per ton at the picker elasticity and Rp 7.15 million (US$431) at the household elasticity. A bank that separates the types through its menu passes the subsidy through at 0.66 on the grades pickers deliver and 0.39 on the grades households deliver, so the same uniform subsidy costs Rp 1.60 million (US$96) per ton at such a bank (equation (22)). The bank's menu does most of the planner's targeting, since it raises the price on the kilograms that respond and retains what is paid on those that do not. These are point estimates from imprecise elasticities. The Anderson–Rubin sets for all depositors reach zero and admit no finite upper bound on cost. Only the picker sets exclude zero, and their lower endpoints of 0.80 (clustered) and 0.45 (wild) give upper bounds of Rp 5.03 million and Rp 12.80 million per ton on the targeted payment.

Who receives it. A per-kilogram subsidy is paid on every kilogram delivered, not only on the kilograms it induces. The share of spending that buys new collection is $\varepsilon^S \rho\, \tau / p$. For a subsidy of Rp 500 per kilogram at a pooling bank, this is about 26 percent, with the remaining 74 percent paid on the quantity that was delivered at the old price. The resource cost of an induced ton is therefore roughly the price, US$108, and the gap to the US$115 fiscal cost is a transfer to the sellers and operators who were already in the market. Who receives that transfer follows from the kilogram shares and the pass-through benchmark. Of a rupiah of subsidy, the operator retains $1 - \rho \approx 0.46$ and passes ρ to depositors in proportion to their kilograms, 0.44 to pickers and 0.10 to households.[29] The household share buys nothing, since household deliveries do not respond to price, but it is small because households supply few kilograms. The picker share goes to the sellers whose supply does respond, and to a population whose earnings sit below those of every occupation in the national labor force survey (Sasaki, Choi and Watanabe, 2022). A per-kilogram payment is therefore self-targeting in the sense of Besley and Coate (1992), since no one who is not poor will collect and carry sacks of waste for Rp 1,787 a kilogram. A planner who values income in pickers' hands at or above its fiscal cost reads the US$115 as an upper bound on

[28] The pass-through is taken from the estimated aggregate elasticity rather than re-derived at the network's composition, since it is the elasticity a bank observes when it sets one price for all its sellers. The kilogram-weighted average of the type elasticities, 1.73, implies a pass-through of 0.63 and a cost of US$98 per ton. The aggregate estimate is the response of the typical bank, whose plastic is more household-heavy than the network's kilograms, so the two are consistent.

[29] The shares are ρs_W, ρs_H, and $1-\rho$ with $\rho = 0.54$ and $s_W = 0.82$ (Table 5). The split between depositors and the operator rests on the constant-elasticity benchmark rather than on an estimate of pass-through. The split among depositors rests only on their kilogram shares. At a bank that separates the types through the menu, the shares are $\rho_W s_W = 0.54$ to pickers, $\rho_H s_H = 0.07$ to households, and 0.39 retained by the operator.

the social cost of induced collection.

Design. The operator's share is the only part of the incidence a design rule can move. Because a bank can offset a depositor-side payment by lowering its own posted price, the payment reaches depositors one for one only under a rule that fixes the bank's posted component over the policy horizon. Without such a rule, a payment made to depositors enters the bank's problem exactly as an operator-side subsidy does, since the bank chooses the total receipt and reoptimizes its own component, and the depositor's receipt rises only at the pass-through rate. Under the pass-through benchmark, that rule moves just under half of the subsidy from operators to depositors. Conditioning the payment on depositor type moves far less at the network mix, the roughly 10 percent of the payment that households retain at a pooling bank and 7 percent at a segmenting one, though more at household-heavy banks. Its value is mainly distributional. A payment that reaches only pickers is a cash transfer to a self-selected population of the very poor, and its only leakage is that household share.

5.3 What a Recovered Kilogram Is Worth

Whether any of these payments buys environmental value depends on e, and by equation (17) with $\omega_H \approx 0$, on $e_W = \delta_W \bar{d}_W$ alone, the share of the picker response that is new collection and the damage avoided per newly collected kilogram given where it would otherwise have been. This subsection prices each of the alternative fates in Table 17, assembles what is known about where pickers' marginal kilograms come from, and closes with the break-even conditions that follow.

Landfill. A ton that would have entered a sanitary landfill avoids mainly disposal, airspace, and local disamenity costs, since plastic generates little methane. The ministry costs sanitary disposal at Rp 250,000 (US$15) per ton, and the landfill externality literature adds US$5–10 for disamenity and airspace, giving d_L of US$20–25 per ton (Tempo, 2025*a*; Eshet, Ayalon and Shechter, 2006; Nahman, 2011). That benchmark assumes that plastic in an engineered landfill is inert, and most Indonesian disposal sites are not engineered landfills. Of 550 sites counted by the Ministry of Environment in 2025, 343 operate as open dumps and about 110 meet sanitary standards (Tempo, 2025*b*). The ministry has ordered dozens closed for discharging directly into rivers, and the sites serving the largest cities are at or past capacity. Piyungan closed at capacity in 2024, Sarimukti is near it, and Bantar Gebang holds some 40 meters of accumulated waste (Kumparan, 2024; The Jakarta Post, 2024; CNN Indonesia, 2022). Their failure modes are those of open dumps. There were 35 site fires in the 2023 dry

season, several burning for weeks and releasing the dioxins and particulates that open burning does (Liputan6, 2023; Katadata, 2026), and slope collapses killed 143 people at Leuwigajah in 2005 and seven, three of them pickers, at Bantar Gebang in March 2026 (Lavigne et al., 2014; Mongabay, 2026). A ton diverted from such a site avoids not only disposal cost but a share of the burning, leachate, and collapse risk that the site imposes, and light plastic is among the fractions most likely to be blown or washed off an unmanaged face. The household kilogram, whose counterfactual I place at landfill, is therefore not worthless at the margin. Its value lies between the sanitary benchmark and the open-burning range, depending on the site its district operates. I keep the sanitary figure as the lower bound in the calculations that follow.

Open burning and dumping. A ton burned in the open imposes health damages through particulate and black-carbon emissions and climate damages through combustion CO_2e. The mortality component starts from the estimate of roughly 270,000 deaths a year attributable to open waste burning worldwide over about one billion tons burned (Kodros et al., 2016; Wiedinmyer, Yokelson and Gullett, 2014), valued at an Indonesian value of a statistical life of US$0.84 million (Vital Strategies, 2024), which gives about US$230 per ton of mixed waste. Because plastic yields more particulate and black carbon per ton than mixed waste, I scale this by a factor of one to five (Pathak et al., 2023; Wang et al., 2023). Combustion of a ton of plastic releases 1.2–1.6 tCO_2e, worth US$60–300 at a social cost of carbon of US$51–190 (U.S. Environmental Protection Agency, 2023*b*). Together these give d_U of US$300–1,500 per ton.

Waterways and the ocean. A ton that leaks into drains, rivers, and eventually the sea has no per-ton damage estimate comparable to the open-burning range, so I price the one damage that has been causally identified, the effect of marine microplastic exposure on low birth weight. Du, Zhang and Zou (2024) match three million births within 200 kilometers of a coast in fifteen countries to satellite measurements of marine microplastic concentration in the nearest ocean grid cell and find that a doubling of in-utero exposure raises the incidence of low birth weight by 0.37 per 1,000 births (0.53 per log point). The proportional effects are larger in the developing-country samples drawn from Demographic and Health Surveys, about 8 percent of the baseline rate per log point, and aerosolization of particles from the sea surface is the principal pathway. The pathway matters here, because it means the exposure of Indonesian coastal residents depends on the concentration in Indonesian waters, which Indonesian leakage largely determines. To convert the estimate into a value per ton I assume that the microplastic concentration in a country's coastal waters is proportional

to the stock of plastic it has leaked, and that a ton leaked today persists in that stock as long as the average ton already there. Under those assumptions the lifetime damage of a marginal ton equals the annual damage per log point of concentration divided by the annual leakage flow.[30] Indonesia records roughly 4.5 million births a year, nearly all of them within 200 kilometers of a coast, and a low-birth-weight rate of 6 to 12 percent depending on the survey (United Nations, 2024; Kementerian Kesehatan, 2023; Ekoriano et al., 2025). One log point of concentration therefore accounts for 2,400 cases a year at the pooled coefficient and between 22,000 and 43,000 at the developing-country coefficient, depending on the baseline rate. The national action plan puts leakage at about 620,000 tons a year (SYSTEMIQ and the Indonesia National Plastic Action Partnership, 2020), within the 0.48–1.29 million tons estimated by Jambeck et al. (2015). Valuing a case at the human-capital estimate of Alderman and Behrman (2006), US$832 in 2006 dollars at a 3 percent discount rate, about US$1,300 today, gives US$5–90 per ton leaked, and US$2–120 across the leakage range. That valuation prices the infant mortality attached to low birth weight at forgone earnings. At the Indonesian VSL used for the open-burning row, each percentage point of excess infant mortality among low-birth-weight births adds roughly US$30–600 per ton, so a differential of five percentage points would place the marine figure at US$150–3,000 per ton, in and above the open-burning range. The infant-health channel is one of several. Beaumont et al. (2019) put the loss of marine ecosystem services at US$3,300–33,000 per ton of plastic in the ocean per year, a figure built from valuation studies rather than from an identified damage, and one that would dominate any health estimate if taken at face value. I therefore record the marine fate in Table 17 as US$5–90 per ton on the identified channel alone, and as comparable to or above the open-burning range once mortality is valued at the VSL or non-health damages are admitted. The maintained ordering $d_{\text{marine}} \geq d_U$ in the break-even calculations is a judgment rather than a measurement.

Avoided virgin production. Recycling may also displace virgin production, worth US$60–250 per ton at full displacement for 1.0–1.3 tCO_2e avoided per ton recycled (U.S. Environmental Protection Agency, 2023*a*; Association of Plastic Recyclers, 2020). A recycled ton need not displace a full ton of virgin output, and no empirical displacement rate exists for plastics. At the rates estimated for aluminum and copper, the figure is US$6–125 (Zink and

[30] If the stock is $S = FT$, with F the annual flow and T the residence time, a marginal ton raises log concentration by $1/S$ for T years, and undiscounted lifetime damage is $T \times$ (annual damage per log point)$/S =$ (annual damage per log point)$/F$. Discounting at 3 percent over a 30-year residence lowers the figure by about a third, and by more if microplastic formation lags macroplastic leakage by decades. The assumption that concentration is proportional to the domestic stock ignores material that drifts in from, or out to, other countries' waters, which for an archipelago astride major currents can go either way.

Geyer, 2017; Lowe et al., 2025). I report this benefit separately rather than folding it into e, since it is independent of the counterfactual fate and accrues only on material that would not otherwise have been recycled, and so not on kilograms redirected from another buyer.

Where pickers collect. The fates above fix the value of a picker's marginal kilogram once its source is known, and the transaction data cannot say where that kilogram would otherwise have been. A deposit is recorded identically whatever its origin, and the entry design does not resolve it. What is known about where Indonesian pickers work narrows the range. The only national decomposition, assembled for the government's plastic action plan, attributes roughly half of the plastic the informal sector collects, about 500,000 tons a year, to collection directly from residential areas, by itinerant buyers and pickers working streets, bins, and neighborhood collection points, and the other half, about 560,000 tons, to material recovered from the collected stream in transit to disposal sites and at the sites themselves (SYSTEMIQ and the Indonesia National Plastic Action Partnership, 2020). The second half was already on its way to a landfill. The first half is drawn from the stream before collection, and in a country in which 57 percent of households burn their waste and formal collection reaches a fifth of villages, its alternative fate is often burning, dumping, or a drain (Kementerian Kesehatan, 2023; Badan Pusat Statistik, 2022). Dumpsite pickers, the best-studied group, are also the most specialized. The pickers at Bantar Gebang recover on the order of 120 tons a day, or 1.6 percent of the roughly 7,000 tons that arrive (Sasaki et al., 2020), and a five-city survey finds mean collection of 44 kilograms per picker per day over eight-hour workdays (Kristanto, Kemala and Nandhita, 2022). Pickers working river banks and drains are documented on the Citarum and Ciliwung but not counted. No source reports the composition of a picker's load by origin, and none reports it for the marginal kilogram, which is the object that matters here. A picker who extends her effort at a higher price works the sources that were not worth working before, which are more likely to be dispersed litter than the face of a landfill.

Break-even. Setting $C/(e\,\Delta Q) = 1$ in equation (24) with $e = \delta_W \bar{d}_W$ gives the condition under which a payment pays for itself on avoided damage alone, $\delta_W \bar{d}_W \geq C/\Delta Q$, that is, US$115 per ton for a uniform subsidy at a pooling bank, US$96 at a segmenting bank, and US$82 for a payment targeted at pickers. Table 18 evaluates the condition against each fate in Table 17 and reports the share of the picker response that must be new collection for each payment to break even. If new picker collection would otherwise have been burned in the open, the uniform subsidy pays for itself when at least 38 percent of the picker response is new collection at the lower end of the range and 8 percent at the upper end. At a segmenting

bank the thresholds are 32 and 6 percent, and for the targeted payment 27 and 5 percent. In resource-cost terms the uniform thresholds are 36 and 7 percent. If new collection is drawn instead from bins, transfer points, or the landfill, so that $\bar{d}_W$ is close to the sanitary value of US$20–25, no diversion share makes either payment break even on avoided damage alone, though the open-dump caveat above raises that floor where the local site is unmanaged. For collection whose counterfactual is a drain or waterway, the answer depends on which marine valuation is admitted. On the identified infant-health channel alone neither payment breaks even at any diversion share. With low-birth-weight mortality valued at the VSL, the targeted payment breaks even if 55 percent of the picker response is new collection at the lower end of the range and 3 percent at the upper end. The uniform subsidy requires 77 and 4 percent, and the subsidy at a segmenting bank 64 and 3 percent. On the ecosystem-service valuation both payments break even at diversion shares of 5 percent or less. The environmental case for a price instrument therefore rests on an empirical question the data cannot resolve, namely where the marginal kilogram a picker collects would otherwise have gone. In future work, I plan to survey the pickers in the depositor registry on the sources of their loads and on the buyers they would otherwise sell to, which would measure the source shares μ_k and the diversion share δ_W directly and allow these conditions to be calibrated rather than bounded.

6 Conclusion

In much of the developing world, the recovery of recyclable material is organized by a private chain of intermediaries rather than by the state, and the prices those intermediaries post are the margin by which recovery can be raised. The impact of subsidizing recovery along this chain depends on how much market power the intermediaries hold over the people who sell to them, and on which sellers it binds. This paper measures that market power for one part of the chain using the transaction records of a national network of waste banks and the opening dates of its members, and asks what it implies for policy.

I find that a waste bank faces two kinds of sellers whose supply responds very differently to the price it posts. The supply of plastic from waste pickers has an elasticity of about two, while the supply from households is statistically indistinguishable from zero. The marginal kilogram a higher price buys is a picker's, and the markdown a bank earns falls on households. A bank cannot pay the two groups different prices for the same grade, but it does not need to. It can refine its menu, adding finer grades that carry a premium and that only a seller who sorts can claim, and pickers, who sort, take them up. When a rival opens nearby, an

incumbent loses about a third of its plastic, all of it from pickers, and responds by refining its menu two to three times as often as before while letting the posted prices of its existing grades drift down. Competition in this market raises the price available to the sellers who can leave and widens the markdown on those who cannot.

These results have important implications for recycling policy. Because households do not respond to price, a per-kilogram payment paid on every kilogram delivered buys new collection only on the picker margin, and a payment restricted to pickers buys a ton of induced collection for about 30 percent less than a uniform one, roughly US$82 against US$115 at the point estimates. A bank that separates pickers from households through its menu captures most of that saving without any targeting by the government. It passes a uniform subsidy through to the grades pickers deliver and retains what is paid on the grades households deliver, so the market power that marks down the inelastic seller is also what directs a subsidy to the elastic one. Most of either payment is inframarginal, though the part that reaches pickers is a transfer to some of the poorest workers in the economy. Whether the collection it buys is worth its cost on environmental grounds depends on where a picker's marginal kilogram would otherwise have gone. If it would have been burned in the open or washed into a waterway, a targeted payment pays for itself even if most of the picker response is material redirected from other buyers. If it would have reached a landfill, no payment does.

The mechanism I document is not specific to waste. Studies of monopsony typically ask how a buyer sets its wage. Yet many buyers post one wage for a given task and let sellers choose among tasks, and the menu of tasks gives them a second instrument. A buyer facing sellers with different outside options can attach a premium to a task that its elastic sellers are willing to perform and the others are not, and so pay the elastic sellers more without paying everyone more. The examples are familiar. Hospitals and call centers pay premia for night and weekend shifts, ride-hailing platforms pay premia for driving at peak hours, insurers pay physicians by procedure, and school districts pay teachers stipends to coach a team or supervise a club. The waste banks studied here show that a buyer reaches for this instrument when competition for its elastic sellers intensifies. Incumbents add finer grades when a rival opens, and the sellers who sort into them are those the rival would otherwise take. Measuring market power in such markets therefore requires following the menu of tasks as well as the wage.

References

Alderman, Harold, and Jere R. Behrman. 2006. "Reducing the Incidence of Low Birth Weight in Low-Income Countries Has Substantial Economic Benefits." *World Bank Research Observer*, 21(1): 25–48.

Anderson, Robert C., and Richard D. Spiegelman. 1977. "Tax Policy and Secondary Material Use." *Journal of Environmental Economics and Management*, 4(1): 68–82.

Anderson, T. W., and Herman Rubin. 1949. "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations." *Annals of Mathematical Statistics*, 20(1): 46–63.

Andrews, Isaiah, James H. Stock, and Liyang Sun. 2019. "Weak Instruments in Instrumental Variables Regression: Theory and Practice." *Annual Review of Economics*, 11: 727–753.

Association of Plastic Recyclers. 2020. "Virgin vs. Recycled Plastic Life Cycle Assessment: Energy Profile and Environmental Burdens." Association of Plastic Recyclers. Summary of Franklin Associates (2018), *Life Cycle Impacts for Postconsumer Recycled Resins: PET, HDPE, and PP*.

Atkin, David, and Dave Donaldson. 2015. "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs." National Bureau of Economic Research Working Paper 21439.

Badan Pusat Statistik. 2022. "Statistik Potensi Desa Indonesia 2021." Badan Pusat Statistik, Jakarta. Village census (Podes) 2021.

Barrozo, Marcos. 2024. "Market Power and Carbon Emissions in the Amazon." Harvard University Working Paper.

Beaumont, Nicola J., Margrethe Aanesen, Melanie C. Austen, Tobias Börger, James R. Clark, Matthew Cole, Tara Hooper, Penelope K. Lindeque, Christine Pascoe, and Kayleigh J. Wyles. 2019. "Global Ecological, Social and Economic Impacts of Marine Plastic." *Marine Pollution Bulletin*, 142: 189–195.

Bergquist, Lauren Falcao, and Michael Dinerstein. 2020. "Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya." *American Economic Review*, 110(12): 3705–3747.

Berkouwer, Susanna B., and Joshua T. Dean. 2022. "Credit, Attention, and Externalities in the Adoption of Energy Efficient Technologies by Low-Income Households." *American Economic Review*, 112(10): 3291–3330.

Besley, Timothy, and Stephen Coate. 1992. "Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs." *American Economic Review*, 82(1): 249–261.

Bils, Mark. 2009. "Do Higher Prices for New Goods Reflect Quality Growth or Inflation?" *Quarterly Journal of Economics*, 124(2): 637–675.

Blomberg, Jerry, and Patrik Söderholm. 2009. "The Economics of Secondary Aluminium Supply: An Econometric Analysis Based on European Data." *Resources, Conservation and Recycling*, 53(8): 455–463.

Broda, Christian, and David E. Weinstein. 2010. "Product Creation and Destruction: Evidence and Price Implications." *American Economic Review*, 100(3): 691–723.

Budiyarto, Arie, Beverley Clarke, and Kirstin Ross. 2025. "Overview of Waste Bank Application in Indonesian Regencies." *Waste Management & Research*, 43(3): 306–321.

Caldwell, Sydnee, and Emily Oehlsen. 2026. "Outside Options and Labor Supply: Evidence from the Gig Economy." National Bureau of Economic Research Working Paper 35626.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics*, 225(2): 200–230.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics*, 90(3): 414–427.

Card, David. 2022. "Who Set Your Wage?" *American Economic Review*, 112(4): 1075–1090.

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *Quarterly Journal of Economics*, 131(2): 633–686.

Casaburi, Lorenzo, and Tristan Reed. 2022. "Using Individual-Level Randomized Treatment to Learn about Market Structure." *American Economic Journal: Applied Economics*, 14(4): 58–90.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *Quarterly Journal of Economics*, 134(3): 1405–1454.

Chatterjee, Shoumitro. 2023. "Market Power and Spatial Competition in Rural India." *Quarterly Journal of Economics*, 138(3): 1713–1760.

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with Zeros? Some Problems and Solutions." *Quarterly Journal of Economics*, 139(2): 891–936.

CNN Indonesia. 2022. "Timbunan Sampah di Bantargebang 40 Meter, Setara Gedung 16 Lantai." *CNN Indonesia, 28 June 2022*, `https://www.cnnindonesia.com/nasional/20220628123659-20-814459/timbunan-sampah-di-bantargebang-40-meter-setara-gedung-16-lantai`.

Cook, Ed, Kremena Ionkova, Perinaz Bhada-Tata, Sonakshi Yadav, Frank van Woerden, et al. 2026. *What a Waste 3.0: Global Snapshot of Solid Waste Management toward Circularity until 2050.* Washington, DC:World Bank Group.

Cullen, Zoë B., and Bobak Pakzad-Hurson. 2023. "Equilibrium Effects of Pay Transparency." *Econometrica*, 91(3): 765–802.

Damanhuri, Enri, and Tri Padmi. 2012. "The Role of Informal Sector on Solid Waste Management – A Case Study of Indonesia." In *Waste Management – An Integrated Vision*. Rijeka:IntechOpen.

de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review*, 110(9): 2964–2996.

DellaVigna, Stefano, and Matthew Gentzkow. 2019. "Uniform Pricing in U.S. Retail Chains." *Quarterly Journal of Economics*, 134(4): 2011–2084.

Diewert, W. Erwin. 1976. "Exact and Superlative Index Numbers." *Journal of Econometrics*, 4(2): 115–145.

Domínguez-Iino, Tomás. 2026. "Efficiency and Redistribution in Environmental Policy: An Equilibrium Analysis of Agricultural Supply Chains." *Journal of Political Economy*, 134(4): 1075–1109.

Du, Xinming, Shan Zhang, and Eric Zou. 2024. "Marine Microplastics and Infant Health." National Bureau of Economic Research Working Paper 33094.

Edwards, Ron, and David Pearce. 1978. "The Effect of Prices on the Recycling of Waste Materials." *Resources Policy*, 4(4): 242–248.

Ekoriano, Mario, Anugerah Widiyanto, Muthmainnah Muthmainnah, Yuli Puspita Devi, Bambang Eko Cahyono, Izatun Nafsi, and Teguh Widodo. 2025. "Profile and Factors Associated with Low Birth Weight in Indonesia: A National Data Survey." *Rural and Remote Health*, 25(1): 9170.

Eshet, Tzipi, Ofira Ayalon, and Mordechai Shechter. 2006. "Valuation of Externalities of Selected Waste Management Alternatives: A Comparative Review and Analysis." *Resources, Conservation and Recycling*, 46(4): 335–364.

Fullerton, Don, and Ann Wolverton. 2000. "Two Generalizations of a Deposit-Refund System." *American Economic Review*, 90(2): 238–242.

Fullerton, Don, and Thomas C. Kinnaman. 1995. "Garbage, Recycling, and Illicit Burning or Dumping." *Journal of Environmental Economics and Management*, 29(1): 78–91.

Fullerton, Don, and Thomas C. Kinnaman. 1996. "Household Responses to Pricing Garbage by the Bag." *American Economic Review*, 86(4): 971–984.

Gagnon, Etienne, Benjamin R. Mandel, and Robert J. Vigfusson. 2014. "Missing Import Price Changes and Low Exchange Rate Pass-Through." *American Economic Journal: Macroeconomics*, 6(2): 156–206.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics*, 225(2): 254–277.

Gordon, Matthew, Anna Papp, and Monica Shandal. 2026. "Open Dumps and the Global Trade in Garbage." Working paper, July 2026.

Hazell, Jonathon, Christina Patterson, Heather Sarsons, and Bledi Taska. 2022. "National Wage Setting." National Bureau of Economic Research Working Paper 30623. Forthcoming, *American Economic Review*.

Iacovidou, Eleni, Spyridoula Gerassimidou, David C. Wilson, Jessika Richter, Susan Jobling, and Eddy Soedjono. 2025*a*. "System-Wide Assessment of Indonesia's Plastic Value Chain: Mapping Flows and Stakeholder Dynamics." *Journal of Cleaner Production*, 520: 146082.

Iacovidou, Eleni, Spyridoula Gerassimidou, Norman Ebner, Mike Webster, Adrian Whyle, Susan Jobling, Fabrizio Ceschin, and Muhammad Reza Cordova. 2025*b*. "System-Wide Analysis of the Plastics Value Chain in Indonesia: The Five Levels of Information." *Journal of Circular Economy*, 3(1): 1–49.

Jambeck, Jenna R., Roland Geyer, Chris Wilcox, Theodore R. Siegler, Miriam Perryman, Anthony Andrady, Ramani Narayan, and Kara Lavender Law. 2015. "Plastic Waste Inputs from Land into the Ocean." *Science*, 347(6223): 768–771.

Katadata. 2026. "Jejak Panjang Kebakaran TPA: Saat Gunung Sampah Menjadi Gunung Api Baru." *Katadata, 3 July 2026*, `https://katadata.co.id/ekonomi-hijau/ekonomi-sirkular/6a479cf2cbb5a/jejak-panjang-kebakaran-tpa-saat-gunung-sampah-menjadi-gunung-api-baru`.

Kaza, Silpa, Lisa C. Yao, Perinaz Bhada-Tata, and Frank van Woerden. 2018. *What a Waste 2.0: A Global Snapshot of Solid Waste Management to 2050.* Washington, DC:World Bank.

