Easy to Collect, Costly to Bear: Mobile Money Taxation and Network Spillovers in Tanzania*

Youssef Benzarti Santiago Garriga Dario Tortarolo

November 9, 2026

Abstract

This paper provides evidence that a tax on mobile money transactions shrinks its base by a third, that the base responds to whether it is taxed rather than to how much, and that the damage spreads through the payment network to users and payments the law never taxed. First, using the complete transaction records of Tanzania's three largest providers around the introduction of a notched levy on person-to-person transfers and cash withdrawals in July 2021, we show that the levy reduces the volume of taxed transactions by 23 percent and their average value by 13 percent relative to untaxed transactions on the same platform, an elasticity of about 10 with respect to the share of a transfer the sender keeps. Second, we show that a cut of about 30 percent seven weeks later brings none of it back, nor do two further cuts, while the removal of the levy on transfers in 2023 does, slowly. Third, we document that relationships between senders and receivers survive at two-thirds the normal rate, that new ones form at less than half the pre-levy pace, and that senders stop rerouting activity around recipients who go quiet; none of these margins recovers when the tax rate falls. Fourth, we find that the levy reaches activity it does not tax, as payments from businesses into wallets, which include wages, fall by about a tenth at the two providers that carry nearly all of them, and that senders of every size cut back by a similar share, whether or not they have a bank account. Fifth, we estimate that the levy destroys 47 cents of surplus for every shilling it raises, several times the cost of a value-added tax in comparable settings.

JEL Classification: H20, H22, H23.
Keywords: levy, taxation, mobile money, Tanzania.

*Benzarti: UCSB and NBER, benzarti@ucsb.edu. Garriga: CEFIP, IIE-FCE-UNLP, santiago.garriga@econo.unlp.edu.ar. Tortarolo: World Bank DECRG, dtortarolo@worldbank.org. We thank Alfred Shauri, Twivwe Siwale, Saugata Mittra, and Revocatus William, as well as staff from the Bank of Tanzania. We are grateful to Monica Singhal for suggesting the salary-channel analysis. The views expressed are those of the authors and do not necessarily represent those of the World Bank or its affiliated organizations.

1 Introduction

The classical constraint on taxation in low-income countries is observability. With large informal sectors, weak third-party reporting, and limited administrative capacity, governments cannot reach most of the bases the standard public finance toolkit would recommend, and tax-to-GDP ratios have moved little for decades (Bachas et al., 2026). Mobile money breaks this constraint completely. Forty percent of adults in sub-Saharan Africa had a mobile money account in 2024, transactions equal 5 percent or more of GDP in several East African economies, and every transaction is captured electronically, in real time, by a handful of licensed providers under direct central-bank supervision (Klapper et al., 2025; GSMA, 2025). A tax on this base is withheld at source, so compliance is mechanical; collection requires no auditors or field enforcement; and evasion at the platform level is essentially impossible. It is what that toolkit would prescribe for a low-capacity state, and at least fifteen sub-Saharan governments have reached for it (Barczay et al., 2025). Removing the administrative constraint, however, may come with large behavioral responses: how much of the base disappears when it is taxed, and what brings it back?

We answer this question with transaction-level administrative data covering 1.3 billion mobile money transactions in Tanzania between April and November 2021, provided by the Bank of Tanzania together with three of the country's six operators: Vodacom (M-Pesa), Tigo (Tigo Pesa), and Airtel (Airtel Money), which together hold approximately 85 percent of registered accounts (TCRA, 2021). Identification comes from two policy shocks seven weeks apart. On July 15, 2021, the government introduced a notched 23-bracket levy on person-to-person transfers and cash withdrawals; for a representative 40,000 TZS transfer, the total cost of sending money rose from 1.03 to 4.78 percent of the amount, placing Tanzania at the top of the global rate distribution (Innovations for Poverty Action, 2023). On September 1, 2021, after public backlash, the levy was cut by approximately 30 percent in every bracket above the smallest, taking the cost of the same transfer to 3.65 percent; M-Pesa and Tigo left their own fees unchanged in every bracket, so the cut reached users in full. Transfers and withdrawals are taxed; deposits and merchant, business, and government payments are not. The untaxed types therefore serve as a control, subject to every shock of 2021 but reached by the levy only indirectly (Section 4).[1] We estimate the response to each shock by difference-in-differences on the daily volume and average value of taxed against untaxed transactions.