Kementerian Kesehatan. 2023. "Survei Kesehatan Indonesia 2023." Kementerian Kesehatan, Jakarta. Indonesian Health Survey; reports household waste-disposal method.

Kementerian Lingkungan Hidup. 2024. "Sistem Informasi Pengelolaan Sampah Nasional (SIPSN)." National Waste Management Information System; 2024 figures as reported by 317 of 514 districts and cities.

Kinnaman, Thomas C., and Don Fullerton. 2000. "Garbage and Recycling with Endogenous Local Policy." *Journal of Urban Economics*, 48(3): 419–442.

Kochhar, Rajat. 2024. "Does Market Power in Local Agricultural Markets Hinder Farmer Climate Change Adaptation?" Working paper.

Kodros, John K., Christine Wiedinmyer, Bonne Ford, Rachel Cucinotta, Ryan Gan, Sheryl Magzamen, and Jeffrey R. Pierce. 2016. "Global Burden of Mortalities Due to Chronic Exposure to Ambient $PM_{2.5}$ from Open Combustion of Domestic Waste." *Environmental Research Letters*, 11(12): 124022.

Kristanto, Gabriel Andari, Dini Kemala, and Paras A. C. Nandhita. 2022. "Challenges Confronting Waste Pickers in Indonesia: An On-Field Analysis." *Waste Management & Research*, 40(8): 1381–1389.

Kumparan. 2024. "Akhir Cerita TPA Piyungan: Overload sejak 2012, Dipaksa Beroperasi sampai 2024." *Kumparan*, https://kumparan.com/pandangan-jogja/akhir-cerita-tpa-piyungan-overload-sejak-2012-dipaksa-beroperasi-sampai-2024-22fl96OInN3.

Lau, Winnie W. Y., Yonathan Shiran, Richard M. Bailey, Ed Cook, Martin R. Stuchtey, Julia Koskella, Costas A. Velis, et al. 2020. "Evaluating Scenarios Toward Zero Plastic Pollution." *Science*, 369(6510): 1455–1461.

Lavigne, Franck, Patrick Wassmer, Christopher Gomez, Timothy A. Davies, Danang Sri Hadmoko, T. Yan W. M. Iskandarsyah, Jean-Christophe Gaillard, Monique Fort, Pauline Texier, Mathieu Boun Heng, and Indyo Pratomo. 2014. "The 21 February 2005, Catastrophic Waste Avalanche at Leuwigajah Dumpsite, Bandung, Indonesia." *Geoenvironmental Disasters*, 1: 10.

Lee, Seunghoon. 2026. "Does a Food Waste Tax Work? Evidence from Quasi-Experiments in South Korea." Working paper, University of Missouri.

Liputan6. 2023. "35 TPA Terbakar Sepanjang 2023, KLHK Usul TPA Dijadikan Objek Vital." *Liputan6*, https://www.liputan6.com/lifestyle/read/5493149/35-tpa-terbakar-sepanjang-2023-klhk-usul-tpa-dijadikan-objek-vital.

Lowe, Benjamin H., Andrea Genovese, David Font Vivanco, and Trevor Zink. 2025. "Revisiting Circular Economy Rebound: Market Dynamics, Policy Implications, and Future Research Directions." *Journal of Industrial Ecology*, 29(6): 1936–1945.

Macchiavello, Rocco, and Ameet Morjaria. 2015. "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports." *American Economic Review*, 105(9): 2911–2945.

Macchiavello, Rocco, and Ameet Morjaria. 2021. "Competition and Relational Contracts in the Rwanda Coffee Chain." *Quarterly Journal of Economics*, 136(2): 1089–1143.

MacKinnon, James G., Morten Ørregaard Nielsen, and Matthew D. Webb. 2023. "Cluster-Robust Inference: A Guide to Empirical Practice." *Journal of Econometrics*, 232(2): 272–299.

Manning, Alan. 2003. *Monopsony in Motion: Imperfect Competition in Labor Markets.* Princeton, NJ:Princeton University Press.

Mansikkasalo, Anna, Robert Lundmark, and Patrik Söderholm. 2014. "Market Behavior and Policy in the Recycled Paper Industry: A Critical Survey of Price Elasticity Research." *Forest Policy and Economics*, 38: 17–29.

Medina, Martin. 2008. "The Informal Recycling Sector in Developing Countries: Organizing Waste Pickers to Enhance Their Impact." Public-Private Infrastructure Advisory Facility (PPIAF), World Bank Gridlines Note 44, Washington, DC.

Meijer, Lourens J. J., Tim van Emmerik, Ruud van der Ent, Christian Schmidt, and Laurent Lebreton. 2021. "More Than 1000 Rivers Account for 80% of Global Riverine Plastic Emissions into the Ocean." *Science Advances*, 7(18): eaaz5803.

Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria. 2018. "Asymmetric Information and Middleman Margins: An Experiment with Indian Potato Farmers." *Review of Economics and Statistics*, 100(1): 1–13.

Mongabay. 2026. "Deadly Indonesia Landfill Collapse a 'Serious Warning' of Systemic Failure." *Mongabay News*, March 2026. Bantar Gebang waste-face collapse; seven fatalities.

Mussa, Michael, and Sherwin Rosen. 1978. "Monopoly and Product Quality." *Journal of Economic Theory*, 18(2): 301–317.

Nahman, Anton. 2011. "Pricing Landfill Externalities: Emissions and Disamenity Costs in Cape Town, South Africa." *Waste Management*, 31(9): 2046–2056.

Nakamura, Emi, and Jón Steinsson. 2012. "Lost in Transit: Product Replacement Bias and Pricing to Market." *American Economic Review*, 102(7): 3277–3316.

Palmer, Karen, and Margaret Walls. 1997. "Optimal Policies for Solid Waste Disposal Taxes, Subsidies, and Standards." *Journal of Public Economics*, 65(2): 193–205.

Pathak, Gauri, Mark Nichter, Anita Hardon, Eileen Moyer, Aarti Latkar, Joseph Simbaya, Diana Pakasi, Efenita Taqueban, and Jessica Love. 2023. "Plastic Pollution and the Open Burning of Plastic Wastes." *Global Environmental Change*, 80: 102648.

Quach, Simon. 2025. "Intra-Firm Competition and Uniform Wages." Working paper, University of Southern California.

Ramadan, A. H., E. Sembiring, B. Rahardyan, and H. Kardhana. 2026. "Mapping Mismanaged Plastic Waste in Indonesia: Subdistrict-Level Analysis through Material Flow from Sources to the Environment." *Scientific Reports*, 16: 13341.

Republik Indonesia. 2008. "Undang-Undang Nomor 18 Tahun 2008 tentang Pengelolaan Sampah." Law No. 18 of 2008 on Waste Management, enacted 7 May 2008.

Republik Indonesia. 2012. "Peraturan Pemerintah Nomor 81 Tahun 2012 tentang Pengelolaan Sampah Rumah Tangga dan Sampah Sejenis Sampah Rumah Tangga." Government Regulation No. 81 of 2012 on household waste management, enacted 15 October 2012.

Rubens, Michael. 2023. "Market Structure, Oligopsony Power, and Productivity." *American Economic Review*, 113(9): 2382–2410.

Sasaki, Shunsuke, Kohei Watanabe, Nurul Widyaningsih, and Tetsuya Araki. 2020. "Recycling Contributions of Dumpsite Waste Pickers in Bantar Gebang, Indonesia." *Journal of Material Cycles and Waste Management*, 22: 1662–1671.

Sasaki, Shunsuke, Yunhee Choi, and Kohei Watanabe. 2022. "Economic Status of Waste Pickers in Bantar Gebang Compared to Other Workers in Indonesia." *Habitat International*, 119: 102501.

Sharma, Garima. 2023. "Monopsony and Gender." Working paper, MIT.

Sokolova, Anna, and Todd Sorensen. 2021. "Monopsony in Labor Markets: A Meta-Analysis." *ILR Review*, 74(1): 27–55.

SYSTEMIQ and the Indonesia National Plastic Action Partnership. 2020. "Radically Reducing Plastic Pollution in Indonesia: A Multistakeholder Action Plan." SYSTEMIQ. `https://www.systemiq.earth/wp-content/uploads/2023/02/NPAP-Indonesia-Multistakeholder-Action-Plan_April-2020.pdf`.

Tempo. 2025*a*. "Hentikan Open Dumping, KLH: TPA Sistem Sanitary Landfill Perlu Biaya Besar dan Lahan Baru." *Tempo.co*, 10 March 2025. Reports the Rp 250,000 per ton conversion cost and the 0.1–0.6 percent APBD allocation figures.

Tempo. 2025*b*. "Menteri LH: 343 dari 550 TPA di Indonesia Masih Open Dumping." *Tempo.co*.

The Jakarta Post. 2024. "Bandung Landfill on the Verge of Overcapacity." *The Jakarta Post, 9 October 2024*, `https://www.thejakartapost.com/indonesia/2024/10/09/bandung-landfill-on-the-verge-of-overcapacity.html`.

United Nations. 2024. "World Population Prospects 2024." United Nations, New York.

U.S. Environmental Protection Agency. 2023*a*. "Documentation for Greenhouse Gas Emission and Energy Factors Used in the Waste Reduction Model (WARM), Version 16: Containers, Packaging, and Non-Durable Goods Materials." U.S. Environmental Protection Agency EPA-530-R-23-021.

U.S. Environmental Protection Agency. 2023*b*. "Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances." U.S. Environmental Protection Agency EPA-430-R-23-001.

Vidyaningrum, Widita. 2020. "Study of Solid Waste Management Financing in Indonesia: Final Report." Strategic Sector Cooperation, Indonesia–Denmark, Jakarta.

Vital Strategies. 2024. "Cost-Benefit Analysis for Air Pollution Control Strategies in Jakarta." Vital Strategies, New York.

Wang, Xiaoliang, Hatef Firouzkouhi, Judith C. Chow, John G. Watson, Steven Sai Hang Ho, Warren Carter, and Alexandra S. M. De Vos. 2023. "Chemically Speciated Air Pollutant Emissions from Open Burning of Household Solid Waste from South Africa." *Atmospheric Chemistry and Physics*, 23: 15375–15393.

Weyl, E. Glen, and Michal Fabinger. 2013. "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy*, 121(3): 528–583.

Widayat, Prama, Marnis, Afred Suci, Samsir, and Jahrizal. 2025. "It Is Not About Money: Exploring Motives for Public Participation in Indonesian Waste Banks." *VOLUNTAS: International Journal of Voluntary and Nonprofit Organizations*, 36(6): 880–893.

Wiedinmyer, Christine, Robert J. Yokelson, and Brian K. Gullett. 2014. "Global Emissions of Trace Gases, Particulate Matter, and Hazardous Air Pollutants from Open Burning of Domestic Waste." *Environmental Science & Technology*, 48(16): 9523–9530.

World Bank. 2018. "Indonesia Marine Debris Hotspot Rapid Assessment: Synthesis Report." World Bank, Washington, DC.

Zahrah, Yasmin, Jie Yu, and Xin Liu. 2024. "How Indonesia's Cities Are Grappling with Plastic Waste: An Integrated Approach towards Sustainable Plastic Waste Management." *Sustainability*, 16(10): 3921.

Zavala, Lucas. 2022. "Unfair Trade? Monopsony Power in Agricultural Value Chains." Working paper.

Zink, Trevor, and Roland Geyer. 2017. "Circular Economy Rebound." *Journal of Industrial Ecology*, 21(3): 593–602.

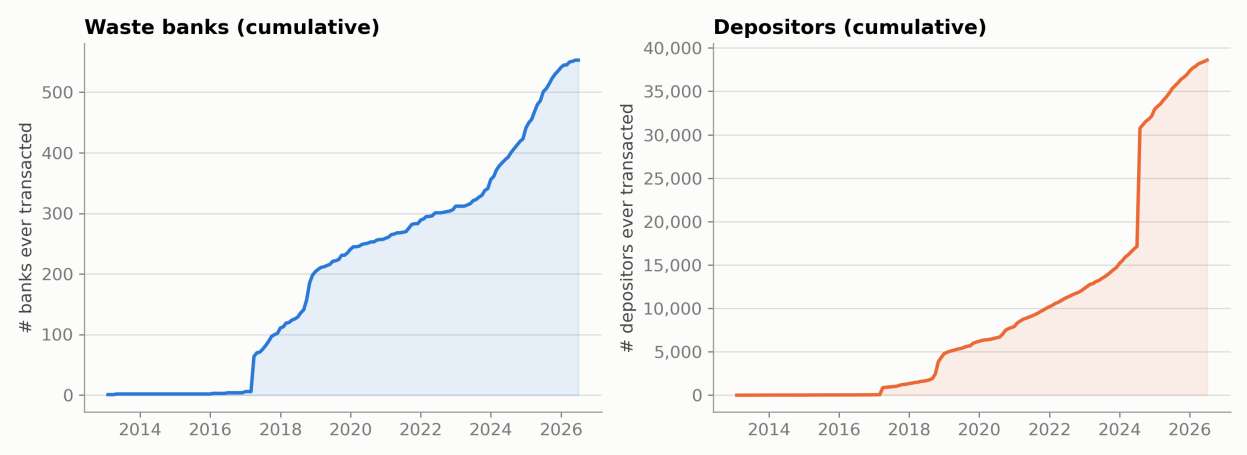
Zink, Trevor, Roland Geyer, and Richard Startz. 2018. "Toward Estimating Displaced Primary Production from Recycling: A Case Study of U.S. Aluminum." *Journal of Industrial Ecology*, 22(2): 314–326.

Figures and Tables

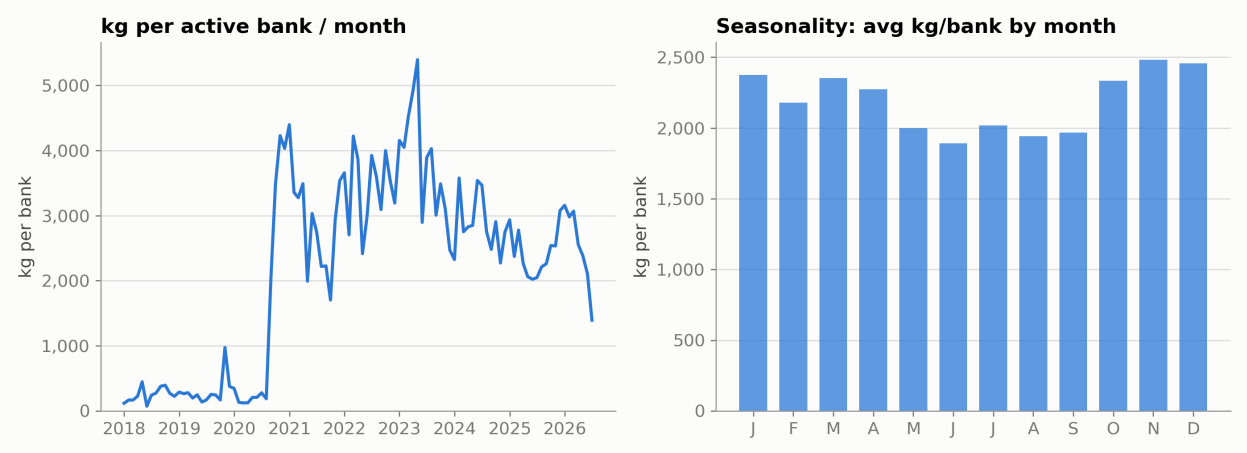
Figures

Figure 1: The network over time: banks, depositors, and volume

Panel A. Cumulative banks and depositors



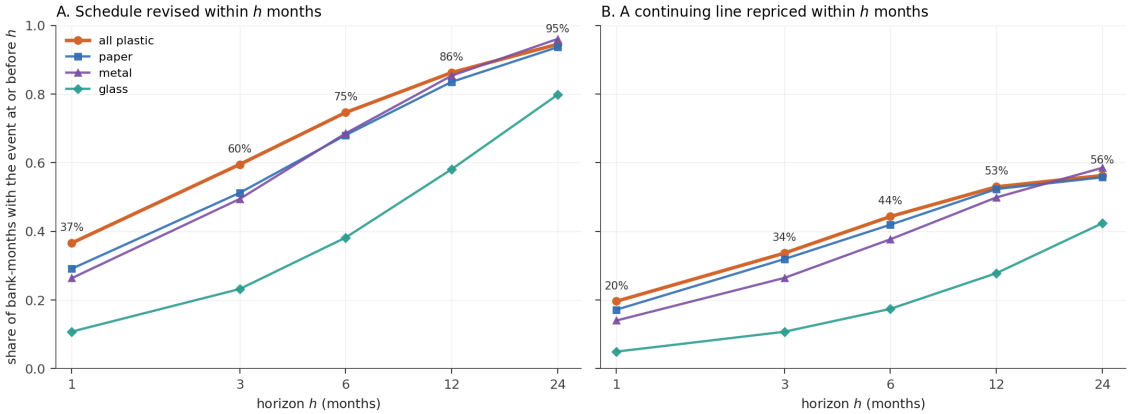
Panel B. Kilograms per active bank



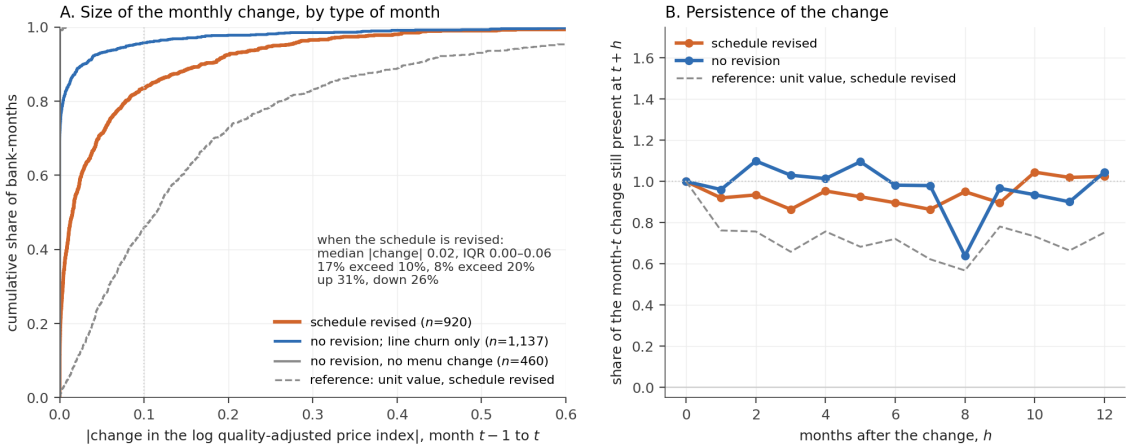
Notes: Panel A plots the cumulative number of distinct banks (left) and depositors (right) that have ever recorded a transaction, by month, February 2013–July 2026 (the partial final month of the extract, August 2026, is excluded). Panel B plots kilograms collected per active bank per month (left), where a bank is active in a month if it records at least one transaction, and the same statistic averaged by calendar month (right), January 2018–July 2026. Panel B excludes one bank, Bank Sampah Langgeng Lumintu, which recorded approximately 1.57 million kilograms of undifferentiated "Residu" at a nominal price of 1 Rp/kg between January 2023 and September 2024; including it produces a spike to about 9,400 kilograms per active bank in May 2024 that reflects this one account rather than the network. I read these entries as a bulk municipal hand-off logged for volume-tracking purposes rather than household transactions priced on market terms, and retain the bank in the estimation sample (see Section 3).

Figure 2: Schedule revisions: how often they occur, and how large and persistent they are

Panel A. Probability that a bank revises its schedule within h months, by material

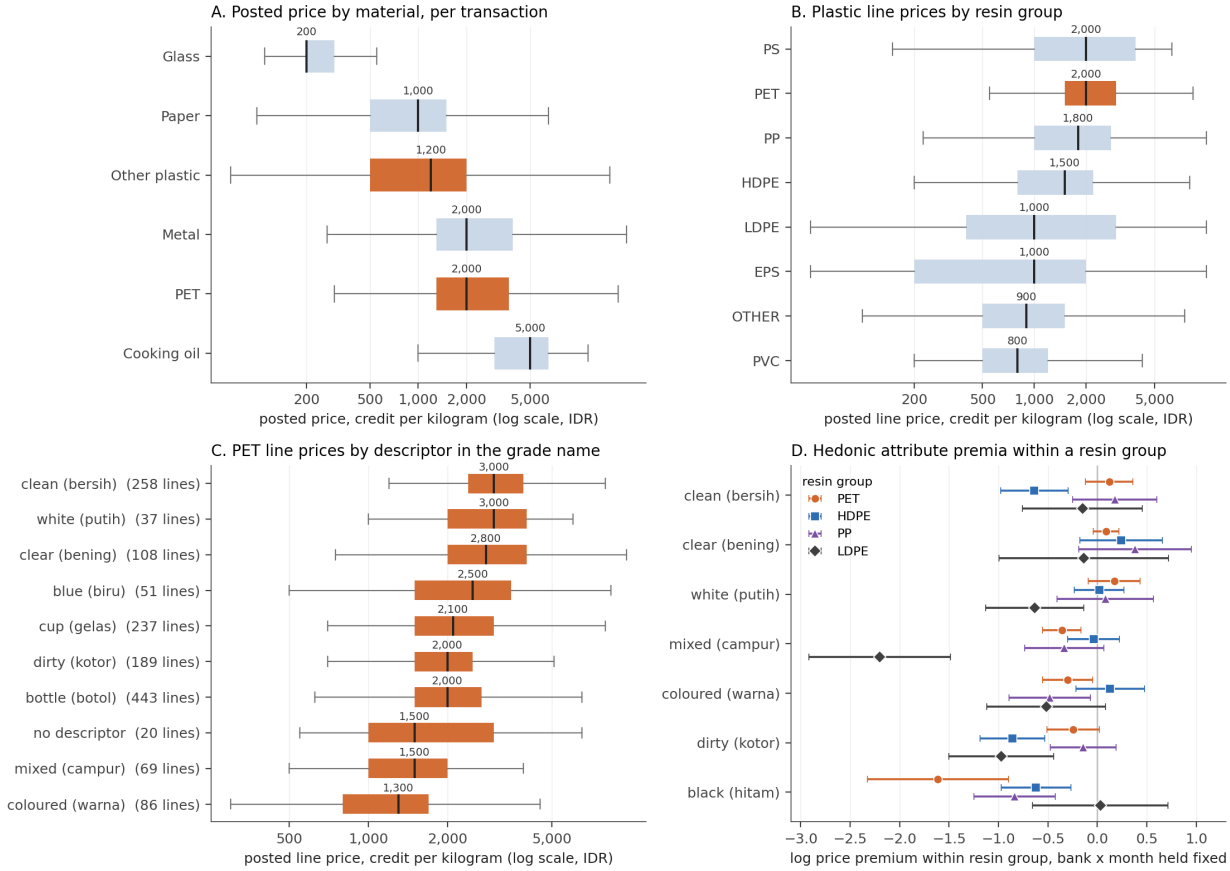


Panel B. Size and persistence of the change in the quality-adjusted price when the schedule is revised



Notes: Panel A: the unit is a bank–month window in which every month from $t-h$ to t is observed. A bank's schedule is revised in a month if it changes the price on a continuing line, transacts a line for the first time that goes on to be transacted in at least three months, or transacts a line for the last time while continuing to record for at least six further months. The left plot shows the share of windows with at least one revision at or before h; the right plot the share with at least one continuing line repriced. Paper, metal and glass are built by the same rule on the same banks' non-plastic lines. At $h = 1$, 2,511 bank–months from 238 banks for plastic; at $h = 24$, 638 bank–months from 48 banks. Panel B, left: cumulative distribution of the absolute change in the log quality-adjusted plastic price index between adjacent observed months, by whether the schedule is revised, only the set of transacting lines changes (churn), or neither; the dashed gray curve is the all-plastic unit value in revision months. Right: slope from a regression of the cumulative change from $t-1$ to $t+h$ on the initial change from $t-1$ to t, among months with an initial change above 0.01 in absolute value, so that one means the initial change is fully persistent; 2,517 bank–months from 241 banks. Sample: January 2018–July 2026.

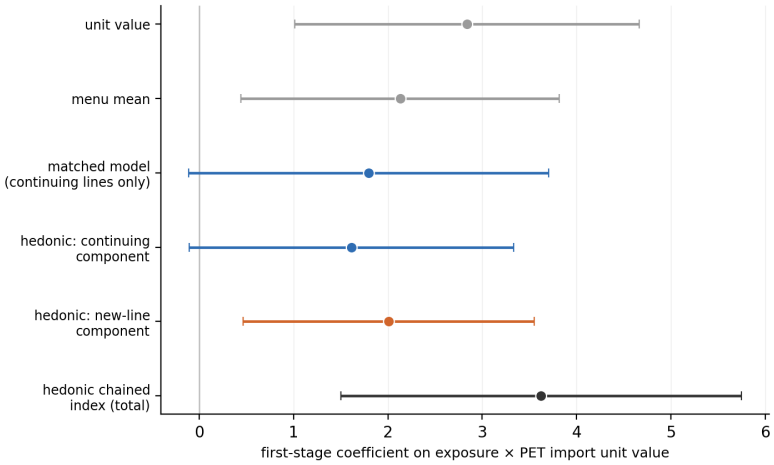
Figure 3: Price levels by material and by plastic grade



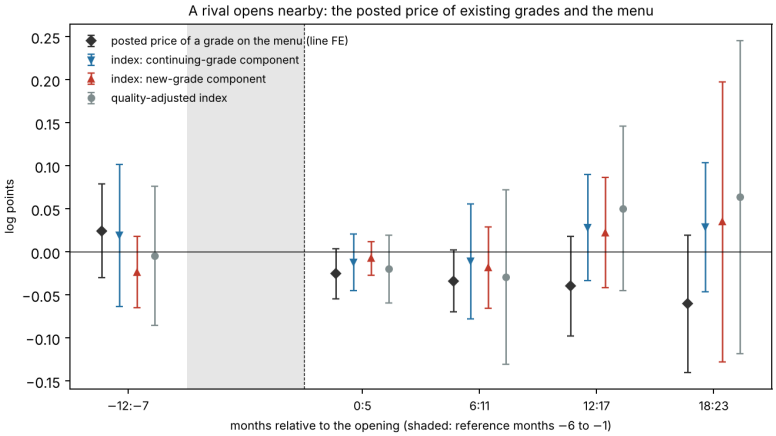
Notes: Panel A: distribution of the per-transaction price paid, by material, on a log scale; boxes span the interquartile range, whiskers the 2nd–98th percentiles, and the line marks the median, which is printed above each box. Panel B: the same for plastic line–month posted prices by resin group, where a line is a bank-specific posted category and its monthly price is the modal credit per kilogram; Table 2 reports the number of distinct lines in each group. Panel C: PET line–months by the descriptor carried in the grade name (a line can carry several), with the number of distinct lines in parentheses. Panel D: the attribute premia of Appendix Table A1 as log-point coefficients with 95 percent intervals: the attribute's main effect plus its interaction with the resin group, from the hedonic model with bank × month effects and credit weights, standard errors clustered by line. Sample for Panels A–C: January 2018–July 2026.

Figure 4: Repricing a continuing grade versus changing the menu, under two shocks

Panel A. An upstream value shock

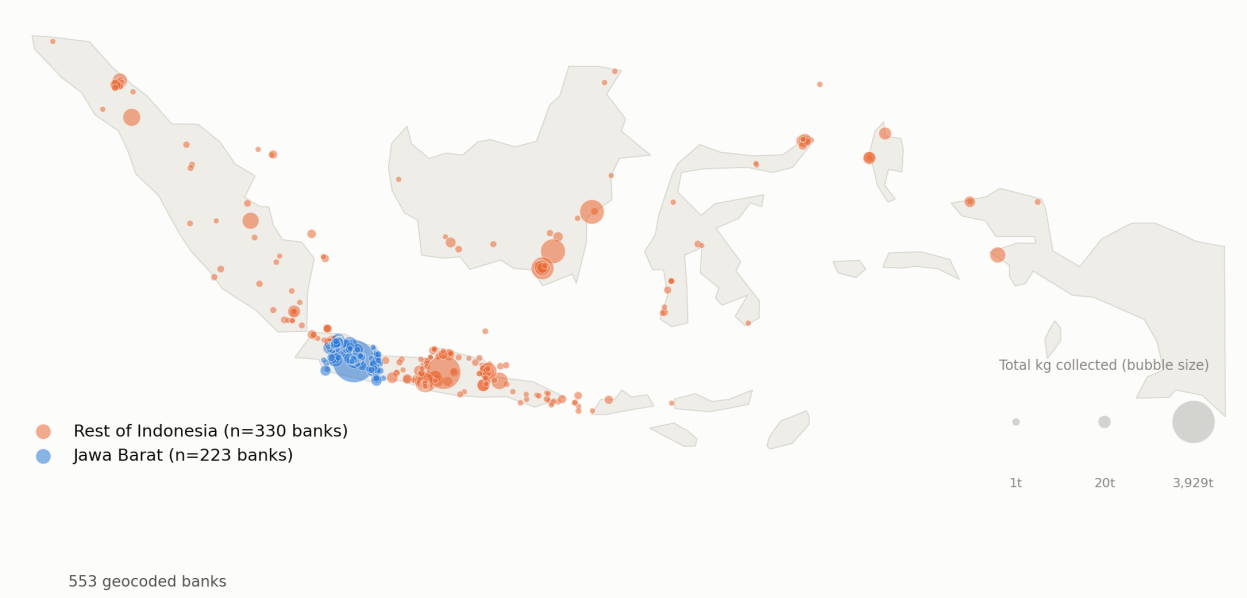


Panel B. A rival opens nearby



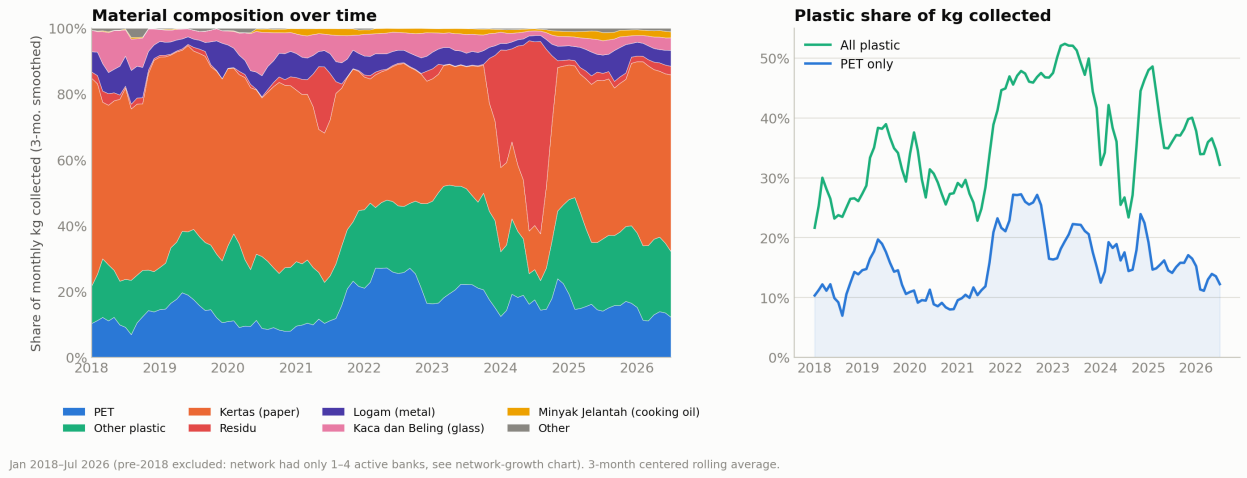
Notes: Panel A orders six measures of a bank's plastic price by how much menu change each admits and reports the coefficient on baseline plastic exposure interacted with the real-rupiah PET import unit value, in a bank–month regression with bank and calendar-month fixed effects on the 578 bank–months and 21 banks of Table 5; bars are 95 percent confidence intervals from standard errors two-way clustered by bank and month. Gray marks measures that mix the posted price with delivered composition, blue those that hold the line fixed, orange the component that prices lines new to the menu, and black the chained index that combines the two, under the six-month reappearance rule of Section 3.2. Panel B reports the entry design of Table 11 over the same margins, pooled over all treated incumbents: the posted price of a plastic grade on the menu (stacked line fixed-effects regression) against the chained index and its continuing- and new-grade components (stacked OLS on the bank–month index); bars are 95 percent confidence intervals clustered by the entrant. Across bank–months the standard deviation of the quality-adjusted level is 0.470 log points against 0.479 for the unit value, so the adjustment changes the dynamics of the price rather than its cross-section.

Figure 5: Geographic distribution of waste banks



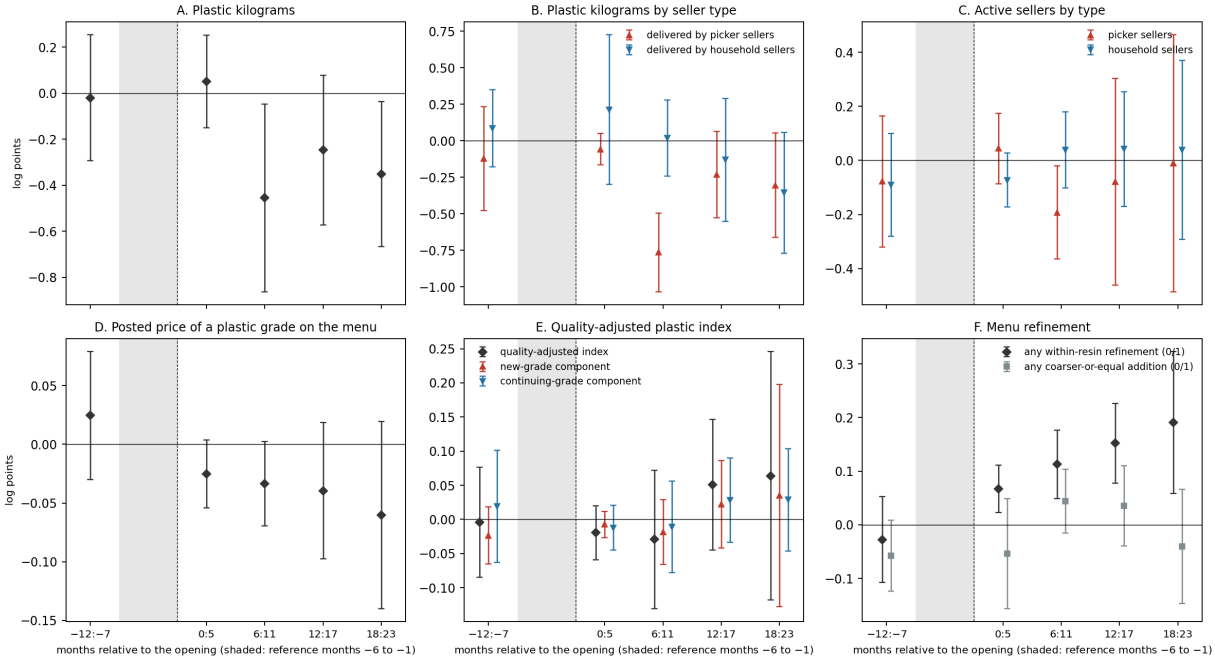
Notes: Each point is one waste bank, located at the median geocoded coordinate of its transactions; marker area is proportional to the bank's total kilograms collected over the sample. All 553 banks in the analysis sample are plotted. Banks in Jawa Barat province are shown in blue; all other provinces are shown in orange.

Figure 6: Material composition of collected waste over time



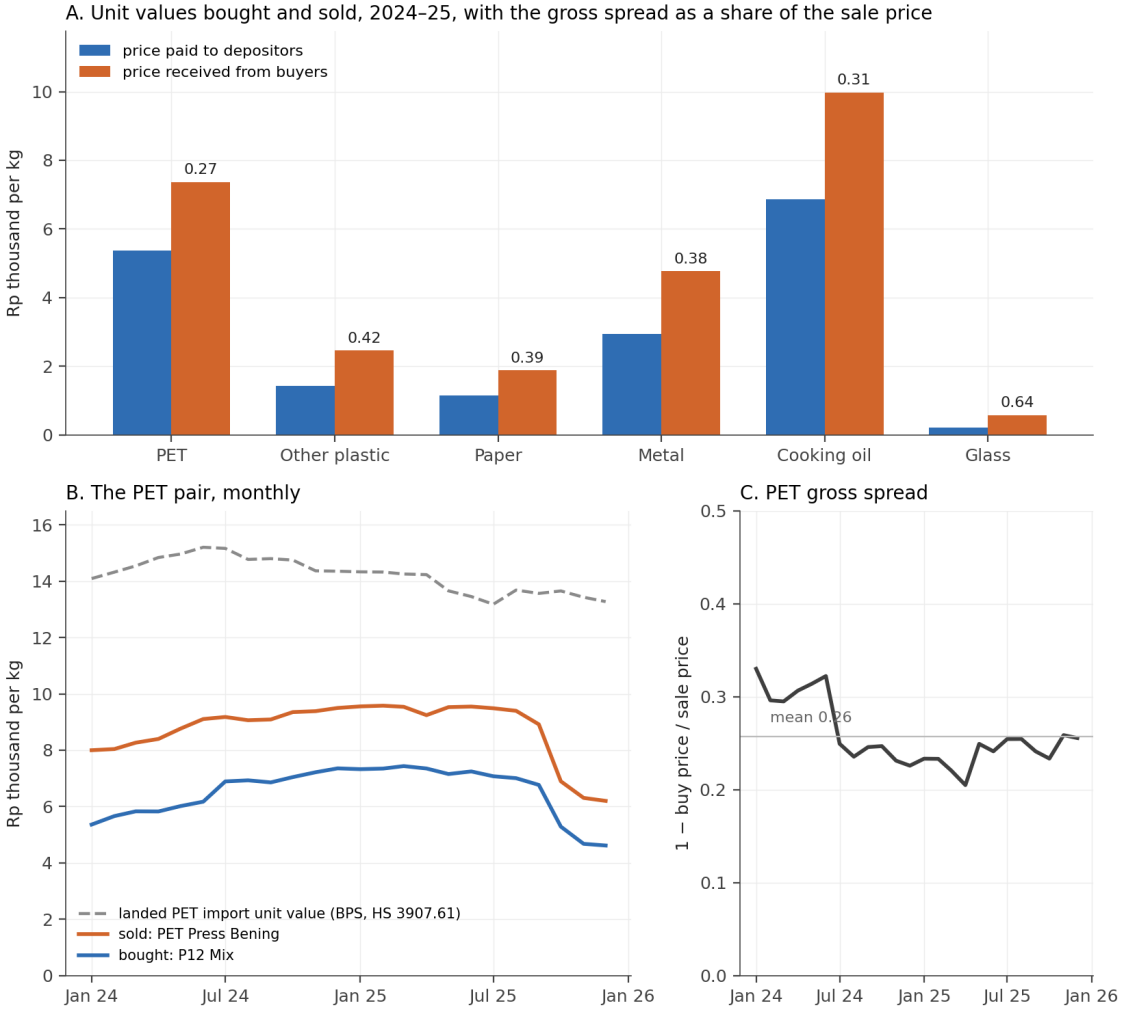
Notes: Left panel: share of monthly kilograms collected by material group, 3-month centered rolling average. Right panel: combined plastic (PET plus other resins) and PET-only share of monthly kilograms, same smoothing. Sample restricted to January 2018 onward; pre-2018 months are excluded because fewer than fifteen banks were active in a typical month before 2018, and one month, April 2017, contains a bulk onboarding of about fifty banks, so a monthly composition share is unrepresentative (see Section 3).

Figure 7: Incumbent responses to the opening of a rival waste bank



Notes: Design as in Tables 9 and 11, pooled over all 97 treated incumbents. Panels A–C are stacked Poisson estimates on individual-depositor quantities over recorded months; Panels D and E are stacked OLS on the log posted price of a plastic grade on the menu and on the bank–month quality-adjusted index and its components, the index following the six-month reappearance rule of Section 3.2; Panel F is a stacked linear probability model on recorded bank–months for a within-resin refinement and for a coarser-or-equal addition in the month (Table 11), in probability units. The shaded band marks the reference period, months −6 to −1. Bars are 95 percent confidence intervals, clustered by the entrant; the number of incumbents behind a bin falls with the horizon (Tables 9 and 11).

Figure 8: Onward sales at one large operator: unit values, the PET pair, and the gross spread



Notes: The operator and the matching are those of Table 15. Panel A: the price the bank paid its depositors and the price it received from buyers for the same material, value over kilograms across January 2024 to December 2025; the number above each pair is the gross spread as a share of the price received. Panel B: the bank's largest sale product (pressed clear PET bales, "PET Press Bening") and its largest purchase grade (loose cap- and label-free PET bottles of mixed colour, "P12 Mix"), monthly unit values, together with the unit value of imported PET resin (BPS, HS 3907.61, converted at the monthly exchange rate), the upstream series of Section 4.1. Panel C: one minus the ratio of the two series in Panel B.

Tables

Table 1: Summary statistics

Panel A. Sample overview

Transactions	Waste banks	Depositors	Provinces	Sample period	Total kg collected
757,559	553	38,597	30	Feb 2013–Aug 2026	9,581,876

Panel B. Transaction-level variables

Variable	N	Mean	Median	SD	Min	P99	Max
Weight (kg/tx)	757,559	12.6	2.0	253.8	0.01	148.9	76,000
Posted price (Rp/kg)	757,559	1,813.2	1,300.0	2,652.4	0.20	10,000.0	500,000
Credit (Rp/tx)	757,559	20,166.3	2,400.0	166,209.2	0.04	254,988.4	32,400,000

Panel C. Composition by material

Material	Transactions	Total kg	% of kg	Unit value (Rp/kg)	Median price (Rp/kg)
Kertas (paper)	236,032	3,584,406	37.4%	1,216	1,000
Other plastic	223,010	1,956,029	20.4%	1,596	1,200
Residu	15,250	1,707,037	17.8%	128	2,500
PET	145,541	1,419,206	14.8%	4,052	2,000
Logam (metal)	75,577	399,130	4.2%	2,509	2,000
Kaca dan Beling (glass)	38,871	318,941	3.3%	305	200
Minyak Jelantah (cooking oil)	16,992	115,756	1.2%	5,463	5,000
Organik (organic)	2,139	63,475	0.7%	365	100
E-waste	2,707	12,187	0.1%	5,415	7,000
Other materials	1,440	5,709	0.1%	1,028	350

Notes: The sample is every transaction surviving the filters of Section 3: parseable timestamp, positive recorded weight, and positive credit. Panel B reports the distribution of the three continuous transaction-level variables; all three are heavily right-skewed, and Min and Max are the sample extremes rather than a trimmed range. The posted price is credit divided by recorded weight. Panel C sorts materials by total kilograms collected; "Other plastic" pools all plastic resin categories other than PET, "E-waste" pools a family of subcategories (batteries, televisions, phones, and similar) recorded separately in the source taxonomy, and "Other materials" pools rubber and leather, textiles, and garden waste. The unit value is total credit divided by total kilograms; the median price is the median of credit per kilogram across transactions. The two differ most for Residu, where a few high-weight, low-value hand-offs dominate the kilograms. Percent-of-kg shares are rounded to one decimal and may not sum to exactly 100.

Table 2: Plastic grades and menus, by resin group

	PET	HDPE	PP	LDPE	PVC	PS	EPS	OTHER	All
Lines (bank × grade)	1,326	609	417	169	144	99	25	895	3,677
Line–months	9,752	4,570	3,568	1,263	955	480	362	4,974	25,924
Banks posting the group	344	200	146	80	82	54	12	200	395
Share of plastic credit	64.9%	7.0%	9.0%	5.6%	0.2%	0.2%	0.4%	12.6%	100.0%
Posted price, Rp/kg: median	2,000	1,500	1,800	1,000	800	2,000	1,000	900	1,500
25th percentile	1,500	800	1,000	400	500	1,000	200	500	800
75th percentile	3,000	2,200	2,800	3,000	1,200	3,900	2,000	1,500	2,500
Cross-bank SD of log line price, within month	0.69	0.86	0.82	1.15	0.71	1.01	1.01	0.91	0.81
Lines on the menu per bank–month, median	2	2	2	1	1	1	1	2	5
Months a line is observed, median	3	4	4	3	2	2	7	2	3
Lines observed in one month only	31%	26%	29%	26%	33%	34%	32%	33%	31%
Continuing line–months repriced	9%	6%	8%	9%	8%	8%	2%	5%	7%
Line–months on a Rp 500 multiple	59%	49%	47%	44%	36%	60%	40%	44%	51%
Within-R^2 of the hedonic model	0.83	0.83	0.88	0.90	0.79	0.97	0.25	0.27	0.82

Notes: A line is a bank-specific posted category (bank × `category_id`); its price in a month is the modal credit per kilogram over its transactions. Resin groups follow the platform's category groups; OTHER is the residual "Plastik Lainnya" group and EPS is expanded polystyrene. The cross-bank standard deviation is of the log line price demeaned by resin group and month, over months with at least five banks posting the group. A continuing line–month is one whose line was also posted in the previous calendar month; it is repriced if its log price moves by more than 0.01. The hedonic within-R^2 is from the attribute model of Section 3.2 re-estimated group by group (Appendix Table A1). Sample: January 2018–July 2026, 25,924 line–months from 3,677 lines and 395 banks.

Table 3: Two kinds of depositors: scale, composition, and plastic variety

	Household	Picker	
Panel A. Scale and composition (depositor–bank pairs)			
Depositor–bank pairs	25,088	3,804	
Kilograms per visit, median	3.0	41.0	
mean	8.5	102.2	
Plastic kilograms per visit, median	1.6	17.2	
Materials per visit	1.9	2.4	
Share of all individual-depositor kilograms	0.29	0.64	
Composition of the group's kilograms: plastic	0.26	0.60	
paper	0.57	0.31	
metal	0.07	0.03	
glass	0.05	0.03	
	Household	Picker	Within-bank gap
Panel B. Variety of plastic			
Plastic grades per visit	1.54	2.81	0.95*** (0.08)
Resin groups per visit	1.34	1.96	0.43*** (0.03)
Share of visits with two or more grades	0.27	0.63	0.20*** (0.03)
Distinct grades ever delivered	3.03	5.00	1.01*** (0.14)
Distinct resin groups ever delivered	1.83	2.51	0.31*** (0.04)
	Household	Picker	Picker at treated incumbents
Panel C. Share of each material's individual-depositor kilograms delivered by the group			
Plastic	0.16	0.84	0.87
Paper	0.41	0.51	0.61
Metal	0.41	0.44	0.51
Glass	0.38	0.56	0.48

Notes: Individual depositors (unit-bank and institutional accounts excluded), 2013–2026. Pickers and households are the two groups of depositors of Section 3.6. Because the paper's object is the supply of plastic, the 8,275 depositor–bank pairs that never deliver plastic (7 percent of individual-depositor kilograms) are excluded from both columns, which is why the two groups' shares in Panels A and C do not sum to one. A visit is all rows sharing depositor, bank, and timestamp. Panel A: composition shares are kilogram-weighted over the group's deliveries; oil and other materials make up the remainder. Panel B: a grade is a bank-specific plastic line (bank × `category_id`) and resin groups are the platform's plastic category groups; the within-bank gap is the coefficient on the picker indicator in a regression of the pair-level statistic on bank fixed effects, with standard errors clustered by bank, so that menu width is held fixed; the gaps are unchanged among depositors with at least three visits. Panel C: the group's share of the network's individual-depositor kilograms of each material; the last column is the picker share at the 97 treated incumbents of the entry design over the twelve months before the rival's opening. Stars: 10, 5 and 1 percent.

Table 4: What households and pickers deliver: plastic categories, their value, and each group's basket

			Price paid	Share of group's plastic kg		Picker share
	Lines	Banks	Rp/kg	Household	Picker	of category
Panel A. Supply-design sample: 21 banks, January 2022–July 2026						
PET, sorted (cap- and label-free, clear or by colour)	92	12	5,910	0.09	0.71	0.98
PET, mixed or unsorted	121	20	1,366	0.35	0.11	0.61
HDPE (jerrycans, buckets)	91	16	1,134	0.15	0.04	0.58
PP (cups, caps, hard plastic)	72	12	2,158	0.14	0.04	0.59
LDPE, PVC, PS, EPS (film, cassettes, cartons)	62	11	878	0.14	0.05	0.66
Mixed or generic plastic	56	15	821	0.14	0.04	0.58
Within PET, at banks posting both a sorted and an unsorted grade in the month						
Share of the group's PET kg on the sorted grade				0.23	0.87	
Share of the group's PET sellers ever on a sorted grade				0.42	0.52	
Panel B. Network: all banks, January 2018–July 2026						
PET, sorted (cap- and label-free, clear or by colour)	446	184	5,465	0.12	0.33	0.93
PET, mixed or unsorted	886	308	1,672	0.30	0.13	0.70
HDPE (jerrycans, buckets)	608	199	2,376	0.15	0.07	0.72
PP (cups, caps, hard plastic)	417	146	1,875	0.10	0.09	0.83
LDPE, PVC, PS, EPS (film, cassettes, cartons)	436	137	1,672	0.11	0.10	0.83
Mixed or generic plastic	889	200	1,561	0.23	0.26	0.86
Within PET, at banks posting both a sorted and an unsorted grade in the month						
Share of the group's PET kg on the sorted grade				0.37	0.74	
Share of the group's PET sellers ever on a sorted grade				0.49	0.59	

Notes: Individual depositors, pickers and households as in Table 3. Panel A is the sample of the supply design, 951,709 plastic kilograms from 5,005 household and 1,329 picker depositor–bank pairs at 21 banks. Panel B is the network from January 2018, 388 banks, 2,394,047 plastic kilograms, 24,388 household and 3,551 picker pairs. Categories are assigned from the line's label and the platform's resin group. A PET line is *sorted* if its label names a colour, clarity, cleanliness, or pressing (*bening*, *biru*, *bersih*, *press*) and *mixed or unsorted* otherwise, since a bank that posts a single PET line takes bottles as they come; the hub's product codes are assigned from its own code descriptions. Mixed or generic plastic is the platform's residual group, almost entirely lines labelled *plastik* or *plastik campur*; lines labelled *PET campur* that the platform files there are counted as unsorted PET. Price paid is the category's credit divided by its kilograms. Basket shares are kilogram-weighted and sum to one down each column within a panel; the picker share is pickers' kilograms over all individual-depositor kilograms in the category. The within-PET rows condition on bank–months in which the bank transacts on both a sorted and an unsorted PET grade, so that the sorted grade is available to every seller; 12 banks in Panel A and 141 in Panel B.

Table 5: The supply elasticity facing a waste bank, by depositor type

	All individual depositors	By depositor type	
		Picker	Household
Panel A. First stage, reduced form, and the supply elasticity			
First stage: quality-adjusted price index	3.623** (1.083)	4.250** (1.232)	3.645*** (1.081)
Webb wild p-value	0.011	0.019	0.008
Two-way-clustered Wald F	11.2	11.9	11.4
Reduced form: log plastic kilograms	4.264 (2.265)	8.390** (2.589)	2.322 (1.982)
Webb wild p-value	0.174	0.033	0.386
OLS elasticity	0.433 (0.257)	0.680 (0.332)	0.307 (0.266)
2SLS supply elasticity	1.177 (0.738)	1.974** (0.777)	0.637 (0.585)
Webb wild Anderson–Rubin 95% set	[−0.45, 5.25]	[0.45, 9.65]	[−1.20, 3.30]
Stacked test of the difference (bank×type, month×type effects)			
Picker − household	1.337** (0.628), $p = 0.046$		
	940 bank–months, 21 clusters		
Panel B. Where the first stage comes from (all individual depositors)			
Continuing grades (repricing)	1.614 (0.877)		
Webb wild p-value	0.105		
Grades entering the menu	2.009** (0.786)		
Webb wild p-value	0.035		
first transactions	1.141 (0.647)		
re-listings after a silence of > 6 months	0.869 (0.629)		
Share through entering grades	0.55		
Bank–months	578	370	570
Banks	21	14	21
Mean active depositors per bank–month	51.9	12.0	44.8
Mean plastic kilograms per bank–month	1,612	2,063	295

Notes: The shifter is the real-rupiah unit value of Indonesian PET imports (HS 39076100); Appendix Table A7 reports the virgin-resin producer-price index as an alternative. Baseline exposure is the plastic share of the bank's 2018–2021 transaction credit; the endogenous variable is the quality-adjusted plastic price index of Section 3.2. The outcome is plastic bought from individual depositors; unit-bank, institutional, corporate-client, and community-facility accounts are excluded throughout (Section 3.6), so transfers inside the network do not enter. The estimation sample is January 2022–July 2026, conditional on positive plastic transactions. All specifications include bank and calendar-month fixed effects; fixed-effect singletons are dropped. Standard errors in parentheses are two-way clustered by bank and month. The two type columns partition column 1's kilograms exactly; picker depositors are about a fifth of depositor–bank pairs and 82 percent of the plastic kilograms in the outcome. The stacked test identifies the difference from within-bank, within-month variation across the two types; Appendix Table A7 reports its leave-one-bank-out range. In Panel B the continuing- and entering-grade coefficients sum exactly to the first stage, because the entering-grade term is an additive component of the chained index; it is in turn the sum of the first-transaction and re-listing terms. Stars are from the null-imposed Webb wild-cluster p-value (9,999 draws clustered by bank, fixed effects re-estimated in every replication), because with 14–21 clusters the clustered p-value is anti-conservative: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The 2SLS row carries the stars of the reduced form it is built from.