In the weeks after the introduction the daily volume of taxed transactions falls by approximately 23 percent relative to the untaxed control, the average value of a taxed transaction falls by approximately 13 percent, and the two margins compound to a 33 percent decline in the total value of taxed activity. The implied elasticity is approximately 10 with respect to the share of a transfer the sender keeps. In concrete terms, a levy averaging 2.1 percent of the value of a transfer removes a third of the taxed base. The response is to the levy, not to its schedule: sorted by the size of their usual transfer, senders whose typical transfer carries a levy of 0.7 percent of the amount cut back by nearly as much as senders whose typical transfer

[1] M-Pesa books transfers sent to other networks as merchant payments, and 80 percent of those rows carry the levy. Appendix A.2 moves them to the taxed side; after the fix, 0.3 percent of control rows are levied. Transfers arriving from other networks, which the receiving provider books as business-to-person payments and the sending provider levies, stay in the control; row (7) of Appendix Table A4 drops business-to-person payments.

carries 3.3 percent. Against merchant payments alone, the comparison most favorable to a large response, it is 19, and it reaches 21 if the control's own decline is attributed to the levy.

The September cut does not undo the July response. A cut that removed 30 percent of the levy should, if users respond to the price, have restored about 30 percent of what the introduction took away; it restored none of the lost transactions or value. Using aggregate data, we show that two further cuts in 2022 left transfer value at 68 percent of its 2020 level. In July 2023 transfers were removed from the levy and cash withdrawals were not. Transfer value then recovered, closing half of the gap to its pre-levy trend within two years; withdrawal volumes stayed at their 2020 level through 2025. A levy averaging 2.1 percent of a transfer's value removed a third of the base. Cutting it to 1.5 percent brought none of it back, nor did the larger cuts that followed. Only the end of the levy on transfers did.

The missing recovery is a puzzle. In the standard model of a tax on transactions, only the rate in force today matters: a user who found a transfer worth making at a lower price should find it worth making again when the price moves back toward that level, whatever happened in between. A base that a price rise removed should return, in part, for a price fall; this one did not, at any of the three cuts. Several explanations could reconcile the fall with the missing recovery, and the data rule out most of them: providers capturing part of the cut, habits or lapsed relationships that those who left would have to overcome, and a network that tipped into disuse, with each user waiting for the others to return. Two explanations survive, and our data cannot separate them: a base that responds to the presence of a visible tax rather than to its rate, as users who perceive a schedule coarsely would, or a country in which everyone had moved back to cash and expected everyone else to stay there. Either way, a rate cut gives up revenue without buying back the base.

On a payment network, every transaction leads to another, so a tax on some of them reaches money the law never taxed. Deposits, which are untaxed, fall by a third, and most where the network was densest: accounts that received transfers from several counterparties cut their deposits by a third or more, isolated accounts by about a quarter. Wages paid into mobile wallets are not taxed, but everything a worker can do with the wage afterward, sending it on or cashing it out, is. A wage paid into a wallet is therefore worth less than the same wage paid in cash, and payments from businesses into wallets fall by about a tenth relative to the untaxed control at M-Pesa and Tigo, which carry 95 percent of them. Part of what mobile money had done over a decade, turning wages paid in cash into documented payments, is undone. Merchant payments, the one untaxed way to spend a wallet balance, are the exception: they recover after the cut and end above their pre-levy level. What leaves the platform does not move into the banking system either: users with a bank account cut back as much as those without one (Section 5.5).

The network dissolution is directly visible in the data. The levy cuts the ties between users: of the pairs who sent money to each other between June 1 and July 14, 2021, only about one in six do so again between July 15 and August 31, a third fewer than they normally would. New ties form at less than half the pre-levy pace. The relationships that were forming when the levy arrived were the most fragile: senders whose circles included more newly active counterparties fell more and did not return.