Table 6: How a bank's menu responds to the upstream value shock

	Mean of outcome	Coefficient on instrument
Panel A. The price of a grade, and the grades that enter (log points)		
Posted price within a grade (line and month effects)	—	1.43** (0.61)
kilogram-weighted	—	1.08* (0.57)
First-transaction term of the first stage	—	1.14 (0.65)
finer grade of a resin already bought (refinements)	—	0.85 (0.52)
coarser-or-equal additions	—	0.47 (0.31)
new resins	—	−0.18 (0.18)
Panel B. How often a bank acts on its menu (0/1 per bank–month)		
Any continuing grade repriced	0.30	0.01 (1.53)
Any grade entering the menu	0.40	0.25 (1.09)
first transaction	0.16	0.26 (0.58)
re-listing after a silence of more than six months	0.31	−0.19 (1.19)
Return after 1–6 months of silence (resupply placebo)	0.60	0.86 (1.02)
Any within-resin refinement	0.12	−0.02 (0.48)
Any coarser-or-equal addition	0.06	0.23 (0.49)
Any new-resin addition	0.01	−0.12 (0.28)
Panel C. The response of a continuing line's price by who supplies it (line and bank × month effects)		
	Instrument × household share	Instrument × PET
Household share of the line's kilograms, 2022–26	−0.73* (0.37)	
with the PET interaction	−1.17** (0.54)	−0.80 (0.46)
Household share measured in 2022, outcomes 2023–26	−1.60** (0.63)	
with the PET interaction	−1.83* (0.91)	−0.48 (0.44)
Bank–months (banks), Panels A and B	578 (21)	
Line–months, Panel C	6,783; 3,554 with the 2022 share	

Notes: The instrument is baseline plastic exposure interacted with the real-rupiah unit value of PET imports, as in Table 5; standard errors two-way clustered by bank and month. Panel A: the within-grade rows regress the log posted price of a plastic grade on the instrument with line and month fixed effects over every plastic line–month at the sample banks (7,444 line–months on 493 lines), unweighted and weighted by kilograms; the first-transaction term is that of Table 5, Panel B, split by the kind of entering line, classified against the bank's same-resin menu as in Table 12, the pieces being exact. Panel B: bank–month regressions with bank and calendar-month fixed effects; the repricing row uses the 549 bank–months and 18 banks with a continuing grade, and a return after a silence of one to six months is resupply that the reappearance rule keeps as continuing, a placebo for the re-listing row. Panel C: the log posted price of continuing lines on the instrument interacted with the household share of the line's kilograms (demeaned within bank; standard deviation 0.30) and with a PET indicator, with line and bank-by-month fixed effects, so that the level of the instrument is absorbed; the 2022 share is predetermined for outcomes over 2023–26 (lines transacting in 2022; 11 banks). Stars from the null-imposed Webb wild-cluster p-value in Panel B and from the clustered p-value in Panels A and C: 10, 5 and 1 percent.

Table 7: Margins of the supply response

	All individual depositors	By depositor type	
		Picker	Household
Total plastic kilograms	1.177 (0.738)	1.974** (0.777)	0.637 (0.585)
Webb wild p-value	0.174	0.033	0.386
Webb wild Anderson–Rubin 95% set	[−0.45, 5.25]	[0.45, 9.65]	[−1.20, 3.30]
= Active depositors	0.954 (0.591)	1.087* (0.410)	0.841 (0.560)
Webb wild p-value	0.202	0.065	0.243
Webb wild Anderson–Rubin 95% set	[−0.65, 2.85]	[−0.15, 2.45]	[−0.75, 2.50]
+ Kilograms per active depositor	0.223 (0.544)	0.887 (0.648)	−0.204 (0.551)
Webb wild p-value	0.726	0.199	0.751
Webb wild Anderson–Rubin 95% set	[−1.30, 2.65]	[−0.65, 8.00]	[−2.10, 1.50]
Bank–months	578	370	570
Banks	21	14	21
Stacked picker − household difference, by margin			
Total plastic kilograms	1.337** (0.628), $p = 0.046$		
Active depositors	0.246 (0.598), $p = 0.685$		
Kilograms per active depositor	1.091 (0.633), $p = 0.100$		

Notes: Each cell is a separate 2SLS elasticity with the PET import shifter; log kilograms decomposes identically into log active depositors plus log kilograms per active depositor, so the two components sum exactly to the total in every column. Baseline exposure is the plastic share of the bank's 2018–2021 transaction credit; the endogenous variable is the quality-adjusted plastic price index of Section 3.2. The outcome is plastic bought from individual depositors: unit-bank, institutional, corporate-client, and community-facility accounts are excluded throughout (Section 3.6), so transfers inside the network do not enter. The estimation sample is January 2022–July 2026 and is conditional on positive plastic transactions. All specifications include bank and calendar-month fixed effects; fixed-effect singletons are dropped. Standard errors in parentheses are two-way clustered by bank and month. Depositor types are defined as in Table 5. Note that no individual margin excludes zero under wild inference in any column: the components are less well identified than their sum, so the table supports a statement about which population responds and not about the margin on which it responds. Stars, the wild bootstrap and the 2SLS star convention are as in Table 5.

Table 8: Where the additional picker kilograms arrive

Line status	Share of picker kg	Share of the response
Continuing, posted price unchanged	0.28	0.80
Continuing, repriced up within six months	0.40	0.58
Continuing, repriced down within six months	0.31	−0.61
Transacting for the first time	0.01	−0.05
Returning after a silence of more than six months	0.01	0.28
All plastic kilograms	1.00	1.00
Bank–months (banks)		520 (16)

Notes: Picker deliveries at the banks of the supply design with picker plastic in the window. Each line is classified every month from its own price history under the six-month rule of Section 3.2: a continuing line is repriced up (down) if its posted price exceeds (falls short of) the price it carried six months earlier in the same spell by more than one log point. The outcome for each line status is picker kilograms on lines of that status divided by the bank's mean monthly picker kilograms, regressed on the instrument of Table 5 with bank and calendar-month fixed effects, bank–months weighted equally; because the outcomes sum to total picker kilograms over the bank's mean, the coefficients sum to the total, and the table reports each as a share of it. The first column is the share of picker kilograms on lines of each status in the window. Nearly all picker kilograms on repriced-down lines are at the hub of Section 4.3. The components are individually imprecise; Appendix Table A15 reports the coefficients, the banks-equal version, and the household decomposition.

Table 9: Incumbent quantities and sellers after a rival opens

	Pre-opening level	Months relative to the opening				
		−12 to −7	0–5	6–11	12–17	18–23
All treated incumbents (97)						
Kilograms, all materials	1,170	0.245***	0.058	−0.117	0.003	−0.182*
		(0.088)	(0.078)	(0.138)	(0.133)	(0.109)
Plastic	488	−0.020	0.051	−0.455**	−0.247	−0.352**
		(0.140)	(0.102)	(0.208)	(0.166)	(0.161)
delivered by picker sellers	409	−0.121	−0.057	−0.764***	−0.232	−0.304*
		(0.181)	(0.055)	(0.137)	(0.151)	(0.183)
delivered by household sellers	79	0.086	0.214	0.019	−0.131	−0.355*
		(0.135)	(0.261)	(0.133)	(0.215)	(0.211)
Paper	567	0.041	0.050	0.022	0.142	−0.042
		(0.159)	(0.100)	(0.064)	(0.153)	(0.158)
Glass	35	−0.402*	−0.000	0.020	0.275	0.018
		(0.243)	(0.201)	(0.323)	(0.446)	(0.438)
Metal	40	0.418	−0.251	−0.258	−0.335	−0.572
		(0.495)	(0.260)	(0.296)	(0.312)	(0.377)
Active sellers	26.3	−0.106	−0.112	−0.037	−0.004	−0.049
		(0.099)	(0.069)	(0.073)	(0.099)	(0.164)
Picker	2.4	−0.078	0.044	−0.192**	−0.079	−0.011
		(0.124)	(0.066)	(0.088)	(0.195)	(0.242)
Household	22.4	−0.090	−0.072	0.039	0.043	0.039
		(0.097)	(0.051)	(0.072)	(0.108)	(0.169)
Incumbents in cell		33	91	88	57	30

Notes: Treatment is the incumbent's first roster-dated opening within 15 kilometers (unit banks excluded as entrants). A bank qualifies at an opening if it records in at least one of the three months before it, has at least three recorded months by then, and records at least once in months 6–11 after it; treated and comparison banks face the same rule, and comparison banks are never treated or not yet treated at the horizon, with a not-yet-treated bank's months kept only while it is untreated. Outcomes are measured on months in which the bank records. Every regression carries bank×stack, month×stack and bank-age fixed effects, with age measured from the bank's opening date in the network roster in bins of 0–2, 3–5, 6–11, 12–23, 24–47 and 48 and more months; the reference period is the average of months −6 to −1. Standard errors, in parentheses, are clustered by the entrant. Stars: 10, 5 and 1 percent. Quantities are monthly totals over individual depositors (unit-bank and institutional accounts excluded), estimated by stacked Poisson pseudo-maximum likelihood; coefficients are log-point changes in expected quantity relative to comparison banks and to the reference period. Plastic delivered by picker and by household sellers sums to plastic. The pre-opening level is the mean over treated incumbents of the monthly value averaged over recorded months −6 to −1 (kilograms, or number of depositors). The Bandung March-2019 cohort is included (Appendix Table A9 excludes it).

Table 10: Quantities after a rival opens, by the incumbent's seller mix before the opening

	Pre-opening level	Months relative to the opening				
		−12 to −7	0–5	6–11	12–17	18–23
Panel A. Picker-intensive incumbents (47)						
Kilograms, all materials	2,209	0.207**	−0.048	−0.228*	0.094	−0.183
		(0.082)	(0.047)	(0.133)	(0.166)	(0.147)
Plastic	944	−0.085	−0.064	−0.625***	−0.188	−0.350
		(0.159)	(0.049)	(0.170)	(0.207)	(0.216)
delivered by picker sellers	842	−0.182	−0.097**	−0.846***	−0.216	−0.369
		(0.194)	(0.043)	(0.110)	(0.206)	(0.294)
delivered by household sellers	102	0.004	−0.116	−0.097	−0.196	−0.337
		(0.145)	(0.095)	(0.178)	(0.293)	(0.251)
Active sellers	26.4	−0.120	−0.074	0.022	0.054	0.042
		(0.107)	(0.062)	(0.109)	(0.154)	(0.259)
Picker	4.7	−0.084	0.029	−0.242***	−0.143	−0.189
		(0.125)	(0.069)	(0.087)	(0.224)	(0.300)
Household	20.3	−0.117	−0.066	0.119	0.108	0.123
		(0.111)	(0.062)	(0.116)	(0.153)	(0.238)
Incumbents in cell		21	43	42	26	12
Panel B. Household-intensive incumbents (48)						
Kilograms, all materials	196	0.207	0.250	0.059	0.125	−0.083
		(0.293)	(0.238)	(0.146)	(0.223)	(0.240)
Plastic	61	−0.260	0.470	0.127	−0.280	−0.276
		(0.351)	(0.420)	(0.168)	(0.218)	(0.256)
delivered by picker sellers	2	−1.388*	0.928**	0.796	0.922	1.159*
		(0.819)	(0.438)	(0.568)	(0.780)	(0.661)
delivered by household sellers	59	0.097	0.562	0.155	−0.143	−0.423
		(0.333)	(0.490)	(0.172)	(0.228)	(0.268)
Active sellers	26.8	−0.067	−0.149*	−0.144*	−0.119	−0.189
		(0.171)	(0.083)	(0.078)	(0.162)	(0.193)
Picker	0.2	−0.154	0.146	0.183	0.557**	0.880***
		(0.764)	(0.225)	(0.208)	(0.235)	(0.314)
Household	25.3	−0.022	−0.089	−0.056	−0.072	−0.055
		(0.143)	(0.067)	(0.061)	(0.168)	(0.223)
Incumbents in cell		12	46	44	29	16

Notes: Design as in Table 9. An incumbent's type is the share of its individual-depositor plastic kilograms delivered by picker sellers over its months before the opening; incumbents above the treated median share of 0.119 are picker-intensive and those at or below it household-intensive (2 incumbents have no plastic seller of either type before the opening). Comparison banks are the same pool for both panels. Picker sellers are few at household-intensive incumbents (0.16 per month before the opening), so their coefficients describe small absolute changes. Appendix Table A12 reports every outcome by type. Standard errors, in parentheses, are clustered by the entrant. Stars: 10, 5 and 1 percent.

Table 11: Incumbent posted prices, the quality-adjusted index, and menu refinement after a rival opens

	Months relative to the opening				
	−12 to −7	0–5	6–11	12–17	18–23
Panel A. All treated incumbents (97)					
Posted price of a grade on the menu: plastic	0.024	−0.025*	−0.034*	−0.040	−0.060
	(0.028)	(0.015)	(0.018)	(0.030)	(0.041)
Paper	0.014	−0.017	−0.018	0.006	−0.004
	(0.022)	(0.016)	(0.016)	(0.032)	(0.055)
Glass	−0.027	−0.060	−0.063	−0.011	0.021
	(0.048)	(0.036)	(0.069)	(0.067)	(0.073)
Metal	−0.012	−0.009	0.004	0.026*	0.060***
	(0.034)	(0.009)	(0.014)	(0.014)	(0.019)
Quality-adjusted plastic index	−0.004	−0.020	−0.029	0.051	0.064
	(0.041)	(0.020)	(0.052)	(0.049)	(0.093)
New-grade component	−0.023	−0.007	−0.018	0.023	0.035
	(0.021)	(0.010)	(0.024)	(0.033)	(0.083)
Continuing-grade component	0.019	−0.012	−0.011	0.028	0.029
	(0.042)	(0.017)	(0.034)	(0.031)	(0.038)
Any within-resin refinement in the month (0/1)	−0.028	0.067***	0.113***	0.152***	0.191***
	(0.041)	(0.022)	(0.033)	(0.038)	(0.067)
Number of refinements	−0.151	1.749***	1.723***	2.047***	2.191***
	(0.505)	(0.384)	(0.618)	(0.430)	(0.588)
Any coarser-or-equal addition in the month (0/1)	−0.058*	−0.054	0.044	0.036	−0.040
	(0.034)	(0.052)	(0.030)	(0.038)	(0.054)
Incumbents with a plastic line in cell	32	90	85	56	28
Panel B. Rival within 5 kilometers (29)					
Posted price of a plastic grade on the menu	−0.038***	0.011	−0.020	−0.056**	−0.034**
	(0.012)	(0.012)	(0.018)	(0.022)	(0.014)
Quality-adjusted plastic index	−0.093**	−0.000	−0.029	0.109	0.021
	(0.035)	(0.025)	(0.067)	(0.082)	(0.101)
New-grade component	−0.032	−0.023	−0.047	0.068	0.002
	(0.034)	(0.022)	(0.067)	(0.081)	(0.106)
Continuing-grade component	−0.060***	0.022**	0.019	0.041	0.019
	(0.013)	(0.010)	(0.011)	(0.027)	(0.021)
Any within-resin refinement in the month (0/1)	−0.069	0.044	0.155*	0.250**	0.094
	(0.099)	(0.033)	(0.076)	(0.089)	(0.056)
Any coarser-or-equal addition in the month (0/1)	−0.048	−0.110	0.081	0.085	−0.090
	(0.056)	(0.089)	(0.084)	(0.133)	(0.167)
Incumbents with a plastic line in cell	10	27	25	19	8
Panel C. Rival between 5 and 15 kilometers (68)					
Posted price of a plastic grade on the menu	0.044	−0.037*	−0.037	−0.041	−0.063
	(0.033)	(0.018)	(0.022)	(0.034)	(0.045)
Quality-adjusted plastic index	0.025	−0.028	−0.031	0.035	0.066
	(0.047)	(0.025)	(0.054)	(0.055)	(0.103)
New-grade component	−0.019	−0.004	−0.010	0.012	0.039
	(0.027)	(0.013)	(0.025)	(0.037)	(0.091)
Continuing-grade component	0.044	−0.024	−0.021	0.023	0.027
	(0.051)	(0.026)	(0.050)	(0.041)	(0.047)
Any within-resin refinement in the month (0/1)	−0.014	0.078**	0.097**	0.120***	0.199**
	(0.044)	(0.030)	(0.037)	(0.042)	(0.076)
Any coarser-or-equal addition in the month (0/1)	−0.066	−0.030	0.026	0.018	−0.037
	(0.040)	(0.041)	(0.034)	(0.033)	(0.054)
Incumbents with a plastic line in cell	22	63	60	37	20

Notes: Design as in Table 9. Posted prices are estimated from every bank–grade line by stacked OLS with line×stack, month×stack and bank-age fixed effects; the quality-adjusted index and its components by stacked OLS on the bank–month hedonic index. The posted price of a grade is the modal credit per kilogram of a bank–grade line over individual depositors; every line enters, including grades added after the opening, and the line fixed effect measures each price against the line's own mean. Coefficients are log points relative to the reference period. The quality-adjusted index follows the six-month reappearance rule of Section 3.2. A within-resin refinement is a first transaction on a plastic grade in a resin the bank already buys whose hedonic attribute index exceeds that of the coarsest same-resin grade with kilograms in the prior six months; a coarser-or-equal addition is a first transaction in a resin the bank already buys that does not (Table 12). The 0/1 rows are stacked linear probability models on recorded bank–months, with a refinement in 11 percent of recorded bank–months; the count row is stacked Poisson pseudo-maximum likelihood. Standard errors, in parentheses, are clustered by the entrant. Stars: 10, 5 and 1 percent.

Table 12: Within-resin refinements: what they are, what they do to the coarse grade, and who supplies them

Panel A. What a refinement is			
First transactions on a plastic grade, banks with three or more prior recorded months	1,274		
within-resin refinements (finer than the coarsest same-resin grade)	850		
coarser-or-equal additions	283		
new-resin additions	141		
Share of refinements posted above the coarse grade's current price	0.73		
at the same price	0.07		
below it	0.19		
Median posted-price gap over the coarse grade (log points)	0.69		
refinements within twelve months after a rival opening (285)	0.69		
Median gap in the hedonic attribute index over the coarse grade (log points)	0.62		
Panel B. The coarse grade after the addition (event-level; months relative to the addition)			
	0–2	3–5	6–9
Coarse-grade kilograms (asinh)			
after a refinement (pre-period mean 2.51)	−0.518** (0.200)	−0.843*** (0.276)	−0.859*** (0.309)
after a coarser-or-equal addition (pre-period mean 2.58)	−0.190 (0.181)	−0.204 (0.193)	0.053 (0.267)
Coarse grade's share of the resin's kilograms			
after a refinement (pre-period mean 0.27)	−0.102*** (0.012)	−0.120*** (0.014)	−0.122*** (0.017)
after a coarser-or-equal addition (pre-period mean 0.56)	−0.240*** (0.031)	−0.241*** (0.034)	−0.229*** (0.055)
New grade's share of the resin's kilograms			
after a refinement (pre-period mean 0.00)	0.127*** (0.016)	0.098*** (0.015)	0.109*** (0.019)
after a coarser-or-equal addition (pre-period mean 0.00)	0.294*** (0.038)	0.271*** (0.038)	0.309*** (0.040)
Coarse and new grade together, kilograms (asinh)			
after a refinement (pre-period mean 2.51)	0.646*** (0.111)	0.150 (0.201)	0.113 (0.255)
after a coarser-or-equal addition (pre-period mean 2.58)	0.915*** (0.113)	0.936*** (0.172)	1.173*** (0.225)
Panel C. Who supplies the refinement (months 0 to 9 after it)			
Share of the new grade's kilograms from sellers who had delivered on the coarse grade	0.08		
from sellers whose first plastic deposit at the bank is in or after the event month	0.69		
from picker sellers	0.88		
Share of their resin kilograms that continuing sellers of the coarse grade move to the new grade: pickers	0.03		
households	0.04		

Notes: Panel A classifies every first transaction on a plastic grade (a bank–grade line first observed) at banks with at least three prior recorded months, 2013–2026, by its position in the bank's menu. A refinement is finer than the coarsest same-resin grade that had kilograms in the prior six months, meaning its hedonic attribute index from equation (2) is higher; a coarser-or-equal addition is not; a new-resin addition is in a resin the bank did not carry. The posted-price gap is the new grade's log posted price in its first month minus the coarse grade's current posted price, its price that month or its last observed price when it is silent that month. The attribute gap is the difference in the two grades' hedonic attribute indices, which is the premium the quality-adjusted index nets out by construction. Panel B is an event study at the level of the coarse grade with event and calendar-month fixed effects, reference months −3 to −1, and standard errors clustered by bank; outcomes are the coarse grade's monthly kilograms (inverse hyperbolic sine), its share of the bank's kilograms in that resin, the new grade's share, and the two grades' combined kilograms. Panel C uses individual-depositor transactions in months 0 to 9 after a refinement, with shares weighted by kilograms across events; seller types are defined in Section 3.6. Stars: 10, 5 and 1 percent.

Table 13: Who delivers on the grades added after a rival opens

	Picker sellers	Household sellers
Share of the type's plastic kilograms delivered on new grades	0.55	0.31
median across incumbents	0.09	0.04
Share of the type's sellers delivering on a new grade at least once	0.59	0.47
Kilograms per seller-month on new grades	201	4.3
Grades delivered on per seller-month	3.8	2.5
Sellers	656	4,299
Picker share of kilograms: new grades	0.88	
continuing grades	0.72	
Picker share of sellers: new grades	0.16	
continuing grades	0.12	
Share of post-opening plastic kilograms on new grades	0.51	
Price premium of new grades over the bank's continuing grades (log points)	0.21 (0.08)	

Notes: Individual-depositor plastic transactions at the 63 treated incumbents (of 90 with post-opening plastic) that post at least one plastic grade not on their menu before the opening, over months 0 to 23 after it. A grade is new if the bank–grade line is first observed at or after the opening. The premium is the kilogram-weighted mean log credit per kilogram on new grades minus that on continuing grades within the bank, averaged over the 62 incumbents with both; its standard error is across incumbents. It is a raw comparison of posted prices and makes no adjustment for the attributes of the grades, so it is not the quality-adjusted premium of the index.

Table 14: Entry effects by the distance of the rival

Distance band	Incumbents	Plastic-grade price 6–11 m	Plastic-grade price 12–17 m	Plastic-grade price 18–23 m	Quality-adj. index 12–17 m	New-grade component 12–17 m	Any refinement 12–17 m	Plastic kg 6–11 m
0–2 km	9	0.031	−0.007	−0.037	0.229**	0.200*	0.194	−0.197
		(0.045) [6]	(0.038) [4]	(0.029) [3]	(0.089) [4]	(0.098) [4]	(0.100) [5]	(0.283) [7]
2–5 km	20	−0.029*	−0.066**	−0.035**	0.047	−0.012	0.291**	−0.592
		(0.016) [19]	(0.021) [15]	(0.012) [5]	(0.111) [15]	(0.106) [15]	(0.128) [15]	(0.455) [19]
5–10 km	42	−0.028*	−0.027*	−0.065**	−0.011	0.007	0.058	−0.434*
		(0.015) [37]	(0.014) [21]	(0.024) [10]	(0.053) [21]	(0.031) [21]	(0.060) [21]	(0.263) [38]
10–15 km	26	−0.051	−0.061	−0.063	0.085	0.011	0.211***	−0.209
		(0.049) [23]	(0.074) [16]	(0.097) [10]	(0.094) [16]	(0.068) [16]	(0.048) [16]	(0.260) [24]

Notes: Design as in Tables 9 and 11, with treated incumbents grouped by the distance of the rival; the index follows the six-month reappearance rule and the refinement indicator is defined in Table 11. Brackets give the number of incumbents identifying the cell. Quantity cells by distance are reported for completeness; the close-rival quantity samples carry positive pre-period coefficients (Appendix Table A13) and are not interpreted. Standard errors, in parentheses, are clustered by the entrant. Stars: 10, 5 and 1 percent.