To learn whether the levy's damage travels from one user to the next, we investigate what happens when the people a sender transacts with go quiet. Every sender has a circle of regular recipients, and in any period

some of them go dormant, transacting with no one for weeks. Before the levy, a sender losing part of that circle sends more to everyone else. The money still has somewhere to go: a relative still needs support, a supplier still needs paying, so when one relationship lapses, activity moves to another. After July 15 the opposite happens: a sender losing part of the circle sends less to everyone else, at all three providers, and the September cut leaves this unchanged.[2] Instead of routing around quiet counterparties, senders follow them into inactivity, and disuse spreads from user to user along the web of who pays whom. The levy's damage therefore does not stop with the user who pays it, because on a payment network each user's reason to transact depends on whether her counterparties still do. The cost of the levy then includes losses to users who never changed their own behavior, which the standard cost calculation omits, and the network has lost the resilience that used to absorb a lost relationship: what was a shock absorber has become a conduit, and a lower rate did not restore it.

The objection is that no one may be infecting anyone. People transact with people like themselves, a pattern network economists call homophily, so a sender and a quiet counterparty may have withdrawn for the same reason: everyone faced the same tax, and people who resemble each other may have responded to it alike. We run a battery of tests against this reading. The first is the one that separates it from a spillover, because the two predict opposite things about which counterparties matter. If homophily is doing the work, the counterparties who most resemble the sender carry the strongest signal of her own type, so the pattern should be strongest among them; if contagion is doing the work, what matters is that the counterparty stopped sending, so the pattern should be no weaker among counterparties unlike her. Measuring resemblance by whether the two live in the same district, at Tigo, where the registry records location, and by how far apart their typical transfer sizes are, the pattern fades only mildly with distance: counterparties in another district carry over nine tenths of the contagion of those in the sender's own district, and counterparties whose transfers differ from hers by more than a factor of four still carry most of it; if homophily were doing the work, the pattern should fade with resemblance, and it barely does. We run additional tests to bound what homophily could contribute. Second, there are no pre-trends: the pattern is flat in the six weeks before July 15 and switches on with the levy, so a shared trait would have to be triggered by the levy rather than already at work.[3] Third, the traits we can see do not carry it: holding fixed how much each sender's counterparties transact and how large their transfers are, the two traits on which users most resemble their counterparties and along which the levy bit hardest, removes between a fifth and two fifths of the pattern, and a trait we do not observe would have to select users more strongly than these two do to remove the rest.[4] What survives is a trait shared by two users who decide together whether to keep using the platform, a household or a business with two accounts, and a circle of five or more counterparties leaves little room for it. We therefore

[2] For each ten percentage points of a sender's circle gone quiet, activity with third parties is higher by about 10 percent at Tigo, 8 at M-Pesa, and 7 at Airtel before the levy; after July 15 the association falls by 15, 11, and 9 percentage points, and the change after September 1 is under one point at each provider (Table 4).

[3] Figure 13 estimates the association fortnight by fortnight; it is flat over the three pre-levy fortnights at Tigo and Airtel, rises by two points in the last of them at M-Pesa, and reverses within two fortnights of July 15.

[4] The controls are deciles of the counterparties' pre-levy transaction volume and transfer size, each interacted with the fortnight. By the coefficient-stability ratio of Altonji et al. (2005) and the bound of Oster (2019), selection on an unobserved trait would have to be one to four times the selection on these two to remove the reversal (Appendix A.18).

read the pattern as causal: one user's withdrawal spread to the next.

Network effects account for a large part of the levy's damage. By October, about a fifth of the fall in taxed activity has run through contacts going quiet, counting the contacts of those contacts. They carry the levy into the activity it does not tax: deposits fall by a third, payments from businesses into wallets by about a tenth. Because the activity the levy does not tax falls too, measuring the levy against it understates the damage; counting the whole fall, the levy destroys 47 cents of surplus for every shilling it raises (Section 7). The losses of users who kept transacting while their counterparties left, which the standard calculation omits, raise the levy's cost of public funds by more than a fifth at both rates (Section 7.2).

The size of the response makes the levy an expensive way to raise money. Every transfer the levy causes to exit the mobile money platform is a loss to its sender with no revenue against it. We estimate that this loss is 47 cents for every shilling the levy raises (Section 7). A typical VAT in a low-income country costs well under 10 cents per shilling. The property tax is the levy's mirror image, hard to collect but with a tax base that cannot flee in response to the tax. Over the three years in which the revenue authority collected both taxes, the mobile money levy raised seven times more revenue, 288 against 39 TZS billion a year (Figure 5). It raised more because it is easier to collect, not because it is cheaper. Cutting the rate made it only slightly cheaper, 41 cents instead of 47: the cost per shilling depends on how far the tax shrinks the base, not on the rate, and the base stayed down through the September cut and the two further cuts of 2022, so the revenue the cuts gave up went to users who were transacting anyway, and the damage stayed. Section 7 also sets out the designs the government passed over. The main one is a levy on withdrawals alone, the design Tanzania itself arrived at in 2023: it spares transfers and keeps two-thirds of the revenue at the same rates, but withdrawals respond to price nearly as strongly as transfers do, so matching the levy's revenue would have needed rates about 60 percent higher. A higher tax on the fees providers charge would add only about half a percentage point to the cost of a transfer, against nearly four for the levy, but several times more to a withdrawal, whose fees are higher. A flat monthly charge per account taxes nothing per transaction, but some light users would leave the platform to avoid it, taking their contacts' transactions with them, and it falls hardest on the accounts that pay least today.

Tanzania's experience with taxing mobile money reflects a continental pattern (Bachas et al., 2026). At least fifteen sub-Saharan governments have taxed mobile money transactions, and most entered as Tanzania did, with a broad levy on transfers, adopted under fiscal pressure because it could be legislated within weeks on a base with no organized constituency; statutory rates elsewhere have mostly run from a fraction of one percent to one and a half percent of the transaction, against a Tanzanian schedule whose effective rate reached 4.8 percent (Barczay et al., 2025; Hearson et al., 2024). Nearly every broad, high-rate levy on transfers has since been cut, narrowed, or abolished, and the designs that survive are the narrow ones: taxes on provider fees, flat or capped charges, and low uncapped rates. In spite of these cuts, five more countries have legislated levies since 2025, at low rates or on narrow bases, with little direct evidence on what the levy costs.[5] Our estimates say what each of these steps does: a broad levy on transfers removes a third of

[5]Uganda retreated to a withdrawal-only levy within five months of its 2018 adoption; Tanzania removed transfers from its levy in 2023; Ghana repealed outright in 2025 and Chad in 2026. The designs that persist are taxes on provider fees (Kenya, Togo, Madagascar), flat charges (Nigeria), and low uncapped rates (Cameroon). Since 2025 Senegal has adopted a levy of 0.5 percent on

the base, a partial cut recovers none of it, and only removal brings it back, slowly.

The paper speaks to three bodies of work. First, the literature on taxation in low-capacity states has framed the choice as one between bases that are easy to administer but distortionary and bases that are inert but hard to reach (Best et al., 2015; Brockmeyer et al., 2024). Mobile money is the easiest base to administer that a low-income state has, and we measure what that ease costs. Its closest precedent is the bank debit taxes of Latin America, whose base eroded as their rates rose (Coelho et al., 2001; Baca-Campodónico et al., 2006). The base falls by a third, and the cost is not undone when the tax is cut: the base that a rate increase removes does not return when the rate falls. Static comparisons of instruments miss this second cost, the quantity counterpart of the price asymmetry that Benzarti et al. (2020) document for value-added taxes. Here the cost is decided by the levy's presence rather than its rate: within the levied range, cutting the rate gives up revenue and undoes no damage, and only removal restores the base.