Table 15: Onward sales at one large operator, 2024–25

Panel A. Kilograms and unit values bought and sold, by material

	Bought (tons)	Sold (tons)	Price paid (Rp/kg)	Price received (Rp/kg)	Gross spread $1 - p/P$	Ratio $(P - p)/p$
PET	530	375	5,376	7,371	0.27	0.37
Other plastic	389	296	1,422	2,456	0.42	0.73
Paper	706	683	1,149	1,885	0.39	0.64
Metal	64	68	2,950	4,770	0.38	0.62
Cooking oil	37	32	6,859	9,975	0.31	0.45
Glass	62	72	212	581	0.64	1.74
All materials	1,788	1,526	2,396	3,580	0.33	0.49

Panel B. Pass-through of the price received into the price paid, log on log

	Levels	Levels, with month effects	First difference	Cumulative, three lags	Obs.	Units
PET pair: pressed clear bales sold, P12 Mix bought	1.09*** (0.06)	—	0.98*** (0.06)	1.07*** (0.12)	24	1 pair
Material panel, material effects	0.61* (0.27)	—	0.43* (0.19)	0.97*** (0.16)	144	6 materials
Grade panel, grade effects	0.24** (0.10)	0.39*** (0.14)	0.07** (0.03)	0.55*** (0.18)	493	24 grades

Panel C. Buyers, by material and year

	Number of buyers		Share of the largest		Herfindahl index	
	2024	2025	2024	2025	2024	2025
PET	24	26	0.39	0.46	0.25	0.35
Other plastic	47	52	0.28	0.38	0.16	0.22
Paper	23	24	0.80	0.43	0.65	0.25
Metal	19	36	0.34	0.64	0.19	0.45
Cooking oil	3	5	0.77	0.85	0.63	0.73
Glass	3	5	0.71	0.62	0.59	0.45
All materials	80	92	0.24	0.27	0.12	0.13

Notes: The operator is a large hub bank (*bank sampah induk*) in Bandung Regency whose purchases are in the transaction data and whose sales ledger (date, buyer, product, kilograms, price) covers January 2024 to December 2025. Purchases are the bank's credited deliveries from all depositors, including feeder banks, which supply about half its kilograms. Products are matched to purchase grades by the shared product code and grades to materials by the platform's category groups; PET is bought loose and sold as pressed bales, so it is matched at the material level. Service revenue (1.6 percent of revenue) is excluded. Panel A: unit values are value over kilograms across the two years; the gross spread is one minus the price paid over the price received, and the ratio is the same difference relative to the price paid; "All materials" values the kilograms sold at the bank's own material-level buying price. Kilograms sold and bought balance within 13 percent for every material in 2025; in 2024 the bank bought 366 tons of PET and sold 210. Panel B: regressions of the log price paid on the log price received, at the month. The PET pair is the bank's largest sale product and largest purchase grade, with heteroskedasticity-robust standard errors; the material and grade panels are weighted by kilograms sold, keep grades with at least twelve matched months and 50 kilograms on each side, and cluster by material or grade. The cumulative column sums the contemporaneous and lagged first-difference coefficients (two lags for the pair and the material panel, three for the grade panel). Panel C: buyers are counted within each annual ledger; shares and the Herfindahl index are of sales value. Stars: 10, 5 and 1 percent.

Table 16: Fiscal cost per ton diverted under a small operator-side subsidy

Payment	Pass-through ρ	Cost per ton induced, $p/(\delta\,\varepsilon^S\rho)$		
		$\delta = 1$	$\delta = \frac{1}{2}$	$\delta = \frac{1}{4}$
Uniform subsidy, pooling bank	0.54	Rp 1.91m (US$115)	Rp 3.82m (US$230)	Rp 7.63m (US$460)
Uniform subsidy, segmenting bank	0.66 / 0.39	Rp 1.60m (US$96)	Rp 3.20m (US$193)	Rp 6.40m (US$386)
Payment to pickers only (1.97)	0.66	Rp 1.37m (US$82)	Rp 2.74m (US$165)	Rp 5.47m (US$330)
Payment to households only (0.64)	0.39	Rp 7.15m (US$431)	Rp 14.31m (US$862)	Rp 28.62m (US$1,724)
Anderson–Rubin bounds on the payment to pickers				
Picker clustered AR, lower end (0.80)	0.44	Rp 5.03m (US$303)	Rp 10.05m (US$606)	Rp 20.11m (US$1,211)
Picker wild AR, lower end (0.45)	0.31	Rp 12.80m (US$771)	Rp 25.60m (US$1,542)	Rp 51.20m (US$3,084)

Notes: The benchmark values a per-kilogram operator-side subsidy at the mean bank–month unit value of plastic bought from individual depositors, Rp 1,787 per kilogram, and converts at Rp 16,600 to the dollar. The induced response in every row is the kilogram-share-weighted sum of the picker and household responses, at elasticities of 1.97 and 0.64 (Table 5) and network shares of 0.82 and 0.18. The rows differ in how the subsidy reaches each type. A pooling bank posts one price and passes through $\rho = \varepsilon/(1+\varepsilon) = 0.54$ at the aggregate elasticity of 1.18 on every grade. A segmenting bank separates the types through its menu and passes through $\varepsilon_W/(1+\varepsilon_W) = 0.66$ on the grades pickers deliver and $\varepsilon_H/(1+\varepsilon_H) = 0.39$ on the grades households deliver (equation (22)). A payment restricted to one type is paid on that type's kilograms only. The subsidy is assumed small enough that the elasticities are locally constant. δ is the share of induced kilograms that is not redirected from another buyer, and is not identified here, so the table varies it rather than estimating it. The Anderson–Rubin sets for all depositors reach zero, so no finite upper bound exists for the first two rows; the picker sets exclude zero, and their lower endpoints of 0.80 (clustered) and 0.45 (wild) supply the bounds in the last two rows. The figures are fiscal costs, which scale inframarginal subsidy payments by the local quantity response; they are not resource costs and they are a sensitivity exercise, not a confidence interval. Benchmark: the ministry costs sanitary landfill at Rp 250,000 per ton.

Table 17: Bounds on the environmental value of a diverted ton of plastic

Counterfactual fate of the ton	US$ per ton	Basis
Damage avoided, d		
Landfill (lower bound, d_L)	20–25	Disposal cost Rp 250,000 (US$15) plus landfill disamenity and airspace, US$5–10; no methane or leachate from inert plastic
Open burning (upper bound, d_U)	300–1,500	Sum of the two components below, rounded
mortality	230–1,150	270,000 deaths per year from open waste burning over roughly one billion tons burned, at an Indonesian VSL of US$0.84 million (US$230 per ton of mixed waste), scaled by one to five for the higher particulate yield of plastic
combustion CO_2e	60–300	1.2–1.6 tCO_2e per ton at a social cost of carbon of US$51–190
Waterway or ocean, infant-health channel	5–90	Low-birth-weight effect of marine microplastics, 0.37 per 1,000 births per doubling, applied to Indonesia's coastal births, divided by annual leakage of 620,000 tons, and valued at US$1,300 per case (Section 5.3)
with LBW mortality at the VSL	150–3,000	US$30–600 per ton per percentage point of excess infant mortality; shown for five points
ecosystem services	$\gg d_U$	US$3,300–33,000 per ton-year, from valuation studies
Indonesian average ton	180–900	26 percent landfilled, 58 percent burned or dumped, 16 percent recovered
Avoided virgin production (independent of counterfactual)		
Displacement rate 1.0	60–250	1.0–1.3 tCO_2e per ton recycled at US$51–190 per ton CO_2
Displacement rate 0.5	30–125	
Displacement rate 0.1	6–25	No empirical displacement rate exists for plastics; aluminum 0.1, copper 0.5

Notes: Sources: disposal cost, Tempo (2025*a*); landfill externalities, Eshet, Ayalon and Shechter (2006) and Nahman (2011); open-burning mortality, Kodros et al. (2016) using the burned tonnage of Wiedinmyer, Yokelson and Gullett (2014); plastic-specific emissions, Pathak et al. (2023) and Wang et al. (2023); value of a statistical life, Vital Strategies (2024); marine microplastics, Du, Zhang and Zou (2024), United Nations (2024), Ekoriano et al. (2025), SYSTEMIQ and the Indonesia National Plastic Action Partnership (2020), Alderman and Behrman (2006), and Beaumont et al. (2019); fate shares, Ramadan et al. (2026); life-cycle factors, U.S. Environmental Protection Agency (2023*a*) and Association of Plastic Recyclers (2020); social cost of carbon, U.S. Environmental Protection Agency (2023*b*); displacement, Zink and Geyer (2017) and Lowe et al. (2025). Dollar ranges are rounded. The open-burning range d_U is the mortality component, the mixed-waste figure of US$230 per ton scaled by one to five for the higher black-carbon and particulate yield of plastic, plus the combustion CO_2e component, rounded to US$300–1,500.

Table 18: Break-even conditions: the share of the picker response that must be new collection

Counterfactual fate of new picker collection	US$ per ton	Required share of new collection, δ_W: Uniform subsidy, pooling bank (US$115)	Uniform subsidy, segmenting bank (US$96)	Payment to pickers (US$82)
Sanitary landfill	20–25	never	never	never
Open burning, lower end	300	38%	32%	27%
Open burning, upper end	1,500	8%	6%	5%
Waterway or ocean, infant-health channel	5–90	never	never	never
LBW mortality valued at the VSL, lower end	150	77%	64%	55%
LBW mortality valued at the VSL, upper end	3,000	4%	3%	3%
ecosystem services	3,300–33,000	0.3–3.5%	0.3–2.9%	0.2–2.5%

Notes: Each cell is the smallest share of the picker response that must be newly collected material, rather than material redirected from another buyer, for the payment to pay for itself on avoided damage alone, $\delta_W \geq (C/\Delta Q)/\bar{d}_W$ from equation (24), when all new collection would otherwise have met the fate in the row. The fiscal cost per induced ton is Rp 1.91 million (US$115) for a uniform subsidy at a bank that posts one price to all sellers, Rp 1.60 million (US$96) for the same subsidy at a bank that separates pickers from households through its menu, and Rp 1.37 million (US$82) for a payment restricted to pickers (Table 16). "Never" means the condition fails even when the whole response is new collection. In resource-cost terms, at US$108 per ton, the uniform thresholds are 36 and 7 percent against open burning and 72 and 3.6 percent against the VSL-valued marine range. Damage values are those of Table 17.

Appendix

A Conceptual Framework

This section sets out a simple model of a waste bank that posts prices to two kinds of sellers. The framework is not meant to describe the recycling chain in full. Its purpose is to make precise why the two groups should respond differently to a posted price, how a buyer that cannot price by seller identity uses the grade menu instead, and what a rival's opening does to both. Each part ends with the prediction that Section 4 takes to the data. The residual-supply elasticity and the picker elasticity are those of equations (12) and (13) in Section 5.

I make four simplifying assumptions. The bank takes the downstream price at which it resells as given, so it has market power over its sellers but not over its buyers. The flow of recyclable material generated in its catchment is fixed. The bank's price determines how much of it is collected and how much of what is collected is delivered to the bank, not how much is generated. The bank posts one price per kilogram per grade and buys whatever is delivered at that price. And the bank maximizes operating surplus. Some banks are community organizations, but the surplus funds their continuation and none of the predictions depends on how it is spent. Material, bank, and month subscripts are suppressed.

A.1 Two Kinds of Sellers

A household generates recyclable material as a by-product of its own consumption and delivers to the bank if the return covers the cost of sorting and carrying it. The return has two parts, the posted price p and a non-pecuniary value ϕ per kilogram that stands for the civic and environmental motives surveys identify as the main reason Indonesian households use a waste bank (Widayat et al., 2025). Household i delivers if $p+\phi+\eta_i \geq k(a)$, where $k(a)$ is the cost of reaching the bank, decreasing in its proximity and convenience a, and η_i varies across households with distribution G. Household supply is $H(p;a) = \bar{g}\, N_H\, [1 - G(k(a) - p - \phi)]$, where $\bar{g}$ is the material a participating household generates, with price elasticity

$$\varepsilon_H = \frac{p}{p+\phi} \times \frac{(p+\phi)\, g(k-p-\phi)}{1 - G(k-p-\phi)}, \tag{25}$$

the product of the pecuniary share of the marginal household's return and the elasticity of participation with respect to the return as a whole. When civic motives dominate, the

first term is small and household supply barely responds to the price even if participation responds strongly to the total return. A fall in $k(a)$ shifts H in level, which no feasible price can replicate.

A picker collects for a living and chooses how much to collect, whether to be active at all, and to which buyer to sell each load. Let $\bar{p}$ be the best net price available from other buyers and A the density of such buyers near the bank. Picker supply to the bank is $W(p, \bar{p}; A)$, with the elasticity of equation (13), a redirection term for material another buyer would otherwise have received, which rises with A, and a collection term for material that would not otherwise have been collected. Because the posted price does not depend on who delivers or how much (Section 3.1), residual supply is the sum $Q(p) = H(p; a) + W(p, \bar{p}; A)$, with the share-weighted elasticity of equation (12). A bank's market power therefore depends on its clientele and on its local market, and residual supply is more elastic the larger the picker share and the more alternative buyers pickers have nearby.

The ranking $\varepsilon_W > \varepsilon_H$ is not assumed. It follows from equation (25) only if the pecuniary share of the marginal household's return is small, and from equation (13) only if pickers' collection or choice of buyer responds to the return. The prediction I take to Section 4.1 is that the posted price moves the quantity pickers deliver, on both the participation and the per-seller margin, and does not detectably move the quantity households deliver.

A.2 A Uniform Price and the Grade Menu

The bank sells a kilogram onward at P and incurs a handling cost c per kilogram. Let $v = P - c$ be the net downstream value. Choosing p to maximize $(v - p)\, Q(p)$ gives, for a locally constant elasticity,

$$p^* = \frac{\varepsilon}{1 + \varepsilon}\, v, \qquad \frac{v - p^*}{p^*} = \frac{1}{\varepsilon}, \qquad \frac{d \ln p^*}{d \ln v} = 1. \tag{26}$$

The markdown below downstream value is the inverse of the residual supply elasticity, and the posted price moves one for one in proportional terms with the value of the material. A bank that posts one price per grade cannot mark down each group by its own elasticity. The single markdown it chooses is governed by the mix in equation (12), and it is widest at banks whose sellers are predominantly households.

The bank cannot pay the two groups different prices for the same grade, but it chooses which grades to post. Suppose a share χ of a seller's material can be sorted into a finer grade, such as clear rather than mixed PET, for which downstream buyers pay a premium Δ_f over the coarse grade's value. Sorting a load has a fixed component κ, the time to separate it by

resin and color before it is weighed, and a per-kilogram component ξ_0, so the sorting cost per kilogram of a load of size q is $\xi(q) = \xi_0 + \kappa/q$. A picker's load is an order of magnitude larger than a household's, 41 against 3 kilograms at the median visit (Section 3.6), so sorting is cheaper per kilogram for pickers, $\xi_W < \xi_H$. The bank posts a price p for the coarse grade and may add the fine grade at a premium ζ over it, at a fixed cost F_m per period for handling, storing, and selling an additional grade.

If $\xi_W \le \zeta < \xi_H$, pickers sort and households do not. A picker then receives, per kilogram delivered,

$$p_W = p + \chi\,(\zeta - \xi_W), \tag{27}$$

while a household receives p. The bank pays $p + \chi\zeta$ per picker kilogram and sells it for $v + \chi\Delta_f$, so its surplus is

$$(v - p)\,H(p) + (v_W - p_W)\,W(p_W) - F_m, \qquad v_W \equiv v + \chi\,(\Delta_f - \xi_W), \tag{28}$$

which is separable in the two prices. Under locally constant elasticities the bank sets

$$p_H^* = \frac{\varepsilon_H}{1 + \varepsilon_H}\,v, \qquad p_W^* = \frac{\varepsilon_W}{1 + \varepsilon_W}\,v_W, \tag{29}$$

and implements them with the premium $\zeta^* = \xi_W + (p_W^* - p_H^*)/\chi$, which is positive when $\varepsilon_W > \varepsilon_H$. The fine grade lets a bank that posts one price per grade set the group-specific markdowns it would set if it could price by identity. The bank cannot condition its price on who a seller is, but it offers a choice of grades that the two groups select differently. This is second-degree price discrimination in the sense of Mussa and Rosen (1978), carried out on the buying side of a market.

Three implications follow. First, the menu moves incidence. The price paid to households falls toward p_H^* and the effective price paid to pickers rises toward p_W^*, so the markdown on households widens and the markdown on pickers narrows. Second, the menu is worth posting only where the gain from separating the two groups exceeds F_m. That gain rises with the volume pickers supply and with the gap between ε_W and ε_H, so finer menus should be found at picker-supplied banks, and a bank with a captive picker base has little reason to post one. Third, a change in v moves the household price at the rate $\varepsilon_H/(1 + \varepsilon_H)$, close to zero when household supply is inelastic, and the picker price at the rate $\varepsilon_W/(1 + \varepsilon_W)$. Revising the posted price of a grade already on the board is costly, because the new number applies to every seller and must be communicated to a clientele that visits intermittently and is credited to savings balances. A fixed cost of revision produces an inaction band, inside which a bank leaves a continuing grade's price alone and lets the value of the material reach

sellers through the prices it writes on grades entering or returning to the menu. Section 4.3 tests the prediction that the posted price moves one for one with downstream value at the one bank where that value is observed, and Section 4.1 asks on which grades the response to upstream value arrives and which sellers deliver the kilograms it brings in. Section 3.6 and Section 4.2 examine where fine grades appear and who delivers on them.

A.3 Competition for Sellers

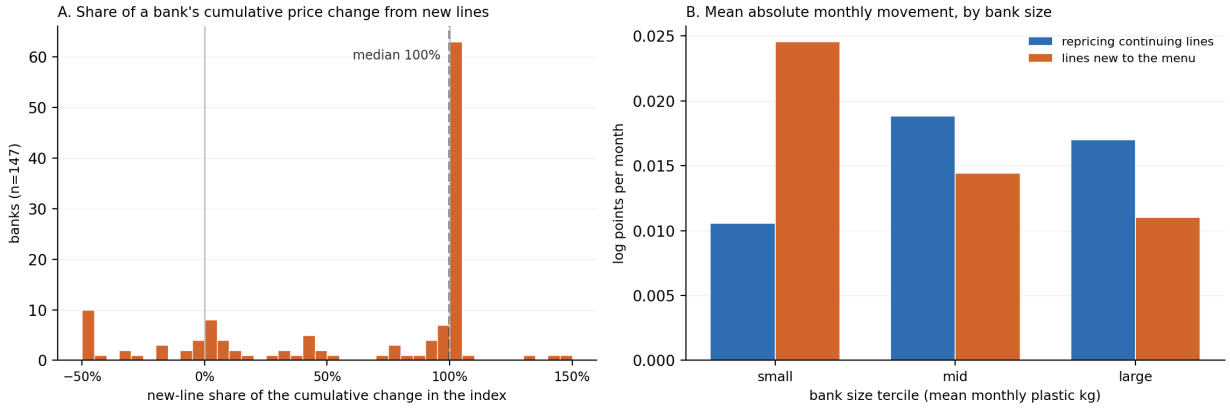
A rival's opening raises A for nearby incumbents and gives their pickers a new outside option at net price $\bar{p}$. Pickers for whom $\bar{p}$ exceeds the incumbent's effective price redirect their loads once the entrant is buying, while households move only if the entrant is more convenient, through $k(a)$, and not in response to its price. On impact the incumbent therefore loses volume and sellers among pickers and in the material pickers carry, and not among households. Because a picker's range covers every buyer within a few kilometers, the loss need not fall with distance inside that range.

The rival also raises ε_W through the redirection term in equation (13), which widens the gap between the two groups' elasticities and pushes the return to separating them above F_m at incumbents where it had been below. The incumbent raises p^*_W, which by equation (27) it does through the premium, adding a finer grade at a higher price. It pays more to the sellers it is losing without paying more on the household kilograms it would keep in any case. A quality-adjusted index nets the premium out by construction, so the quality-adjusted price level need not rise even though the price available to a seller who sorts does. On the coarse grade the remaining sellers are disproportionately households, so s_W falls, residual supply on it becomes less elastic, and by equation (26) the posted price of an existing grade falls. Retaining pickers by raising the posted price of an existing grade instead would pay the increase on every household kilogram as well, which is unprofitable when ε_H is near zero.

Section 4.2 tests three predictions. An incumbent loses picker supply and not household supply. It refines its menu, adding a finer grade of a resin it already buys, posted above the coarse grade, and pickers rather than households move onto it. And the posted price of the grades it already carries falls, while the quality-adjusted price level does not rise. Because the fine grade is worth its fixed cost only where pickers supply enough to fill it, picker-supplied incumbents should refine more than household-supplied ones, whose response is confined to the markdown on the coarse grade.

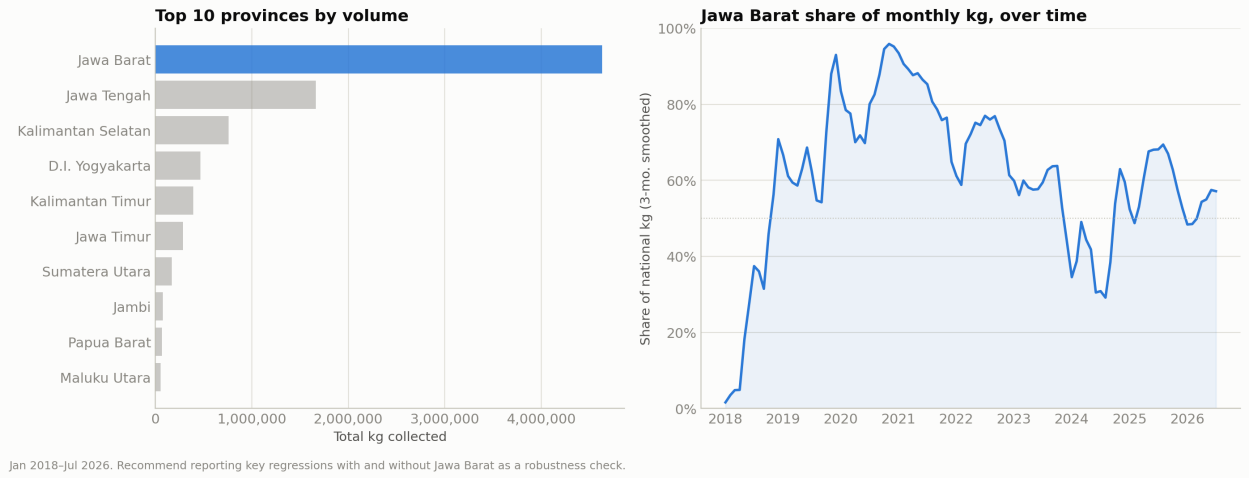
B Additional Figures

Figure A1: What the quality-adjusted price index is made of



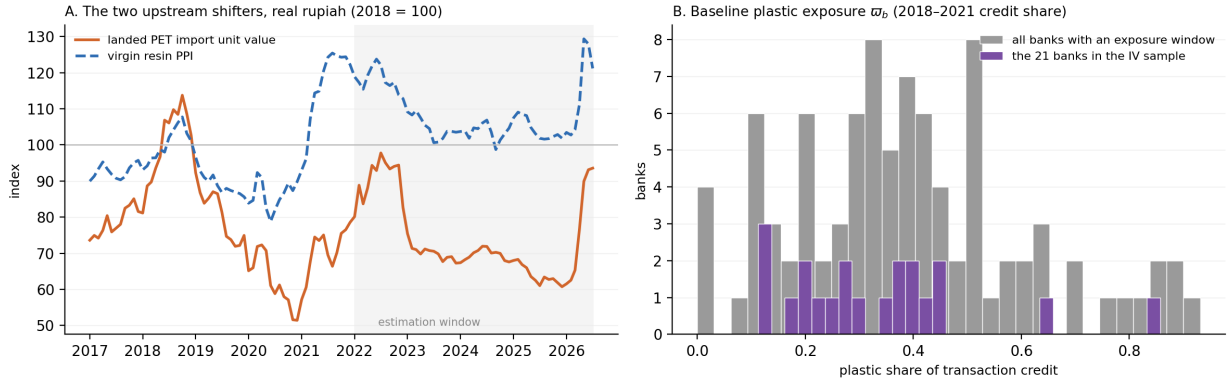
Notes: Panel A plots, across the 147 banks observed for at least six months with a non-trivial cumulative change, the share of the cumulative change in the chained quality-adjusted index that arrives through lines new to the menu, a line being new when it first transacts or returns after a silence of more than six months, rather than through the repricing of continuing grades. The median bank is at 100 percent and the interquartile range is 15.7 to 100 percent; ten banks sit at the clipped left edge, where the two components offset. Panel B reports the mean absolute monthly movement of each component by bank size tercile.

Figure A2: Provincial concentration of collection volume



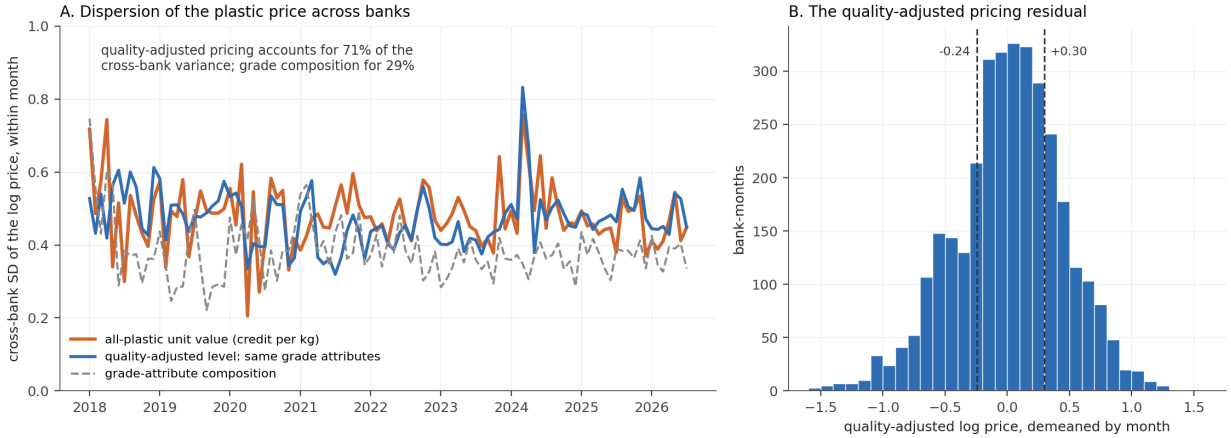
Notes: Left panel: the ten provinces with the highest total kilograms collected, January 2018–July 2026. Right panel: Jawa Barat's share of national monthly kilograms collected over the same window, 3-month centered rolling average; the dotted reference line marks a 50 percent share.