Second, we provide, to our knowledge, the first empirical account of what a tax does to the relationships a payment network is made of.[6] Alvarez et al. (2025) show that a positive shock to network size leaves a lasting mark only when users' benefits depend on each other, and Higgins (2024) shows a shock to one side of Mexico's payment market raises adoption among agents it never touched. Crouzet et al. (2023) show that a temporary shortage of cash left a lasting rise in the adoption of electronic wallets in India, and Björkegren (2019) that ignoring network effects understates the welfare cost of taxing mobile phones. Our findings say what the network externality costs. In an ordinary market a tax makes people buy less, and when the tax is cut they buy more again, because the sellers and the customers are still there. On a payment network, when one user stops, the people who used to send to them lose part of their reason to use the platform, and some of them stop too. This chain accounts for about a fifth of the decline by October. What it changes is who bears the cost: users who never changed their own behavior lose part of their reason to use the platform, a loss the standard calculation of a tax's cost leaves out. The network makes the damage wider, and it makes the platform fragile, since the repair that used to absorb a lost relationship stops and a lower rate does not restart it. The stock that matters is the set of relationships, not the count of users. The literature measures a network by how many people are on it, and by that measure Tanzania's network barely moved: the levy moves the number of active accounts by a few percent at most, 4 percent in our data and under 1 percent in the Bank of Tanzania's national count. The number of pairs actually sending money to each other falls by half. Nearly everyone keeps their account; what they lose are the people they sent to. A study that counted users would have found the network intact when half of its connections were gone. And a mature network is not protected: Hanson et al. (2025) prescribe taxing once adoption is high, and Tanzania did, at 45 percent of adults, yet the network dissolved and rebuilt only once the tax on transfers was gone, so the optimal tax on a network good has to count what the tax does to the relationships, not only to the number of users.

Third, research on mobile money has mostly studied the benefits of the technology, for financial inclu-

transfers, capped at 2,000 CFA francs, and Malawi one of 0.05 percent on transfers above 100,000 kwacha; Mali has adopted a levy of 1 percent on withdrawals, Niger one of 0.5 percent on transfers of 2 million CFA francs or more between operators, and Gabon's 2025 finance law set one of 0.5 percent on transactions of 10,000 CFA francs or more.

[6] Rivadeneyra et al. (2022) study a subsidy to payments on Ecuador's government-run mobile money network. Hanson et al. (2025) model a network as a stock whose growth taxation slows. Garriga and Tortarolo (2024) show that appointing large firms as collection agents can improve tax compliance of their clients and suppliers.

sion, risk-sharing, savings, and the response to shocks (Jack and Suri, 2014; Suri, 2017; Batista and Vicente, 2025; Apeti and Edoh, 2023); we add causal estimates of what happens when it is taxed, from the complete data of a national market, across a rate increase, its partial reversals, and its removal, and for the relationships the network is made of. The closest evidence is Barczay et al. (2025), who use data from one operator to estimate the response to the introduction of levies in Cameroon and the Central African Republic, with Mali as a control, and Paul and Sharma (2025), who study the same Tanzanian levy from the household side, with survey data and a triple difference, and find rural food consumption falling by 10 to 18 percent.

The rest of the paper is organized as follows. Section 2 describes the institutional setting and the 2021 reform. Section 3 introduces the data. Section 4 sets out the empirical strategy. Section 5 presents the results: the main elasticities, the response to the cut and why it is not a price response, the indirect incidence on wages paid through the platform, and the distributional incidence of the levy. Section 6 documents what the levy did to the payment network. Section 7 discusses welfare and revenue implications and computes the marginal value of public funds. Section 8 concludes.

2 Institutional setting and the tax reform

2.1 Mobile money in Tanzania

Tanzania has a population of roughly 60 million and a tax-to-GDP ratio of about 13 percent that has been broadly stable for two decades. As in much of sub-Saharan Africa, a large informal sector and limited reach of the conventional banking system make the standard tax bases administratively expensive to expand: the informal economy is estimated at just over half of GDP (World Bank, 2023), and 23 percent of adults hold an account at a bank or other formal financial institution (Demirgüç-Kunt et al., 2022). Mobile money grew quickly in this context. The technology was introduced in Tanzania in 2008 when Vodacom launched M-Pesa; the Bank of Tanzania reported 0.4 million mobile money transactions that year, 1 billion in 2013, and 3.2 billion in 2021, worth 115.5 trillion TZS. Active wallets rose from 9 million in 2013 to an average of about 32 million during our sample period, across an adult population of roughly 34 million. By the time of the levy in July 2021 mobile money was the country's dominant payment infrastructure. Because users frequently hold accounts with more than one provider, the share of adults with at least one mobile money account is lower: the Findex 2021 survey puts it at 45 percent, well above the sub-Saharan African average of 27 percent (Klapper et al., 2025).[7]

As of December 2021, six operators served the country, with shares of registered accounts of 40 percent for Vodacom (M-Pesa), 24 percent for Tigo (Tigo Pesa), 21 percent for Airtel (Airtel Money), 10 percent for Halotel (HaloPesa), 3 percent for TTCL (T-Pesa), and 2 percent for Zantel (EzyPesa) (TCRA, 2021). We observe transaction-level data for the three largest (Vodacom, Tigo, and Airtel), which together account for roughly 85 percent of registered mobile money accounts. All providers are licensed under the National Payment Systems Act of 2015 and supervised by the Bank of Tanzania, which is the source of the data

[7] The 2021 Findex survey in Tanzania was fielded face to face between August 2 and 29, 2021, inside the July 15 to August 31 period of the levy's introduction rates, with 1,001 respondents.

we use. The GSMA, the industry association of mobile network operators, estimates that mobile money transactions contribute more than 5 percent of Tanzanian GDP, on par with the rest of East Africa (GSMA, 2025). The technology serves both as a person-to-person transfer mechanism and a payment rail for retail commerce, utility payments, salaries, and government transfers.

For identification, three features of the platform matter. First, every transaction is captured electronically, in real time, at the level of the individual sender and receiver, with a direct readout of revenue collected at the source. Second, the small number of regulated providers makes the levy administratively trivial to implement: it is withheld by the operator and remitted to the revenue authority, with no role for the user in compliance. Third, for most Tanzanian households mobile money is the closest functional substitute to cash. An estimated three-quarters of Tanzanian adults have no bank account (Klapper et al., 2025). A transaction that leaves the platform therefore moves in most cases to cash, and the behavioral margin we estimate is not a shift between competing formal payment instruments, but between a regulated digital platform and the informal cash economy.

2.2 The 2021 levy

The mobile money levy was introduced in the Finance Act of 2021, under pandemic-era fiscal pressure. Prior to the reform, mobile money transactions were already subject to a 10 percent excise duty on provider charges (introduced in 2014) and the standard 18 percent VAT on those charges (since 2016). The new levy was layered on top of these as a third instrument, applied directly to the transaction value rather than the provider's fee. The implementing regulations were issued as Government Notice 496A on June 30, 2021 (the National Payment Systems Electronic Mobile Money Transfer and Withdrawal Transactions Regulations). They specified a graduated schedule of fixed-amount levies on a defined subset of mobile money transactions; the regulations came into operation on July 1, 2021, and collection began on July 15, the date at which the levy appears in the transaction data. Seven weeks later, on September 1, 2021, Government Notice 642A reduced the levy amount in every bracket above the smallest by approximately 30 percent, leaving the scope of the tax and the bracket boundaries unchanged. These are the two policy shocks we exploit.

The announcement gave the range of the levy but not what a typical transfer would pay. The budget speech of June 10, 2021 proposed "a Levy of between 10 shillings to 10,000 shillings in each mobile money transaction of sending and withdrawing," varying with the value of the transaction, and gave no schedule (Ministry of Finance and Planning, 2021). The schedule came with the regulations of June 30, and users first paid it on July 15, when the cost of sending 40,000 TZS rose from 1.03 to 4.78 percent of the amount (Section 2.3). The outcry came after collection began. By July 19 the levy had been "roundly criticised as people believe the amount being deducted was just too high for Tanzanians to carry," and the Minister for Finance and Planning said the President had instructed the ministry to work on the complaints (Wambura, 2021). Taxed activity in our data follows the same timing: it does not fall after the speech, and it falls on July 15 (Section 5).

The levy applies to two transaction types only: person-to-person transfers and cash withdrawals (cash-

out). All other transaction types are explicitly exempt: cash deposits (cash-in), person-to-business and business-to-person payments, person-to-government and government-to-person payments, and business-to-business payments. This split of the platform's activity into taxed and untaxed types, shown in Figure 1, is the basis for our identification strategy (Section 4). Exemption does not insulate a type from the levy, however: every type draws on the same wallets, so a user who tops up less or cashes out in response to the levy also transacts less in the exempt types. Section 4 states the no-spillover assumption this violates, and Section 7 bounds the consequence.