Figure A3: The upstream shifters and the exposure that turns them into bank-level variation



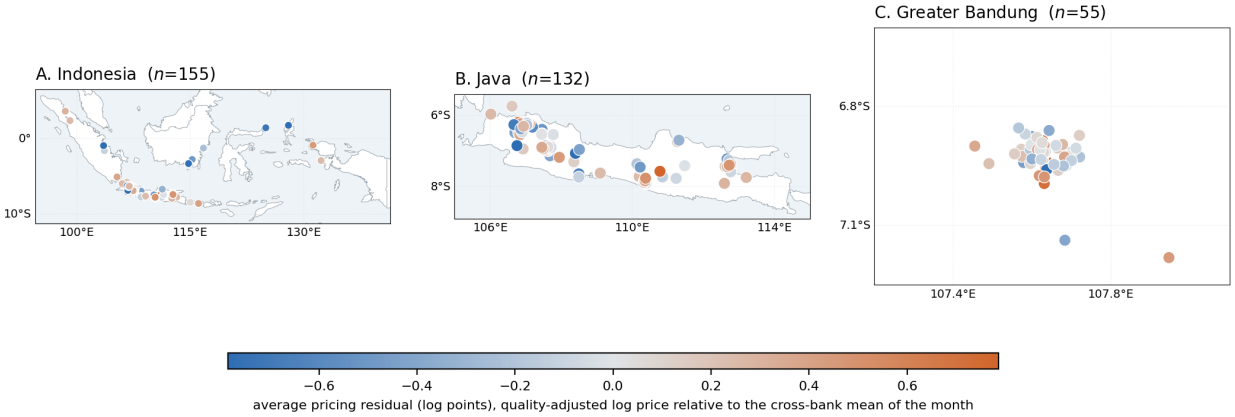
Notes: Panel A plots the two upstream price shifters used in Section 4, each expressed in real rupiah and indexed to its 2018 average: the unit value of Indonesian PET polymer imports and the U.S. producer-price index for plastics materials and resins. The shaded band marks the January 2022 to July 2026 estimation window. The logs of the two series correlate 0.29 over this period, so they are not two versions of the same shifter, which is why both are reported. Panel B plots the distribution of baseline plastic exposure ϖ_b, the plastic share of a bank's 2018–2021 transaction credit, with the banks that enter the IV estimation sample overlaid.

Figure A4: Dispersion of the plastic price across banks: grade mix versus pricing



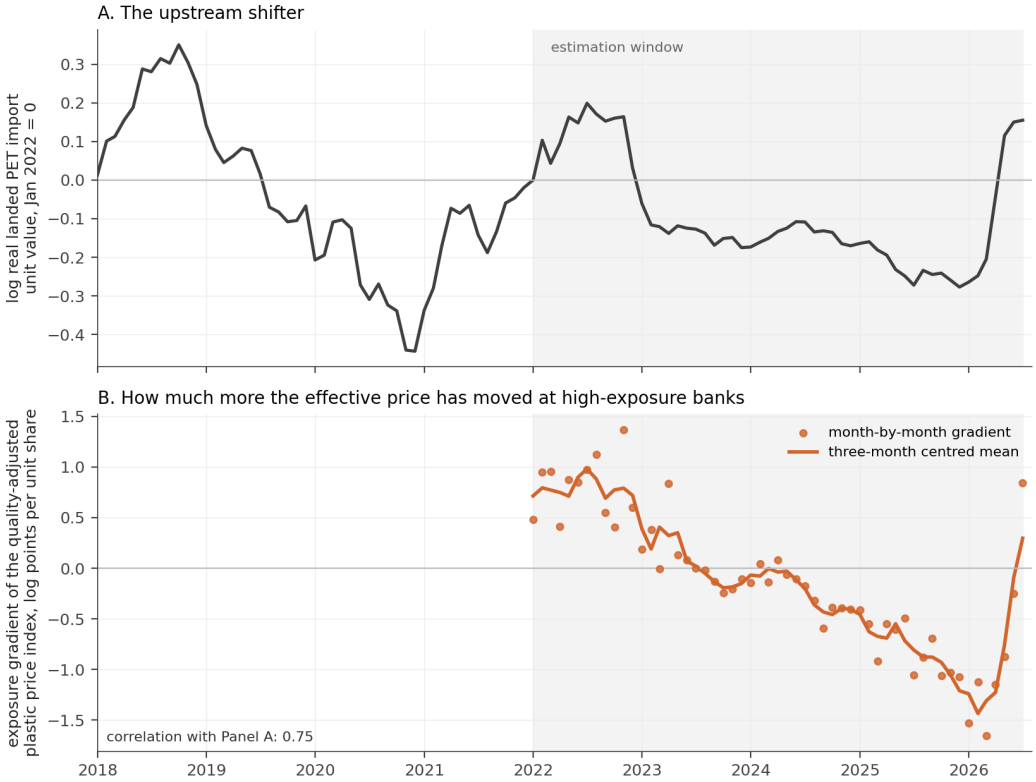
Notes: Panel A: for each month, the cross-bank standard deviation of the log bank–month plastic unit value (credit per kilogram) and of the quality-adjusted price level, the bank–month effect of the hedonic model of Section 3.2, which holds the full set of grade attributes fixed, together with the part attributable to grade-attribute composition. Grade composition accounts for 29 percent of the cross-bank variance in the unit value and pricing for the remaining 71 percent. Panel B: the distribution of the quality-adjusted log price level across bank–months, demeaned by month, with the 25th and 75th percentiles marked. Sample: January 2018–July 2026.

Figure A5: Quality-adjusted pricing residuals across banks in Greater Bandung



Notes: Each marker is a bank in Greater Bandung, shaded by its average quality-adjusted pricing residual over the sample, the bank-level mean of the hedonic bank–month effect net of the month mean. The standard deviation of this residual across the 55 banks shown is 0.38 log points, against 0.42 across the whole network. Sample: January 2018–July 2026.

Figure A6: The supply first stage over time



Notes: Panel A: the lead upstream shifter, the log real-rupiah unit value of Indonesian PET polymer imports (BPS, HS 3907.61; log USD + log IDR/USD − log CPI), normalized to January 2022. The shaded band is the estimation window. Panel B: the exposure gradient of the quality-adjusted price index, month by month: within each month, the slope of the bank-demeaned $\log P_{bt}^{H}$ on baseline plastic exposure across the banks of the 578 bank–month estimation sample (between five and sixteen banks per month), with a three-month centered mean. Under the first-stage specification of Section 4.1.1, $\log P_{bt}^{H} = \alpha_b + \delta_t + \beta\, s_b v_t + \varepsilon_{bt}$, this slope equals βv_t up to a constant, so Panel B is the variation the first stage uses. Its correlation with Panel A over the window is 0.75, and 0.65 after a linear trend is removed from both.

C Additional Tables

Table A1: What the hedonic model prices: attribute premia within a resin group

	PET	HDPE	LDPE	PP
Clean (*bersih*)	+13%	−47%	−14%	+19%
	(0.12)	(0.17)	(0.31)	(0.22)
Dirty (*kotor*)	−22%	−58%	−62%	−13%
	(0.13)	(0.17)	(0.27)	(0.17)
Clear (*bening*)	+9%	+27%	−13%	+46%
	(0.07)	(0.21)	(0.44)	(0.29)
Colored (*warna*)	−26%	+14%	−40%	−38%
	(0.13)	(0.18)	(0.31)	(0.21)
White (*putih*)	+19%	+2%	−47%	+8%
	(0.13)	(0.13)	(0.25)	(0.25)
Black (*hitam*)	−80%	−46%	+3%	−57%
	(0.36)	(0.18)	(0.35)	(0.21)
Mixed (*campur*)	−30%	−4%	−89%	−29%
	(0.10)	(0.13)	(0.37)	(0.20)
Within-R^2 of the hedonic model	0.83	0.83	0.90	0.88

Notes: Each cell is the implied price premium $\exp(\gamma) - 1$ for a grade carrying that attribute relative to one that does not, within the same resin group: the attribute's main effect plus its interaction with the resin group, with the standard error of the combined term below. Estimated on every plastic line–month in the records, 26,135 line–months on 3,811 lines at 462 banks from 2013 onward, with bank × month effects and credit weights, standard errors clustered by line. These are the objects a bank creates when it splits a coarse bin: dirty, mixed, colored and black grades are discounted and clean and clear grades carry premia within PET and polypropylene, with exceptions in HDPE and LDPE, so refining the partition raises what a seller who sorts to the finer grade receives. Individual cells are imprecise and the robust object is the ordering, not any single premium. The final row reports the within-R^2 of the hedonic model estimated separately on each resin group; the fit is poor for the residual OTHER group (0.27) and for EPS (0.25), which is why no composition claim in the paper rests on either.

Table A2: The supply-IV estimation sample against the rest of the network

	IV sample	Other banks
Banks	21	445
Active plastic months (median)	49	2
Months between first and last transaction (median)	66	3
Monthly plastic kilograms (median)	62	36
Monthly plastic sellers (median)	13	4
Plastic grades on the menu (median)	4	2
Share in Jawa Barat	57%	42%
Share of all plastic kilograms in the network	66%	34%

Notes: The supply IV of Section 4 is estimated on bank–months with a plastic price index, positive plastic purchases from individual depositors and a baseline exposure window, between January 2022 and July 2026, after fixed-effect singleton removal. The column headed *IV sample* describes the banks that survive those requirements; *other banks* describes every other bank with plastic purchases from individual depositors. They differ in persistence rather than scale: the IV banks report continuously, which is what a monthly panel with bank and calendar-month effects requires, and they are not the largest or the most plastic-intensive banks in the network. Their median baseline plastic exposure is 0.294 against 0.385 for the rest, so the estimating sample is if anything less plastic-intensive than the network it is drawn from.

Table A3: Reduced-form diagnostics for the supply IV

	First stage (price index)	Reduced form (log kilograms)
Lead placebo: instrument at t and at $t+k$ entered together		
Instrument at t (with t+6 included)	3.76 (1.24)	—
Lead at t+6	0.26 (0.56) [$p = 0.64$]	—
Lead at t+12	0.33 (0.43) [$p = 0.46$]	—
Exposure-specific trends		
Adding exposure × linear trend	1.67 (1.41)	0.79 (2.89)
Adding exposure × year effects	2.56 (0.96)	—
Band-passed instrument (36 mo.)	3.00** (1.28)	5.39 (3.13)
Falsification		
Paper exposure × PET price	−0.72 (1.23)	—
2018–19 exposure, 2020–21 outcome	0.94 (0.56)	—
Exposure balance (joint F on log kilograms, log depositors, Jawa Barat, bank age)		
Wald p	0.67	
Sample splits		
Excluding Jawa Barat	—	8.07 (1.97)
2SLS elasticity	2.79 (1.37)	
Exposure entered continuously (within-month percentile rank)		
Levels instrument	0.13 (0.07)	—
Band-passed instrument (36 mo.)	1.72*** (0.52)	2.69* (1.39)

Notes: All rows use the PET import price unless stated, on the 578 bank–month estimation sample (shorter where a lead, a band-pass window or a split requires it), with bank and calendar-month effects and two-way-clustered standard errors. Leads load when entered alone because the shifter is highly persistent; the rows above enter the contemporaneous instrument and the lead together. The resin PPI behaves the same way under the trend control, its first stage falling to 2.57 (2.83); in a horse race among the three components of either shifter the deflator term dominates, at 20.11 (6.07) for the resin index, and the negative log consumer price index has a correlation of −0.98 with a linear trend over the estimation window, so the trend control and the deflator component are the same test. Bracketed figures are p-values. Stars: 10, 5 and 1 percent.

Table A4: Crude oil and the upstream shifter

	First stage (price index)	Reduced form (log kilograms)	2SLS (elasticity)
Panel A. The PET import price regressed on crude (time series, 55 months)			
Levels, real	elasticity 0.665 (0.106), $R^2 = 0.581$		
Nominal US dollars, no deflator	elasticity 0.672 (0.113), $R^2 = 0.591$		
Band-passed, 24 months	elasticity 0.486 (0.062), $R^2 = 0.459$		
Band-passed, 36 months	elasticity 0.557 (0.060), $R^2 = 0.554$		
Crude at lags 0, 3 and 6 jointly	$R^2 = 0.827$		
Panel B. Crude used as the instrument			
Real crude price	2.52*** (0.82)	3.19 (1.59)	1.26* (0.65)
band-passed, 36 months	1.95*** (0.78)	3.27 (1.56)	1.68* (0.95)
lagged three months	2.80*** (0.89)	3.83* (1.55)	1.37* (0.68)
lagged six months	1.87** (0.99)	−0.13 (1.78)	−0.07 (0.96)
Nominal US dollar crude, no deflator	2.67*** (0.80)	3.58* (1.47)	1.34** (0.60)
Memo: PET import price (Table 5)	3.62** (1.08)	4.26 (2.26)	1.18 (0.74)
Panel C. Both shifters entered in the same first stage			
	PET or resin	Crude	Correlation
PET import price and crude, levels	3.73*** (1.31)	−0.13 (0.66)	0.999
band-passed, 24 months	3.14* (1.56)	0.09 (0.89)	0.694
band-passed, 36 months	3.00** (1.37)	0.20 (0.74)	0.770
Resin PPI and crude, levels	4.78* (2.33)	0.99 (0.65)	—
Panel D. Crude projected out, and crude as a control			
	First stage	Reduced form	2SLS
PET import price, crude projected out	4.05*** (1.38)	4.34 (3.40)	1.07 (0.98)
crude, exchange rate and inflation projected out	3.38** (1.49)	5.51 (4.12)	1.63 (1.63)
crude projected out, band-passed 24 months	1.81 (1.44)	1.52 (2.81)	—
crude projected out, band-passed 36 months	2.53* (1.35)	1.91 (2.88)	—
Crude, PET import price projected out	−0.34 (0.95)	0.15 (2.27)	—
PET import price, controlling for crude	3.73*** (1.31)	3.93 (3.38)	1.05 (1.03)
controlling for crude at lags 0, 3 and 6	2.52 (1.46)	2.21 (6.48)	0.88 (2.73)

Notes: Panel A is a time-series regression of the log real PET import unit value on the log real Brent crude price over the 55-month estimation window, with Newey–West standard errors at six lags; the nominal row deflates neither series. Panels B to D use the 578 bank–month estimation sample of Table 5, with bank and calendar-month effects and two-way-clustered standard errors, and instrument the quality-adjusted price index with baseline plastic exposure interacted with the series named. Band-passing replaces a series with its deviation from a centered moving average. Projecting out is a residual from a regression fitted on the estimation window. In Panel C the correlation column is between the two instruments on the estimation panel: in levels it is 0.999, because the interaction is dominated by the cross-sectional exposure share, so the levels row is close to collinear and the band-passed rows carry the comparison. No 2SLS elasticity is reported where the first stage does not separate from zero. Stars are taken from the Webb wild p-value where one is available, which is every first-stage and reduced-form row in Panels B and D, and from the two-way-clustered p-value elsewhere, which is every 2SLS row, all of Panel C, and the two control rows of Panel D. Stars: 10, 5 and 1 percent.

Table A5: Placebo materials: the upstream plastic shifter and the prices of paper, metal, and glass

	Plastic	Paper	Metal	Glass
Panel A. Quality-adjusted price index of the material (first stage)				
Exposure × PET import price	3.62** (1.08)	1.45 (1.00)	−4.43** (1.58)	0.75 (1.32)
Webb wild *p*-value	0.011	0.177	0.004	0.586
continuing-grade component	1.61 (0.88)	0.39 (0.79)	−0.25 (0.45)	0.82 (0.50)
new-grade component	2.01** (0.79)	1.06* (0.55)	−4.18** (1.49)	−0.08 (1.33)
first-ever grades	1.14 (0.65)	0.03 (0.18)	−0.75 (0.59)	−0.30 (0.62)
re-listings after a silence of over six months	0.87 (0.63)	1.03* (0.56)	−3.43** (1.24)	0.22 (0.78)
Index scoring only first-ever grades as new	2.75** (1.03)	0.42 (0.82)	−1.00 (0.87)	0.52 (0.66)
Panel B. Log kilograms of the material from individual depositors (reduced form)				
Exposure × PET import price	4.26 (2.26)	5.61** (2.59)	3.88 (2.48)	−4.82 (2.80)
Webb wild *p*-value	0.174	0.048	0.147	0.204
picker depositors	8.39** (2.59)	7.23* (3.23)	5.21 (5.62)	2.57 (6.19)
household depositors	2.32 (1.98)	3.92* (1.93)	4.91 (3.49)	−5.98 (2.95)
Bank–months	578	588	500	421
Banks	21	22	16	16

Notes: Each column applies the design of Table 5 to a different material: the endogenous price is the quality-adjusted index of Section 3.2 built on that material's lines under the same six-month reappearance rule, the outcome is that material's kilograms from individual depositors, and the instrument is baseline *plastic* exposure interacted with the real-rupiah PET import price throughout. Each column uses the bank–months in which the material's index and kilograms are both observed over January 2022 to July 2026, after singleton removal; restricting the paper, metal, and glass columns to the 578 bank–months of the plastic column leaves every price coefficient within 0.2 of the value shown. Bank and calendar-month fixed effects; standard errors two-way clustered by bank and month; stars from the null-imposed Webb wild-cluster *p*-value: 10, 5 and 1 percent. The by-type rows of Panel B estimate the reduced form on the material's kilograms from picker and household depositors separately, on the plastic column's banks and window. The new-grade component is split into grades transacting for the first time and grades returning after a silence of more than six months, which the rule scores as new. Paper, metal, and glass menus are short lists of distinct products rather than grades of one material, so a returning line is priced against one or two other products and its re-listing term is not a quality premium. This is why the metal first stage is negative: it is a re-listing term (first-ever grades −0.75 (0.59); continuing-grade component −0.25 (0.45)), 71 percent of its movement comes from such returns, and it is −2.90 (2.23) excluding 2022, −2.98 (1.48) under a twelve-month rule, −1.39 (0.95) under a thirty-six-month rule, and −1.00 (0.87) when only first-ever grades count as new; leaving any one bank out keeps it between −3.4 and −6.0. For plastic the first-ever and re-listing parts are both positive, 1.14 (0.65) and 0.87 (0.63). The metal column is starred from the clustered *p*-value because on its 16 lumpy clusters the wild bootstrap understates the variance (clustered $p = 0.23$ against wild $p = 0.006$ on the band-passed cells), an inversion that appears for no other material. The glass continuing-grade coefficient has a wild *p*-value of 0.138; about 40 percent of glass repricings coincide with a revision of the bank's plastic schedule.

Table A6: Repricing a continuing grade versus changing the menu

	−12 to −7	0–5	6–11	12–17	18–23
Panel A. Unconditional: Where movement in the effective price comes from					
Share of the variance of the monthly index change, continuing grades			57.3%		
new grades			41.1%		
Share of mean absolute monthly movement, new grades			51.9%		
Bank–months containing a new line			31.7%		
Bank–months containing a repriced continuing line			24.3%		
Median bank: share of its cumulative change from new grades			100.0%		
Panel B. An upstream value shock: First stage by price definition					
Unit value (no quality adjustment)			2.84** (0.93)		
Menu mean of line prices (equal-weighted across grades)			2.13** (0.86)		
Continuing-grades index (matched model; omits new grades)			1.80 (0.97)		
Continuing-grade component (repricing margin alone)			1.61 (0.88)		
New-grade component (menu margin alone)			2.01** (0.79)		
Hedonic chained index (holds quality fixed; prices new grades)			3.62** (1.08)		
Panel C. A rival opens nearby: The same two margins					
Posted price of a grade already on the menu	0.024	−0.025*	−0.034*	−0.040	−0.060
	(0.028)	(0.015)	(0.018)	(0.030)	(0.041)
Quality-adjusted index	−0.004	−0.020	−0.029	0.051	0.064
	(0.041)	(0.020)	(0.052)	(0.049)	(0.093)
Continuing-grade component	0.019	−0.012	−0.011	0.028	0.029
	(0.042)	(0.017)	(0.034)	(0.031)	(0.038)
New-grade component	−0.023	−0.007	−0.018	0.023	0.035
	(0.021)	(0.010)	(0.024)	(0.033)	(0.083)
Any within-resin refinement in the month (0/1)	−0.028	0.067***	0.113***	0.152***	0.191***
	(0.041)	(0.022)	(0.033)	(0.038)	(0.067)

Notes: Panel A is descriptive, over 2,957 monthly changes in the chained index at 205 banks; it decomposes movement in the index into the repricing of lines the bank already posts and lines new to its menu. Panel B estimates the exposure first stage of Section 4.1.1 for six price definitions on one common sample, the 578 bank–months and 21 banks of Table 5; standard errors two-way clustered by bank and month, stars from the null-imposed Webb wild-cluster p-value. Panel C is the entry design of Table 11, pooled over all treated incumbents: the first row is the stacked line fixed-effects regression on the posted price of a plastic grade over individual depositors, the remaining rows are stacked OLS on the bank–month index and its components, all with bank-age fixed effects, a reference period of months −6 to −1, and standard errors clustered on the entrant. The index follows the six-month reappearance rule of Section 3.2, and the refinement row is the indicator of Table 11. The three panels make one point from three directions: repricing a continuing line is economically small and statistically weak, and the movement in what a seller is effectively paid arrives through the grades the bank posts, as the premium on lines entering the menu under the upstream shock and as the refinement of the menu itself after entry. Stars: 10, 5 and 1 percent.

Table A7: Robustness of the supply IV

	PET import price	Virgin-resin PPI
First stage: quality-adjusted plastic price index	3.62** (1.08)	6.48** (2.30)
Three-month reappearance rule	4.24*** (1.34)	7.43*** (2.69)
Nine-month reappearance rule	3.64*** (1.08)	6.50*** (2.29)
Pre-rule index (every resupply scored as new)	4.67*** (1.34)	8.03*** (2.67)
2SLS supply elasticity	1.18 (0.74)	1.26 (1.03)
Three-month reappearance rule	1.00 (0.59)	1.10 (0.84)
Nine-month reappearance rule	1.17 (0.73)	1.26 (1.03)
Pre-rule index	0.91 (0.52)	1.02 (0.73)
Weak-instrument-robust inference		
Two-way-clustered Anderson–Rubin 95% set	[0.00, 3.75]	[−0.40, 5.95]
Webb wild Anderson–Rubin 95% set	[−0.45, 5.25]	[−1.00, 7.65]
Webb wild p, first stage	0.011	0.016
Webb wild p, reduced form	0.174	0.353
Webb wild p, new-grade component	0.035	0.034
Construction and influence (PET import price)		
Seven index constructions, first stage	[2.82, 4.19]	—
Seven index constructions, 2SLS	[1.02, 1.51]	—
Leave one bank out, first stage	[3.17, 4.93]	—
Leave one bank out, 2SLS	[0.51, 1.53]	—
Excluding 2022, first stage	2.62** (0.98)	1.84 (2.40)
Paper-index placebo, first stage	1.34 (0.98)	2.63 (2.09)
Bank–months	578	578
Banks	21	21
Robustness of the picker − household gap		
Gap, full sample	1.337** (0.628), $p = 0.046$	
Leave one bank out	[1.136, 1.898] over 21 banks; significant at 5% in 10 of 21	

Notes: The instrument is the bank's 2018–2021 plastic-credit share interacted with the indicated real-rupiah upstream shifter; the endogenous variable is the quality-adjusted plastic price index and the outcome is log plastic kilograms bought from individual depositors, on the estimation sample of Table 5 throughout. Parentheses contain standard errors two-way clustered by bank and month; stars come from the Webb wild-cluster p-value where one is reported and from the clustered p-value in the three-month, nine-month and excluding-2022 rows: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Wild inference uses 9,999 null-imposed Webb draws clustered by bank with the fixed effects re-estimated in every replication. Brackets in the construction block are coefficient ranges, not confidence sets; the Anderson–Rubin rows are confidence sets. The seven constructions vary the attribute vocabulary threshold, the index weighting, the cap on imputed new-grade gaps, and the code and suffix attributes, each under the six-month reappearance rule. The three- and nine-month rows re-chain the index with a line scored as new after a silence of more than three or nine months rather than six; the reduced form is the same in every row. The pre-rule rows report an index without the reappearance rule, which scored every grade transacting after a silent recorded month as new and so counts resupplies as menu changes (Section 3.2). The paper-index placebo applies the same exposure and shifter to a hedonic paper-price index built identically on its own sample of paper-transacting banks; its within-R^2 of 0.239 is far below the 0.813 of the plastic index, so its components are not interpreted. Appendix Table A5 reports the paper, metal and glass indices on their own samples.

Table A8: Cycle-identified estimates of the supply IV

	PET import price			Virgin-resin PPI		
	Levels	Band-passed		Levels	Band-passed	
		24 mo.	36 mo.		24 mo.	36 mo.
First stage: plastic-price index	3.62** (1.08)	3.74** (1.61)	3.00** (1.28)	6.48** (2.30)	7.10** (3.96)	6.91** (3.09)
Wald F	11.2	5.4	5.5	7.9	3.2	5.0
Effective F	11.5	6.1	5.7	8.9	4.3	5.2
Webb wild p	0.011	0.023	0.033	0.013	0.035	0.032
Reduced form: log kilograms	4.26 (2.26)	7.30 (3.29)	5.39 (3.13)	8.19 (5.59)	19.87** (6.15)	16.87 (6.32)
Webb wild p	0.174	0.124	0.277	0.353	0.038	0.148
2SLS supply elasticity	1.18 (0.74)	1.95 (1.19)	1.80 (1.52)	1.26 (1.03)	2.80 (1.61)	2.44 (1.73)
Clustered AR 95% set	[0.00, 3.75]	[0.30, 12.60]	[−0.10, 16.60]	[−0.40, 5.95]	[1.00, +∞)	[0.50, +∞)
Wild AR 95% set	[−0.45, 5.25]	[−0.60, 17.90]	[−1.00, +∞)	[−1.00, 7.65]	[0.20, +∞)	[−0.50, +∞)
Bank–months	578	491	429	578	491	429
Banks	21	21	21	21	21	21

Notes: The instrument is the bank's 2018–2021 plastic-credit share interacted with the indicated upstream shifter; bank and calendar-month fixed effects; standard errors two-way clustered by bank and month. *Levels* uses the real-rupiah shifter as in Table 5. *Band-passed* replaces it with its deviation from a centered moving average of the stated width, which removes the low-frequency component and leaves an instrument that is 93 to 99 percent orthogonal to an exposure-specific linear trend; the centered window is required to be complete, which shortens the sample at both ends. Because the band-passed columns are estimated on those shorter samples, the levels instrument re-estimated on the same samples gives first stages of 3.62 (1.22) and 3.30 (1.20) for the PET import price and 7.29 (2.77) and 6.71 (2.58) for the resin PPI, with 2SLS elasticities of 1.29, 1.25, 1.44 and 1.43, so the shorter sample accounts for part but not all of the difference between the levels and band-passed elasticities. The effective F is the Montiel Olea–Pflueger statistic with the bank-clustered variance. Stars are taken from the null-imposed Webb wild p-value with 9,999 draws clustered by bank and the fixed effects re-estimated in each replication: 10, 5 and 1 percent. The resin 24-month column has a Wald F of 3.2 and an Anderson–Rubin set unbounded above, so its point estimate should not be read as an elasticity.

Table A9: Rival entry: specification, comparison group, sample rule and cohort

Specification	Incumbents	Plastic kg −12 to −7 m	Plastic kg 6–11 m	Picker-delivered plastic, 6–11 m	Household-delivered plastic, 6–11 m	Picker sellers, 6–11 m	Household sellers, 6–11 m	Index 12–17 m	New-grade comp., 12–17 m	Plastic price ≤5 km, 12–17 m
Headline: roster-age fixed effects, reference months −6 to −1	97	−0.020	−0.455**	−0.764***	0.019	−0.192**	0.039	0.051	0.023	−0.056**
		(0.140) [33]	(0.208) [88]	(0.137) [88]	(0.133) [88]	(0.088) [88]	(0.072) [88]	(0.049) [56]	(0.033) [56]	(0.022) [19]
Reference month −1 only	97	0.094	−0.344	−0.659***	0.054	−0.211*	0.142	0.049	0.024	−0.070**
		(0.141) [33]	(0.228) [88]	(0.147) [88]	(0.147) [88]	(0.110) [88]	(0.089) [88]	(0.051) [56]	(0.031) [56]	(0.024) [19]
No bank-age fixed effects	97	−0.051	−0.667**	−1.039***	0.206**	−0.146*	0.146**	0.043	0.009	−0.043**
		(0.125) [33]	(0.322) [88]	(0.269) [88]	(0.102) [88]	(0.079) [88]	(0.069) [88]	(0.041) [56]	(0.032) [56]	(0.018) [19]
No bank-age fixed effects, reference month −1	97	0.006	−0.611*	−0.998***	0.230**	−0.138	0.234***	0.047	0.017	−0.053**
		(0.136) [33]	(0.337) [88]	(0.310) [88]	(0.114) [88]	(0.102) [88]	(0.079) [88]	(0.043) [56]	(0.029) [56]	(0.021) [19]
Six-month age bins	97	−0.116	−0.787*	−1.441***	0.077	−0.230	0.070	0.043	0.016	−0.072**
		(0.167) [33]	(0.418) [88]	(0.307) [88]	(0.140) [88]	(0.161) [88]	(0.077) [88]	(0.051) [56]	(0.036) [56]	(0.030) [19]
Age measured from the first recorded month	97	0.539***	−0.621***	−0.926***	0.120	−0.182	0.166*	0.069	0.028	−0.056
		(0.140) [33]	(0.184) [88]	(0.166) [88]	(0.099) [88]	(0.132) [88]	(0.091) [88]	(0.042) [56]	(0.030) [56]	(0.038) [19]
Comparison banks within ±6 months of age	71	0.267	−1.355**	−2.248***	−0.080	−0.074	0.054	0.143**	0.029	−0.075**
		(0.337) [24]	(0.636) [62]	(0.421) [62]	(0.217) [62]	(0.155) [62]	(0.125) [62]	(0.061) [37]	(0.043) [37]	(0.033) [16]
Three nearest comparison banks on age and size	97	−0.067	−1.160**	–	0.190	–	0.276***	0.045	0.026	−0.012
		(0.222) [33]	(0.509) [88]		(0.147) [88]		(0.082) [88]	(0.051) [56]	(0.035) [56]	(0.046) [19]
Comparison banks reweighted on age and size	97	−0.148	−0.934**	−1.301***	0.193	−0.099	0.180**	0.039	0.002	−0.052**
		(0.133) [33]	(0.413) [88]	(0.357) [88]	(0.131) [88]	(0.105) [88]	(0.083) [88]	(0.048) [56]	(0.033) [56]	(0.018) [19]
Qualification at month −1 only	143	−0.132	−0.452**	−0.776***	0.034	−0.240***	0.127	0.054	0.015	−0.079**
		(0.156) [41]	(0.209) [88]	(0.128) [88]	(0.128) [88]	(0.088) [88]	(0.079) [88]	(0.046) [58]	(0.032) [58]	(0.033) [20]
Also a record in months 12–23	71	0.379	0.106	−0.157	0.146	−0.131	0.124	0.049	0.027	−0.059**
		(0.271) [28]	(0.199) [62]	(0.419) [62]	(0.134) [62]	(0.095) [62]	(0.077) [62]	(0.046) [52]	(0.033) [52]	(0.023) [19]
Bandung March-2019 cohort excluded	71	−0.020	−0.490**	−0.786***	0.030	−0.211**	0.077	0.082*	0.031	−0.060*
		(0.141) [33]	(0.203) [62]	(0.129) [62]	(0.155) [62]	(0.090) [62]	(0.087) [62]	(0.044) [41]	(0.034) [41]	(0.029) [9]
All accounts in quantities	97	−0.087	−0.396**	–	–	−0.192**	0.039	0.051	0.023	−0.056**
		(0.139) [33]	(0.187) [88]			(0.088) [88]	(0.072) [88]	(0.049) [56]	(0.033) [56]	(0.022) [19]

Notes: Each row re-estimates the design of Tables 9–11 under one change; the first row is the headline. Rows 2–4 vary the reference period (month −1 alone) and the bank-age fixed effects; row 5 uses six-month age bins; row 6 measures age from the bank's first recorded month rather than the roster date. Rows 7–9 change the comparison group and use no age fixed effects: banks within six months of the incumbent's roster-clock age at the opening, the construction of Appendix Table A10 (71 incumbents have at least one such bank; the Poisson fit for the picker outcomes under this group drops fixed-effect groups whose outcome is identically zero and uses a looser demeaning tolerance, without which it does not converge); the incumbent's three nearest comparison banks in standardized first-record age and log pre-opening plastic kilograms (month −1 reference); and all comparison banks reweighted by entropy balancing on roster-clock age and log pre-opening plastic kilograms where feasible (33 of 97 stacks) and by a Gaussian kernel on their distance otherwise. Rows 7 and 9 use the headline reference period. Rows 10–11 vary the sample rule; row 12 excludes the March 2019 Bandung opening, which treats 34 incumbents at once, 26 of them in the headline sample; row 13 adds unit-bank and institutional accounts to the quantities (plastic by seller type is not defined for that universe). The index columns follow the six-month reappearance rule of Section 3.2. Brackets give the number of incumbents identifying the cell. Standard errors, in parentheses, are clustered by the entrant. Stars: 10, 5 and 1 percent.

Table A10: Rival entry: treated incumbents and their comparison banks before the opening

		Pooled comparison group (headline)			Age-matched		Reweighted	
	Treated	Comparison	Difference	Norm. diff.	Comparison	Difference	Comparison	Difference
Age at the opening, roster clock (months)	25.60	39.47	−7.00	−0.35	14.96	0.25	34.49	−1.19
			(5.48)			(0.43)		(2.29)
Age at the opening, months since first record	8.92	19.45	−7.01***	−0.54	11.05	−2.52	16.58	−4.81**
			(2.33)			(1.79)		(1.98)
Months recorded before the opening	6.26	13.91	−5.32***	−0.69	9.45	−3.35**	12.61	−4.33***
			(1.61)			(1.61)		(1.53)
Plastic kg per month	488	662	−56.91	−0.08	432	159.10	528	81.57
			(267.67)			(278.68)		(259.00)
Kilograms per month, all materials	1,170	1,577	−232.31	−0.10	908	514.68	1,312	99.05
			(558.96)			(537.89)		(518.05)
Depositors per month	26.29	23.42	−0.95	0.11	20.62	5.54	23.54	0.40
			(3.10)			(5.05)		(3.74)
Picker depositors per month	2.37	4.52	−2.76	−0.30	3.50	−1.19	3.73	−1.51
			(1.70)			(1.25)		(1.31)
Household depositors per month	22.39	17.20	2.16	0.24	15.98	6.17	17.99	2.07
			(2.96)			(4.75)		(3.72)
Picker share of plastic kg	0.22	0.46	−0.24***	−0.70	0.43	−0.22***	0.40	−0.19**
			(0.08)			(0.08)		(0.07)
Plastic grades on the menu	8.59	12.70	−4.33***	−0.48	11.38	−2.67***	12.18	−3.42***
			(1.28)			(0.94)		(0.86)
Mean log posted plastic price	7.23	7.11	0.11	0.31	7.16	0.03	7.11	0.08
			(0.07)			(0.06)		(0.05)
Plastic growth, months −6..−1 over −12..−7 (log points)	0.18	0.30	−0.03	−0.09	0.26	0.07	0.34	0.01
			(0.24)			(0.20)		(0.25)
Located in Java	0.72	0.56	0.13	0.33	0.70	−0.13	0.60	0.05
			(0.12)			(0.11)		(0.10)
Central (induk) bank	0.02	0.08	−0.04**	−0.28	0.02	−0.01	0.05	−0.02
			(0.02)			(0.02)		(0.02)
Incumbents / comparison banks	97	103			44		103	
Comparison banks per stack		9.7			2.8		9.7	
Joint test, p-value			0.000			–		0.000

Notes: Characteristics of the 97 treated incumbents and of the comparison banks in their stacks, measured over the six months before the opening (recorded months) or at the opening. Treated means are over incumbents; comparison means are over comparison bank–stack pairs, so a bank that serves in several stacks counts in each. The pooled group is the headline comparison group of Table 9; the age-matched group keeps comparison banks within six months of the incumbent's roster-clock age at the opening (71 incumbents with at least one); the reweighted group applies the entropy-balancing or kernel weights of Appendix Table A9, row 9. Differences are the coefficient on a treated indicator in a stacked regression with stack fixed effects, weighted as the estimator is, with standard errors in parentheses clustered by the entrant; the normalized difference is the difference in means divided by the square root of the average of the two variances. The joint test is a Wald test, clustered by the entrant, that the standardized characteristics do not predict treatment in a linear probability regression with stack fixed effects; it is not computed for the age-matched group, which has too few comparison banks per stack. Plastic growth is the log change in mean monthly plastic kilograms from months −12 to −7 to months −6 to −1, for banks recording in both. Stars: 10, 5 and 1 percent.

Table A11: Recording after the opening: treated incumbents and their comparison banks

	All treated		Rival within 5 km		Rival beyond 5 km	
Months	Treated	Comparison	Treated	Comparison	Treated	Comparison
0–5	0.94	0.93	0.97	0.93	0.93	0.93
6–11	1.00	1.00	1.00	1.00	1.00	1.00
12–17	0.70	0.66	0.80	0.61	0.66	0.68
18–23	0.47	0.34	0.39	0.23	0.51	0.41

Notes: Share of banks with at least one recorded transaction in the stated months after the opening, among treated incumbents and among the comparison banks qualifying at the same opening under the headline rule (Table 9); comparison shares are averaged over openings. Months 6–11 are one by construction.

Table A12: Every outcome by incumbent type: picker-intensive and household-intensive incumbents

	Picker-intensive incumbents					Household-intensive incumbents				
	Months relative to the opening									
	−12 to −7	0–5	6–11	12–17	18–23	−12 to −7	0–5	6–11	12–17	18–23
Kilograms, all materials	0.207**	−0.048	−0.228*	0.094	−0.183	0.207	0.250	0.059	0.125	−0.083
	(0.082)	(0.047)	(0.133)	(0.166)	(0.147)	(0.293)	(0.238)	(0.146)	(0.223)	(0.240)
Plastic	−0.085	−0.064	−0.625***	−0.188	−0.350	−0.260	0.470	0.127	−0.280	−0.276
	(0.159)	(0.049)	(0.170)	(0.207)	(0.216)	(0.351)	(0.420)	(0.168)	(0.218)	(0.256)
delivered by picker sellers	−0.182	−0.097**	−0.846***	−0.216	−0.369	−1.388*	0.928**	0.796	0.922	1.159*
	(0.194)	(0.043)	(0.110)	(0.206)	(0.294)	(0.819)	(0.438)	(0.568)	(0.780)	(0.661)
delivered by household sellers	0.004	−0.116	−0.097	−0.196	−0.337	0.097	0.562	0.155	−0.143	−0.423
	(0.145)	(0.095)	(0.178)	(0.293)	(0.251)	(0.333)	(0.490)	(0.172)	(0.228)	(0.268)
Paper	0.006	0.037	0.061	0.250*	0.064	0.288	−0.218***	−0.168	0.212	−0.170
	(0.175)	(0.074)	(0.076)	(0.151)	(0.254)	(0.304)	(0.061)	(0.188)	(0.365)	(0.387)
Glass	−0.511*	−0.013	0.183	0.354	0.502	−0.003	0.029	−0.403	0.005	−0.492
	(0.277)	(0.233)	(0.387)	(0.707)	(0.607)	(0.794)	(0.277)	(0.530)	(0.639)	(0.532)
Metal	0.341	−0.427	−0.512	−0.432	−1.219***	0.545	0.205	0.231	0.384	0.639
	(0.529)	(0.276)	(0.335)	(0.412)	(0.458)	(0.786)	(0.321)	(0.398)	(0.413)	(0.420)
Active sellers	−0.120	−0.074	0.022	0.054	0.042	−0.067	−0.149*	−0.144*	−0.119	−0.189
	(0.107)	(0.062)	(0.109)	(0.154)	(0.259)	(0.171)	(0.083)	(0.078)	(0.162)	(0.193)
Picker	−0.084	0.029	−0.242***	−0.143	−0.189	−0.154	0.146	0.183	0.557**	0.880***
	(0.125)	(0.069)	(0.087)	(0.224)	(0.300)	(0.764)	(0.225)	(0.208)	(0.235)	(0.314)
Household	−0.117	−0.066	0.119	0.108	0.123	−0.022	−0.089	−0.056	−0.072	−0.055
	(0.111)	(0.062)	(0.116)	(0.153)	(0.238)	(0.143)	(0.067)	(0.061)	(0.168)	(0.223)
Posted price of a grade on the menu: plastic	0.011	−0.034	−0.043	−0.057	−0.065	0.063	−0.011	−0.017	−0.008	−0.045
	(0.025)	(0.023)	(0.026)	(0.043)	(0.064)	(0.099)	(0.011)	(0.022)	(0.029)	(0.042)
Paper	0.004	−0.042**	−0.015	0.015	0.042	0.026	0.008	−0.023	−0.015	−0.052*
	(0.029)	(0.018)	(0.024)	(0.050)	(0.108)	(0.042)	(0.016)	(0.022)	(0.034)	(0.027)
Glass	−0.075	−0.074	−0.042	−0.014	−0.005	0.061	−0.039	−0.104	−0.016	0.024
	(0.045)	(0.053)	(0.064)	(0.060)	(0.062)	(0.092)	(0.038)	(0.136)	(0.135)	(0.132)
Metal	−0.013	−0.012	−0.007	0.008	0.028**	−0.012	−0.001	0.031**	0.064**	0.090***
	(0.045)	(0.012)	(0.018)	(0.015)	(0.011)	(0.020)	(0.014)	(0.011)	(0.025)	(0.029)
Quality-adjusted plastic index	0.008	−0.006	0.010	0.024	−0.066	−0.017	−0.038	−0.077	0.093	0.217*
	(0.020)	(0.017)	(0.029)	(0.050)	(0.053)	(0.105)	(0.034)	(0.085)	(0.089)	(0.123)
New-grade component	0.003	−0.006	−0.007	−0.004	−0.091***	−0.073	−0.015	−0.037	0.055	0.179
	(0.012)	(0.010)	(0.024)	(0.033)	(0.018)	(0.047)	(0.016)	(0.036)	(0.067)	(0.115)
Continuing-grade component	0.005	−0.001	0.017	0.028	0.024	0.056	−0.023	−0.040	0.038	0.038
	(0.018)	(0.013)	(0.014)	(0.030)	(0.054)	(0.105)	(0.028)	(0.058)	(0.056)	(0.056)
Any within-resin refinement in the month (0/1)	0.000	0.083**	0.168***	0.190***	0.157	−0.070	0.048*	0.055	0.109*	0.144**
	(0.049)	(0.037)	(0.049)	(0.056)	(0.121)	(0.067)	(0.028)	(0.041)	(0.054)	(0.058)
Number of refinements	0.311	2.073***	2.491***	2.629***	2.314**	−0.954	1.421**	0.273	1.265*	1.386
	(0.749)	(0.605)	(0.738)	(0.679)	(1.119)	(0.847)	(0.676)	(0.967)	(0.723)	(0.946)
Any coarser-or-equal addition in the month (0/1)	−0.052	0.003	0.101**	0.110**	−0.039	−0.049	−0.109	−0.020	−0.046	−0.074
	(0.035)	(0.041)	(0.047)	(0.052)	(0.117)	(0.042)	(0.068)	(0.028)	(0.037)	(0.049)
Incumbents in cell	21	43	42	26	12	12	46	44	29	16
Incumbents with a plastic line in cell	20	43	42	26	11	12	46	42	29	16

Notes: Design and type definition as in Tables 9–10: 47 picker-intensive and 48 household-intensive incumbents, split by the seller mix before the opening. Standard errors, in parentheses, are clustered by the entrant. Stars: 10, 5 and 1 percent. Second-year cells rest on the incumbents still recording; picker sellers at household-intensive incumbents average 0.16 per month before the opening, so their coefficients describe small absolute changes. The index follows the six-month reappearance rule and the refinement rows are defined as in Table 11.

Table A13: Quantities after a rival opens, by the distance of the rival

	Months relative to the opening				
	−12 to −7	0–5	6–11	12–17	18–23
Panel A. Rival within 5 kilometers (29)					
Kilograms, all materials	0.291***	0.103	−0.042	0.633***	−0.311
	(0.059)	(0.112)	(0.362)	(0.241)	(0.203)
Plastic	0.101	−0.045	−0.583	0.161	0.266
	(0.075)	(0.115)	(0.427)	(0.508)	(0.209)
delivered by picker sellers	0.112	−0.078	−0.941***	−0.459	0.122
	(0.098)	(0.119)	(0.252)	(0.659)	(0.138)
delivered by household sellers	0.088	0.011	−0.148	0.792***	−0.724***
	(0.468)	(0.090)	(0.191)	(0.300)	(0.154)
Paper	−0.234	−0.043	0.036	0.685***	−0.608***
	(0.314)	(0.133)	(0.158)	(0.118)	(0.231)
Glass	−0.361	0.688*	0.239	1.614***	0.740
	(0.676)	(0.393)	(0.448)	(0.390)	(0.893)
Metal	−1.240**	0.836***	0.390	0.715	−1.254***
	(0.550)	(0.299)	(0.566)	(1.100)	(0.345)
Active sellers	−0.357	−0.132	−0.022	0.353*	−0.023
	(0.295)	(0.095)	(0.077)	(0.196)	(0.337)
Picker	0.004	0.168	−0.239*	−0.142	−0.758***
	(0.187)	(0.146)	(0.136)	(0.508)	(0.180)
Household	−0.404	−0.064	0.031	0.314	0.093
	(0.419)	(0.075)	(0.067)	(0.209)	(0.193)
Incumbents in cell	10	28	26	20	9
Panel B. Rival between 5 and 15 kilometers (68)					
Kilograms, all materials	0.180	0.079	−0.110	−0.095	−0.172
	(0.156)	(0.100)	(0.139)	(0.145)	(0.113)
Plastic	0.179	0.093	−0.396	−0.293*	−0.412**
	(0.265)	(0.145)	(0.245)	(0.172)	(0.170)
delivered by picker sellers	0.630	−0.041	−0.675***	−0.138	−0.317
	(0.424)	(0.077)	(0.165)	(0.135)	(0.225)
delivered by household sellers	0.105	0.267	0.067	−0.196	−0.296
	(0.139)	(0.312)	(0.151)	(0.237)	(0.233)
Paper	0.079	0.074	0.030	0.068	0.014
	(0.179)	(0.110)	(0.081)	(0.181)	(0.161)
Glass	−0.390	−0.131	0.007	0.112	−0.056
	(0.275)	(0.162)	(0.347)	(0.535)	(0.474)
Metal	0.518	−0.493**	−0.370	−0.463	−0.414
	(0.455)	(0.236)	(0.292)	(0.313)	(0.308)
Active sellers	−0.058	−0.104	−0.035	−0.042	−0.052
	(0.104)	(0.066)	(0.088)	(0.122)	(0.176)
Picker	−0.101	−0.002	−0.158	−0.039	0.131
	(0.162)	(0.072)	(0.106)	(0.189)	(0.205)
Household	−0.055	−0.077	0.041	0.018	0.030
	(0.096)	(0.057)	(0.085)	(0.121)	(0.178)
Incumbents in cell	23	63	62	37	21

Notes: Design as in Table 9. The close-rival panel carries positive pre-period coefficients on plastic and total kilograms and few incumbents at long horizons; I report it and do not interpret the quantity response by distance (Section 4.2). Standard errors, in parentheses, are clustered by the entrant. Stars: 10, 5 and 1 percent.

Table A14: A second rival: the next opening after the first

	Months relative to the opening				
	−12 to −7	0–5	6–11	12–17	18–23
Panel A. Second opening within 15 kilometers (64 incumbents)					
Posted price of a plastic grade on the menu	−0.000	−0.036***	−0.039**	−0.071**	−0.038
	(0.035) [39]	(0.012) [59]	(0.018) [58]	(0.031) [25]	(0.061) [16]
Quality-adjusted plastic index	0.028	−0.057***	−0.038	−0.023	0.183
	(0.037) [39]	(0.017) [59]	(0.025) [58]	(0.075) [25]	(0.118) [16]
New-grade component	0.022	−0.021	−0.002	−0.037	0.106
	(0.020) [39]	(0.016) [59]	(0.030) [58]	(0.048) [25]	(0.110) [16]
Plastic kilograms	−0.054	0.273	0.596***	0.929***	1.119***
	(0.185) [40]	(0.229) [60]	(0.165) [60]	(0.163) [27]	(0.250) [17]
Picker sellers	−0.113	−0.087	0.295	0.381**	1.242***
	(0.097) [40]	(0.126) [60]	(0.211) [60]	(0.168) [27]	(0.399) [17]
Household sellers	−0.197*	0.227***	0.354**	0.558***	0.523***
	(0.111) [40]	(0.049) [60]	(0.160) [60]	(0.150) [27]	(0.137) [17]
Panel B. Second opening within 5 kilometers (40 incumbents)					
Posted price of a plastic grade on the menu	−0.043**	0.037*	0.043	−0.012	−0.036
	(0.020) [28]	(0.021) [38]	(0.039) [38]	(0.040) [20]	(0.057) [13]
Quality-adjusted plastic index	−0.001	0.178***	0.274**	0.172**	0.180
	(0.080) [28]	(0.056) [38]	(0.101) [38]	(0.076) [20]	(0.106) [13]
New-grade component	0.032**	−0.054***	−0.073***	−0.004	−0.008
	(0.014) [28]	(0.015) [38]	(0.024) [38]	(0.043) [20]	(0.069) [13]
Plastic kilograms	−0.173	−0.032	0.049	0.408	0.732***
	(0.153) [29]	(0.198) [40]	(0.237) [40]	(0.265) [22]	(0.187) [15]
Picker sellers	−0.001	−0.514**	−0.065	0.242	0.567*
	(0.128) [29]	(0.233) [40]	(0.238) [40]	(0.191) [22]	(0.306) [15]
Household sellers	−0.113	−0.090	−0.170	0.153	0.308**
	(0.090) [29]	(0.079) [40]	(0.186) [40]	(0.152) [22]	(0.121) [15]

Notes: For an incumbent already treated by a first opening, the event is the next roster-dated opening within the stated radius; the incumbent is treated at it if it qualifies under the rule of Table 9. Comparison banks are once-treated incumbents qualifying at that date whose own second rival arrives later or never, with their months kept while they have one rival. Estimators, fixed effects, reference period and clustering (by the second entrant) as in Tables 9 and 11. Brackets give the number of incumbents identifying the cell. The median gap between the first and the second opening is 2 months within 15 kilometers and 5 within 5, so the windows overlap the first event's; the quantity rows show that incumbents receiving a second rival are in areas where the network is expanding. The price cut is therefore a comparison across areas with different rates of entry rather than a second increment of the first effect. The index follows the six-month reappearance rule of Section 3.2. Standard errors, in parentheses, are clustered by the entrant. Stars: 10, 5 and 1 percent.