The schedule is a notched 23-bracket fixed-fee system, reproduced in Appendix Table A12. The smallest taxed transactions (1,000 TZS) carry a 10 TZS levy under both regimes; the largest (above 3 million TZS) carry 10,000 TZS at introduction and 7,000 TZS after the September cut. Figure 2 plots the two schedules. Within each bracket the levy is constant in transaction value, which produces an effective rate that declines smoothly with transaction size inside the bracket and jumps discontinuously upward at each bracket boundary (panel B). The effective rate is about 1 percent at the bottom of the schedule, peaks at 4.8 percent just above 20,000 TZS, and converges toward zero at large values; we document this rate structure in the next subsection.

At the time of the transaction, the provider withholds the levy together with its own service fee, the 10 percent excise duty on that fee, and the 18 percent VAT on the fee plus excise, so the user sees a single deduction at the point of sending. The provider remits the levy directly to the Tanzania Revenue Authority, and our data contain the exact levy amount paid on each transaction. The levy in our data matches the statute. In every one of the 23 brackets, between 85 and 99 percent of taxed transactions between July 15 and August 31, 2021 pay exactly the statutory amount; the remainder are transactions the levy did not reach, such as merchants' withdrawals (Section 5); they remain in the taxed series, which if anything dilutes the estimated response.

The schedule was lowered twice more after our sample ends, and the levy was then removed from transfers. In July 2022 a replacement schedule widened the base to transfers from bank accounts to mobile wallets and to withdrawals at ATMs, and cut every bracket above the smallest by roughly 40 percent relative to the September 2021 schedule. In October 2022 the schedule was lowered again. In July 2023 the Finance Act removed transfers from the levy altogether, leaving cash withdrawals as the only taxed transaction. The withdrawal levy is still in force as of the 2026/27 budget, which pairs it with mandates for digital payment in several sectors. Our sample window, April to November 2021, therefore covers the levy at its highest rate, and one cut. The design we study is still in use. Tanzania and Uganda continue to levy cash withdrawals, and Mali began to in 2025. Since 2025 Senegal and Malawi have adopted levies on transfer values, the base Tanzania taxed in July 2021 and exempted again in July 2023, at 0.5 and 0.05 percent. Our estimates of the introduction and of the September cut therefore describe what these governments can expect from a broad levy on transfers, and what a later partial cut does and does not restore.

2.3 The size of the price shock

The levy nearly quintupled the cost of sending money. For a representative 40,000 TZS transfer, the all-in cost, combining the provider's service fee, the 10 percent excise duty, the 18 percent VAT, and the new levy, rose from 410 TZS to 1,910 TZS, from 1.03 to 4.78 percent of the amount sent. The September levy cut brought it down to 1,460 TZS, or 3.65 percent. Appendix Figure A7 plots the series.

Write $c(A)$ for the all-in cost of a transfer of A shillings. We measure the price of a transfer by the share of the amount that the sender keeps, $p^k(A) = 1 - c(A)/A$, which is the net-of-tax price in the sense of the taxable-income literature; the levy is set by the amount transferred, not by the provider's fee, so this is the price it moves. At 40,000 TZS the share kept fell from 0.990 to 0.952 at the introduction, a log change of $\Delta \ln p^k = -0.0386$, and the September cut raised it by 0.0118. Section 4 states how it scales our estimates.

The effective rate of the levy varies sharply across transaction sizes. Because the levy is constant within each bracket, the effective rate is highest at the lower bound of each bracket and declines smoothly toward its upper bound, with discrete jumps at boundaries. Appendix Table A13 reports the levy as a share of transaction value at the lower bound of representative brackets under the two regimes. The effective rate peaks at around 4–5 percent for transactions in the 20,000–50,000 TZS range (a range that holds about one in six person-to-person transfers by count; half are below 10,000 TZS), and is about 1 percent from 500,000 TZS, falling below it from 800,000 TZS.

3 Data

We use transaction-level administrative data covering the universe of mobile money activity at Tanzania's three largest providers between April and November 2021. The data are provided by the Bank of Tanzania, the regulator of the country's payment systems, through direct extracts from the licensed mobile money operators. The extract covers Vodacom (M-Pesa), Tigo (Tigo Pesa), and Airtel (Airtel Money), which together account for approximately 85 percent of registered mobile money accounts; the three smaller operators, Halotel, TTCL, and Zantel, are not in the extract. The dataset contains every transaction processed by these providers over the period; it is not a sample.