Table A15: Where the quantity response lands: full decomposition, kilogram shares, and the within-seller version

	Picker depositors		Household depositors	
	Unweighted	Banks equal	Unweighted	Banks equal
Panel A. Additive decomposition of the kilogram response by the line's own history (bank and month effects)				
Outcome: kilograms on lines of the kind, divided by the bank's mean monthly kilograms; the rows of each block sum to its first row				
All plastic kilograms	1.78 (1.07)	2.55 (1.90)	1.90 (1.52)	−1.78 (2.41)
continuing lines at an unchanged posted price (28; 70)	1.43 (0.87)	1.39 (1.22)	−0.51 (2.19)	−4.31 (3.06)
continuing lines repriced up within six months (40; 14)	1.03 (0.78)	0.71 (0.60)	0.97 (0.88)	0.93 (0.76)
continuing lines repriced down within six months (31; 14)	−1.09 (1.09)	−0.60 (0.75)	1.65 (1.08)	0.99 (0.71)
lines transacting for the first time (1; 1)	−0.09 (0.16)	−0.33 (0.40)	0.01	−0.07 (0.14)
lines returning after a silence of more than six months (1; 1)	0.50 (0.61)	1.37 (1.60)	−0.23 (0.23)	0.67 (0.75)
Bank–months (banks)	520 (16)		578 (21)	
Panel B. By the line's quality-adjusted price relative to the bank's continuing schedule this month				
lines priced above the continuing schedule (42; 37)	−0.18 (0.55)	0.49 (0.77)	2.78 (1.43)	2.07 (0.98)
lines priced at or below it (58; 63)	1.96* (0.96)	2.06 (1.65)	−0.88 (1.49)	−3.85 (2.15)
Panel C. Share of the type's kilograms on lines of the kind (bank and month effects, banks equal; bank–months in which the type delivers)				
repriced up within six months		1.99 (1.88)		1.14 (0.76)
repriced down within six months		−1.38 (1.71)		0.85 (0.58)
transacting for the first time		0.04 (0.06)		0.15 (0.31)
returning after a silence		−0.20 (0.64)		2.22 (1.75)
priced above the continuing schedule		−1.36 (1.73)		1.16 (1.21)
Bank–months (banks)		370 (14)		570 (21)
Panel D. Within bank–month: asinh kilograms of the type on the line, line and bank-by-month effects (carried menu)				
Instrument interacted with the line's status, relative to continuing lines at an unchanged price				
× repriced up within six months	0.264 (0.155)		−0.158 (0.159)	
× repriced down within six months	0.095 (0.099)		−0.111 (0.121)	
× transacting for the first time	0.827** (0.294)		0.102 (0.161)	
× returning after a silence	0.228 (0.159)		−0.157 (0.308)	
× priced above the continuing schedule (separate regression, bank-clustered)	−0.047 (0.051)		−0.144* (0.077)	
The line's own posted-price change over six months (OLS)	0.019 (0.043)		−0.208*** (0.063)	
Instrument × the line's fixed gap to the resin's coarsest grade	0.101 (0.159)		−0.104 (0.120)	
Line–months (banks)	11,705 (16)		11,800 (21)	
Panel E. Within continuing seller: seller and month effects, log tenure control (active months)				
asinh kilograms	1.30 (1.29)		0.75 (0.73)	
Share of her kilograms on lines repriced up within six months	1.95** (0.75)		0.05 (1.08)	
share of her baseline basket that is currently repriced up (benchmark)	1.42 (0.78)		−0.60 (1.41)	
difference: reallocation toward repriced lines	0.42** (0.15)		0.62 (0.39)	
Share on lines she had not supplied before	0.29 (0.73)		0.75 (0.67)	
Share on lines priced above the continuing schedule	−1.06 (0.93)		0.79 (1.54)	
Log distinct resins delivered in the month	−1.75** (0.65)		−1.45** (0.61)	
Log distinct lines delivered in the month	−1.15 (1.22)		−0.84 (0.74)	
Seller–months (sellers; banks)	3,745 (471; 11)		21,736 (2,232; 15)	

Notes: Plastic deliveries by individual depositors at the 21 banks and 578 bank–months of the supply design, with seller types and account exclusions as in Table 5. Coefficients are reduced forms on the instrument of Table 5; dividing by the first stage of 3.62 gives the response to a unit of the index. Lines are classified each month as in Table 8, a line is above the continuing schedule when its quality-adjusted price exceeds the credit-weighted level of the bank's continuing lines that month, and re-created lines are merged with their predecessor (66 of 598). Row labels in Panels A and B give the percent of picker and household kilograms on lines of the kind. Panels A and B: kilograms on lines of the kind over the bank's window-mean monthly plastic kilograms, so that cells sum to the coefficient on total kilograms, weighting bank–months equally (first column of each pair) or banks equally (second); picker regressions use the 16 banks at which pickers deliver in the window. Panel C: the type's kilogram share on lines of the kind. Panel D: the carried menu holds every line transacted in the prior six months, with zeros when it does not transact; line and bank-by-month effects absorb the instrument, so its interactions measure the tilt of kilograms across lines within a menu, and the coefficient on first-transacting lines describes growth from a base near zero rather than a share of the response. Panel E: continuing sellers with at least three plastic months; the benchmark is the share of the lines a seller used in her first three months that are repriced up this month. Standard errors two-way clustered by bank and month. Stars in Panels A–C from the null-imposed Webb wild-cluster p-value (2,999 draws); Panels D and E use clustered p-values on 11 to 21 clusters and are anti-conservative. 10, 5 and 1 percent.

Table A16: The menu under the upstream value shock: which posted prices move

	Mean of outcome	Coefficient on instrument	Wild p	Bank–months (banks)
Panel A. Bank–month: log posted prices and spreads on the carried menu (bank and month effects)				
Posted price, non-PET resins	7.00	2.95** (0.89)	0.011	528 (18)
Posted price, PET	7.54	−0.38 (0.37)	0.353	528 (18)
PET minus other resins	0.55	−3.32*** (0.91)	0.009	528 (18)
Posted price, bottom fixed price tier	6.41	2.75* (1.53)	0.079	524 (18)
Posted price, top fixed price tier	7.90	0.78 (0.81)	0.402	524 (18)
Top minus bottom tier	1.49	−1.97* (1.02)	0.087	524 (18)
Standard deviation of posted prices	0.69	−0.94** (0.50)	0.039	552 (18)
90th minus 10th percentile of posted prices	1.57	−1.46* (0.76)	0.089	552 (18)
Within resin, finer grades minus the resin's coarsest grade	0.52	1.38 (0.80)	0.267	534 (17)
Credit-weighted mean posted price	7.47	2.80** (0.75)	0.027	552 (18)
Quality-adjusted index (first stage, this weighting)	0.28	2.50*** (0.85)	0.005	578 (21)
Panel B. Bank–month, within resin: the coarsest grade, finer continuing grades, and finer entering grades (transacting lines)				
PET				
posted price, coarsest PET grade	6.99	−2.10 (4.01)	0.759	419 (14)
posted price, finer continuing PET grades	7.58	−0.50 (1.47)	0.804	411 (13)
posted price, finer PET grades entering this month	7.59	−0.75 (4.35)	0.803	88 (10)
gap: finer continuing minus coarsest	0.60	2.02 (3.88)	0.726	411 (13)
gap: finer entering minus coarsest	0.81	10.08 (6.72)	0.117	88 (10)
finer entering grades: gap to the continuing schedule	−0.08	5.69 (3.22)	0.220	85 (9)
any finer PET grade entering (0/1)	0.24	−1.03 (1.48)	0.584	419 (14)
Other resins				
posted price, coarsest grade	6.49	3.06 (1.90)	0.121	427 (14)
posted price, finer continuing grades	7.13	1.43 (2.98)	0.643	413 (11)
posted price, finer grades entering this month	7.22	−0.53 (3.52)	0.846	85 (11)
gap: finer continuing minus coarsest	0.65	−2.39 (2.54)	0.383	413 (11)
gap: finer entering minus coarsest	0.97	−3.80 (4.39)	0.273	85 (11)
Panel C. Line–month: pass-through to a continuing line by who supplies it (line effects and bank-by-month effects)				
		Instrument × household share	Instrument × PET	Line–months
Household share of the line's kilograms, 2022–26		−0.73* (0.37)		6,783
with the PET interaction		−1.17** (0.54)	−0.80 (0.46)	6,783
Household share measured in 2022, outcomes 2023–26		−1.60** (0.63)		3,554
with the PET interaction		−1.83* (0.91)	−0.48 (0.44)	3,554
Instrument × PET alone (line and month effects)			−0.67 (0.39)	6,869
Panel D. The two components of the first stage under alternative treatments of persistence (bank and month effects)				
	Continuing	Entering		
Levels, headline (Table 5)	1.61 (0.88) [0.11]	2.01** (0.79) [0.04]		578 (21)
Monthly change in the component on the instrument	0.02 (0.12) [0.90]	0.13* (0.06) [0.05]		486 (17)
Monthly change in the component on the change in the instrument	1.18 (0.66) [0.23]	0.28 (0.21) [0.18]		486 (17)
Band-passed shifter, 12-month window	1.94 (1.44) [0.33]	1.59*** (0.80) [0.01]		578 (21)
Band-passed shifter, 24-month window	1.73 (1.10) [0.20]	1.50** (0.61) [0.01]		578 (21)
With an exposure-specific linear trend (clustered p)	1.07 (1.00) [0.30]	0.59 (0.46) [0.22]		578 (21)
Range across leave-one-bank-out samples	0.84 to 2.35	1.52 to 2.59		

Notes: Panels A and B are bank–month regressions of a menu statistic on the instrument of Table 5 with bank and calendar-month fixed effects, each bank's observations receiving total weight one, standard errors two-way clustered by bank and month in parentheses, and p-values from the null-imposed Webb wild-cluster bootstrap clustered by bank (999 draws). Panel A uses the carried menu, every line transacted in the prior six months at its last posted price; the fixed price tiers are terciles of the line's window-median posted price among the bank's lines. Panel B uses lines transacting in the month, in bank–months with at least two grades of the resin; the coarsest grade is the grade with the lowest attribute value $x_j'\gamma$ of equation (2), and an entering grade transacts for the first time or returns after a silence of more than six months (Section 3.2). Panel C regresses the log posted price of continuing lines on the instrument interacted with the household share of the line's kilograms (demeaned within bank; standard deviation 0.30) and with a PET indicator, with line and bank-by-month fixed effects, so that the level of the instrument is absorbed; the share is measured over the estimation window, or over 2022 alone with the outcome restricted to 2023–26 so that it is predetermined (11 banks). Panel D re-estimates the first stage of the continuing- and entering-grade components of the index (Table 5) on the month-to-month change in each component, on the band-passed instrument of Appendix Table A8, with an exposure-specific trend, and leaving one bank out at a time; wild p in brackets. Stars from the wild p (Panels A, B, D) or clustered p (Panel C): 10, 5 and 1 percent. With 14 to 21 bank clusters the point estimates in Panel B should be read with their intervals.

Table A17: What a continuing seller delivers when the upstream value rises

	Picker sellers		Household sellers	
	Reduced form	2SLS on index	Reduced form	2SLS on index
Panel A. Seller–month: the price she faces and the quality of what she brings (seller-by-bank and month effects)				
Kg-weighted schedule level of the lines she delivers to (log)	3.83** (1.41)	1.28*** (0.37)	2.91** (1.04)	0.73** (0.24)
Kg-weighted attribute value of her load (log)	−1.21 (1.07)	−0.41 (0.48)	−0.31 (1.01)	−0.08 (0.25)
Log realized price per kilogram (credit/kg)	1.48 (1.67)	0.50 (0.41)	2.25*** (0.67)	0.56*** (0.17)
Share of kg on the top fixed price tercile	0.43 (0.52)	0.14 (0.13)	0.25 (0.46)	0.06 (0.12)
Share of kg on lines above the continuing schedule	−1.06 (0.92)	−0.35 (0.26)	0.79 (1.53)	0.20 (0.39)
Share of kg on grades finer than the resin's coarsest	0.46 (2.14)	0.16 (0.68)	1.70 (1.73)	0.41 (0.35)
asinh kg on sorted PET grades	5.90** (2.50)	1.98*** (0.30)	0.07 (0.56)	0.02 (0.14)
asinh kg on unsorted PET grades	2.87 (4.69)	0.96 (1.19)	4.00** (1.83)	1.00** (0.41)
asinh kg on non-PET grades	−7.01 (5.82)	−2.35* (1.15)	−3.06 (2.66)	−0.76 (0.55)
Log distinct lines in the month	−1.08 (1.22)	−0.36 (0.55)	−0.80 (0.77)	−0.20 (0.19)
Log distinct resins in the month	−1.79** (0.62)	−0.60 (0.35)	−1.42** (0.62)	−0.35*** (0.09)
Lines per visit	−2.69 (2.49)	−0.90 (1.11)	−2.84 (1.90)	−0.71* (0.39)
Log kilograms	1.27 (1.38)	0.43 (0.56)	1.08 (0.84)	0.27 (0.26)
Log plastic credit	2.75* (1.29)	0.92** (0.37)	3.33** (1.15)	0.83* (0.39)
Seller–months (sellers, banks)	3,745 (471, 11)		21,736 (2,232, 15)	

Panel B. Seller–month, pickers with a prior deposit: on which lines her kilograms rise (seller-by-bank and month effects)	Coefficient on instrument	× first month at the bank
asinh kilograms on lines repriced up	13.59*** (3.96)	−0.21 (0.24)
Share of kilograms on lines repriced up	1.86** (0.70)	−0.06 (0.04)
asinh kilograms on entering lines	0.35 (1.23)	0.06 (0.07)
Share of kilograms on entering lines	−0.02 (0.16)	0.02 (0.02)
asinh kilograms, all plastic	1.26 (1.27)	−0.01 (0.08)
Picker–months (pickers, banks)	4,163 (680, 13)	

Panel C. Bank–month: the fine share of the bank's plastic kilograms and its decomposition (bank and month effects)	Coefficient on	Wild p
Fine share among continuing sellers	0.55 (0.90)	0.554
Within-seller component (fixed seller weights)	0.69 (1.02)	0.511
Composition component across continuing sellers	−0.22 (0.23)	0.346
Fine share among sellers new to the bank this month	0.42 (1.28)	0.745
Share of the bank's fine kilograms from new sellers	−0.51 (0.41)	0.237
Share of the bank's kilograms from new sellers	−0.45 (0.35)	0.219
Log number of new sellers	−0.97 (2.15)	0.656
Bank–months (banks)	507–530 (14–17)	

Notes: Panel A: continuing sellers at the 21 banks of the supply design, with at least three plastic months at the bank in the window, as in Table A18, Panel B, with the account exclusions of Table 5. Each row is a regression of the outcome on the instrument (reduced form) or on the quality-adjusted index instrumented by it (2SLS), with seller-by-bank and calendar-month fixed effects and log tenure at the bank, standard errors two-way clustered by bank and month. The schedule level is the kg-weighted mean over the lines she delivers to of the bank's quality-adjusted posted level on the line, and the attribute value the kg-weighted mean of $x_j'\gamma$ from equation (2), so that the two sum to her realized log price. The fixed price tercile and the continuing schedule are as in Table A16; the fine share is defined on resins with at least two transacting grades; the PET categories are those of Table 4. Panel B: every picker with a prior deposit at the bank, the outcome regressed on the instrument, its interaction with an indicator for the seller's first month at the bank, and that indicator; lines are classified as in Appendix Table A15. Panel C: bank–month regressions with bank and month fixed effects; the fine share is the share of the bank's plastic kilograms on grades finer than the coarsest grade of their resin, and the within and composition components hold seller weights fixed at the bank's window means. Stars in Panels A and B from the clustered p-value with 11 to 15 bank clusters, which is anti-conservative; Panel C from the wild p clustered by bank. 10, 5 and 1 percent.

Table A18: How sellers respond to the upstream value shock: pay, kilograms, and where the kilograms land

	Mean of outcome	Coefficient on instrument	Implied response to the index	Bank–months (units)
Panel A. Levels (seller–months at the 21 banks)				
Median monthly plastic credit, picker sellers (Rp)	28,080			4,454
household sellers (Rp)	4,170			25,591
Median plastic kilograms per month, picker sellers	22.9			
household sellers	3.0			
Panel B. Within continuing seller: take-home pay and kilograms (seller and month effects)				
Picker sellers, log monthly plastic credit	10.11	6.78 (2.84)	1.87	3,433 (435)
log kilograms	3.00	4.28 (2.70)	1.18	
Household sellers, log monthly plastic credit	8.67	3.56 (1.90)	0.98	22,048 (2,268)
log kilograms	1.35	2.17 (1.37)	0.60	
Panel C. Bank–month: the bank's plastic payout by seller type (bank and month effects)				
Log payout to picker sellers	11.45	17.12*** (4.19)	4.73	370 (14)
log payout per active picker	10.36	11.42* (3.71)	3.15	
Log payout to household sellers	10.97	−3.51 (3.88)	−0.97	570 (21)
log payout per active household	8.99	−4.53 (3.37)	−1.25	
Picker minus household, log payout per active seller	1.24	10.63** (2.40)	2.94	361 (13)
Picker share of the bank's plastic payout	0.27	0.74 (0.59)	0.20	578 (21)
Panel D. Bank–month: picker kilograms and their composition by category of grade (bank and month effects)				
	asinh kg	share of kg		
All plastic (log kilograms)	11.09** (2.79)			370 (14)
By whether the grade was on the menu before the window				
on the menu before 2022 (67)	4.44 (4.76)	−2.23 (1.56)		
entered since January 2022 (33)	23.81 (10.97)	2.23 (1.56)		
in tons rather than asinh: on the menu before 2022	1.14 (2.01)			
entered since January 2022	−0.55 (3.54)			
By resin				
PET (43)	14.23* (6.13)	2.26 (1.85)		
other resins (57)	5.01 (5.74)	−2.26 (1.85)		
Within resin (bank–resin–month, bank–resin and month effects; resins with two or more grades)				
the resin's coarsest grade (39)	6.38* (1.96)	−0.44 (0.66)		
grades finer than the resin's coarsest grade (61)	8.34 (3.59)	0.44 (0.66)		
Bank–resin–months (banks)		1,158 (13)		

Notes: Plastic deliveries by individual depositors at the 21 banks and 578 bank–months of the supply design, with seller types and account exclusions as in Table 5. Each coefficient is from a regression of the outcome on baseline plastic exposure interacted with the real-rupiah PET import unit value, with standard errors two-way clustered by bank and month; dividing by the first stage of 3.62 gives the response to a unit of the index. All regressions weight banks equally, each bank's observations receiving total weight one, because the hub of Section 4.3 records in every month of the window and would otherwise take a large share of the identifying variation. Panel A reports medians over seller–months; means are far larger for pickers (Rp 816,000) because of a few very large sellers at the hub. Panel B uses continuing sellers, those with at least three plastic months at the bank in the window, with seller-by-bank and calendar-month fixed effects. Panels C and D are bank–month regressions with bank and calendar-month effects; payout is plastic credit paid to individual depositors of the type. Panel D is estimated on picker deliveries only, since household kilograms do not respond in total and the household cells by category are too imprecise to decompose; the two columns are the response of asinh picker kilograms in the category and of the category's share of picker kilograms, and the shares within a pair of categories sum to one, so their coefficients are equal and opposite. Row labels give the bank-weighted mean share of picker kilograms in the category. In Panel D a grade is a bank-specific attribute vector, a grade has entered since January 2022 from the month in which it first transacts or returns after a silence of more than six months, and the coarsest grade of a bank–resin is as in Table 12. The asinh outcomes include bank–months in which pickers deliver nothing in the category (36 percent of picker bank–months for grades entering since 2022 against 4 percent for older grades, with bank-weighted means of 198 and 819 kilograms), so the entering grades grow faster from a small base and the two rows in tons are added for that split. With the band-passed shifter of Appendix Table A8 the entering-grade share responds by 2.1 (1.7). Stars from the null-imposed Webb wild-cluster p-value clustered by bank (2,999 draws in Panels C and D, 999 in Panel B): 10, 5 and 1 percent.

Table A19: The grades that enter a bank's menu, and who delivers on them

Panel A. Entering grade–months at the 21 banks, January 2022–July 2026 (533 grade–months, 346 grades)

	Share of grade–months	Share of entering credit	Median log gap to coarsest grade	Share of credit above schedule	Median rank within resin
Revived: finer than the resin's coarsest carried grade	0.44	0.57	0.92	0.78	0.59
Revived: the resin's coarsest grade or coarser	0.10	0.03	0.00	0.61	0.07
Revived: resin not carried in the prior six months	0.11	0.01		0.21	
First transaction: refinement of a resin the bank carries	0.30	0.39	1.05	0.63	0.58
First transaction: coarser than or equal to the carried grade	0.02	0.00	0.00	0.17	0.00
First transaction: resin not carried	0.02	0.00		0.29	

Panel B. Who delivers on a grade after it enters: shares of picker kilograms on the grade, by months since entry

	Sellers arriving at or after entry	Incumbents who supplied the grade	Incumbents who supplied the resin	Incumbents new to the resin
First-transaction grades (101 grades, 59 tons of picker kilograms in the first year)				
entry month	0.44		0.55	0.02
months 1–3	0.52		0.47	0.01
months 4–6	0.59		0.38	0.02
months 7–12	0.49		0.48	0.03
months 1–6, mean across banks (6 banks)	0.27		0.71	0.02
Revived grades (162 grades, 66 tons)				
entry month	0.33	0.27	0.27	0.13
months 1–3	0.58	0.17	0.20	0.05
months 4–6	0.53	0.18	0.13	0.16
months 7–12	0.68	0.06	0.21	0.06
months 1–6, mean across banks (9 banks)	0.38	0.25	0.23	0.14

Panel C. The resin's picker kilograms before and after a finer grade enters (shares; pre = months −3 to −1, post = months 1 to 6)

	Coarsest grade	Existing finer grades	The entrant	Grades entering later
First-transaction refinement (76 events), pre	0.33	0.67		
post	0.19	0.60	0.04	0.17
Revived finer grade (214 events), pre	0.19	0.81		
post	0.13	0.66	0.07	0.15

Notes: Panel A classifies every grade–month scored as entering under the six-month rule of Section 3.2 at the 21 banks of the supply design. A revived grade transacts after a silence of more than six months (median silence 13 months, interquartile range 9 to 22). A first-transaction grade is a refinement if the bank transacted the resin in the prior six months and the grade's attribute value $x_j'\gamma$ exceeds that of the coarsest grade it carried. The gap to the coarsest grade is in log posted price; the schedule is the credit-weighted quality-adjusted level of the bank's continuing lines; the rank is the share of the resin's carried grades that are coarser than the entrant. The composition of entering credit across these categories does not move with the instrument (all $p > 0.19$). Panel B follows each entering grade for twelve months and splits the picker kilograms delivered on it by the seller's history at the bank: first deposit at or after the entry month, or before it, and for incumbents whether she had delivered that grade, or that resin, before. Pooled shares weight large banks heavily; the bank-mean row averages each bank's own shares. Panel C reports, for each event, the shares of the resin's picker kilograms at the bank on the coarsest carried grade, the other grades carried before the event, the entrant, and grades entering after it, averaged over events. All panels are descriptive.

Table A20: Persistence of grades written during the window

Months after the grade is written	1	3	6	12	18	24
Panel A. Does the grade stay on the menu?						
Share of written grades transacting in that month	0.50	0.43	0.37	0.31	0.27	0.27
transacting in that month or the two before it	1.00	0.64	0.53	0.44	0.39	0.34
first transactions	0.59	0.49	0.43	0.34	0.26	0.33
re-listings	0.43	0.38	0.32	0.28	0.29	0.21
Written grades with the bank recording in that month	512	483	446	407	369	316
Panel B. Among grades still transacting: kilograms and price relative to the month written						
Kilograms, median ratio to the month written	1.26	1.19	1.27	1.68	1.54	1.97
Kilograms, total ratio to the month written	1.61	2.96	1.76	2.03	2.15	4.29
Change in posted log price since written	0.00	0.01	0.00	−0.05	−0.04	−0.05
Premium over the continuing schedule when written	−0.08	−0.13	−0.09	−0.06	−0.06	−0.10
Premium over the continuing schedule in that month	−0.08	−0.13	−0.09	−0.10	−0.07	−0.15
grades written in high-instrument months, when written	−0.03	−0.07	−0.02	0.01	0.00	−0.08
in that month	−0.01	−0.03	−0.03	−0.07	−0.05	−0.19
grades written in low-instrument months, when written	−0.13	−0.17	−0.14	−0.13	−0.11	−0.15
in that month	−0.14	−0.19	−0.14	−0.14	−0.08	−0.05
Grades still transacting	257	208	165	124	101	86

Notes: The 571 plastic grades that enter the menu of one of the 21 supply-design banks between January 2022 and July 2026, either transacting for the first time (232) or returning after a silence of more than six months (339), followed month by month. Panel A conditions on the bank recording any plastic transaction in the month; a grade that does not transact in a given month has not been removed from the schedule (Section 3.2), so the second row counts a grade as present if it transacted in that month or either of the two before it. Panel B is over grades transacting in the month. The premium over the continuing schedule is the gap between the grade's quality-adjusted price and the credit-weighted quality-adjusted level of the bank's continuing grades in the same month, clipped at ±1 log point as in the index; it is the quantity the new-grade component of the first stage accumulates (Table 5, Panel B), and it is unweighted here (credit-weighted it is 0.12 at entry, 0.16 for grades written in high-instrument months and 0.11 in low-instrument months). A high-instrument month is one in which the bank's instrument is above its own mean over the window. Of the kilograms that grades written during the window eventually carry, 5 percent arrive in the month written, 36 percent within the first year and 59 percent after it.

Table A21: The six-month reappearance rule: silences between transactions of a grade, prices on return, and the first stage by rule

Months since the grade last transacted	1	2	3	4–6	7–9	10–12	13–24	>24
Panel A. How common are silences?								
Share of consecutive transactions of a grade	0.83	0.08	0.03	0.03	0.01	0.01	0.01	0.01
Panel B. The price when the grade transacts again								
Share returning at exactly the previous price	0.84	0.76	0.69	0.70	0.58	0.65	0.57	0.68
Mean absolute change in the bank's continuing-grade level during the silence	0.06	0.09	0.08	0.11	0.13	0.14	0.17	0.23
Share repriced in the direction the continuing level moved, where it moved by more than 10 log points	0.16	0.23	0.26	0.25	0.44	0.29	0.16	0.12
Pass-through of the continuing-level change into the return price (slope)	0.15	0.30	0.16	0.17	0.80	—	—	0.34
Returns	5,637	505	198	190	81	43	62	50
Panel C. Does a silent grade come back?								
Silent so far (months)	1	2	3	6	9	12		
Share of such silences that end within the next six months	0.73	0.61	0.50	0.39	0.26	0.20		

Notes: Plastic grades at the 21 supply-design banks, January 2022 to July 2026, where a grade is a bank-specific attribute vector (resin, cleanliness, color, form, code and suffix), so that lines a bank re-creates for a grade it already carries are one grade with it. Panel A is the distribution of the gap, in calendar months, between consecutive transactions of the same grade; measured in the bank's recording months the shares are 0.87, 0.06, 0.02, 0.02, 0.01, 0.01, 0.01 and 0.00. Panel B compares the posted price when a grade transacts again with its own price at the previous transaction, and relates the change to the movement of the bank's credit-weighted quality-adjusted level over continuing grades during the silence; the slope is from a regression of the price change on that movement within the gap class and is omitted where fewer than ten returns move the level. Panel C follows silences from their start among grades at banks with at least twelve recording months after the silence begins. Returns after a silence of up to six months are overwhelmingly at the old price and barely reflect movements in the bank's schedule, whereas returns after a longer silence are repriced to the schedule far more often, which is what the six-month rule of Section 3.2 encodes. Extending the rule to twelve and twenty-four months, or never treating a return as new, gives first stages of 3.13, 2.93 and 2.11 in a re-chained reconstruction of the index whose six-month value is 3.36.