The data cover every transaction in six populated categories: person-to-person transfers, person-to-business and business-to-person payments, payments to government, and cash-in (deposits) and cash-out (withdrawals) at agents. Two further legally defined categories (business-to-business and government-to-person transfers) contain no observations in the available extract and are omitted from the analysis. The treatment and control groups used throughout the paper follow the legal scope of the levy described in Section 2.2, person-to-person transfers and cash-out taxed and merchant payments, business-to-person payments, payments to government, and cash-in untaxed, with two corrections for provider booking conventions (Appendix A.2).

Each transaction carries anonymized sender and receiver identifiers, a timestamp to the second, the value in TZS, the provider's fee, the levy paid, and the transaction type. The anonymized identifiers allow us to track individual users across transactions and over time, producing a user-level panel within the sample

window. Additional attributes available at the user level include user category (person, business, government entity, or agent), date of registration, age, gender, district of residence, and sector of economic activity; we use a subset (gender, district) in the heterogeneity analysis. The data are accessed under a confidentiality agreement with the Bank of Tanzania and the participating operators; individual identifiers are masked before extraction, so it is not possible to recover the underlying identity of any user.

The full sample contains approximately 1.30 billion transactions, averaging around 162 million transactions per month across the three providers.[8] Table 1 reports summary statistics by transaction type. Person-to-person transfers, the directly-taxed category, account for 13.1 percent of total transaction count; cash-out, the other taxed category, accounts for 19.4 percent; together, the two taxed types make up about a third of all activity. The largest single category by count is merchant payments (30.4 percent), followed by cash-in (25.1 percent); both are untaxed.

4 Empirical strategy

The levy applied to a strict subset of transaction types and not to the rest, which gives us a control group. We estimate the response to each of the two policy shocks (the July 15, 2021 introduction and the September 1, 2021 cut) using a difference-in-differences design that compares the evolution of taxed and untaxed transaction series before and after each date. The design delivers the response of taxed activity relative to untaxed activity on the same platform; because the untaxed types share the platform and fall with it, that differential is a lower bound on the absolute response. Section 7 bounds the gap.

Our baseline specification, estimated at the transaction–type-by-day level, is

$$\log Y_{cpt} = \alpha_{cp} + \delta_{pt} + \beta_1 D_c \cdot \mathbf{1}\{t \geq t_1\} + \beta_2 D_c \cdot \mathbf{1}\{t \geq t_2\} + \varepsilon_{cpt}, \tag{1}$$

where c indexes the transaction types described in Section 3 and p the provider, so that each provider's transaction types form separate daily series; t indexes calendar days; D_c is an indicator for the taxed types, person-to-person transfers and cash-out; t_1 and t_2 denote July 15 and September 1, 2021; α_{cp} are series fixed effects and δ_{pt} provider-by-day fixed effects, so that each provider's taxed series are compared with the same provider's untaxed series on the same day; and Y_{cpt} is the daily number of transactions of type c at provider p (how many transactions happen), the daily mean transaction value (how large they are), or the daily total transaction value (the product of the two). The coefficient β_1 measures the response of taxed activity to the levy introduction relative to the pre-policy baseline; the coefficient β_2 measures the additional response to the September rate cut.

[8] Tigo's April 2021 delivery duplicates every transaction (two daily files four-fold) and its November delivery duplicates the transactions of November 28–30; the duplicates are exact copies, timestamp included, and we remove them at the raw-file level before any processing (Appendix A.1). The counts here and in Table 1 are net of this correction, which removes 68.8 million duplicate rows; no difference-in-differences window contains an affected day.

Estimation windows. We estimate β_1 on the narrow window June 1 – August 31, 2021 and β_2 separately on the August 1 – October 13 window with the September 1 indicator.[9]

Scaling. The coefficients of equation (1) are log changes in activity. To express them as elasticities we divide by the log change in the share of the amount kept, $\varepsilon = \beta_1 / \Delta \ln p^k$, with β_1 estimated on daily total value and $\Delta \ln p^k = -0.0386$ for the 40,000 TZS reference transfer (Section 2.3). This is the convention in which the comparator instruments of Section 7 are calibrated. It is not a structural parameter: the levy is a step schedule, so the price change differs by transaction size, and we take it at a single reference transfer.

Parallel trends. The identifying assumption is that, in the absence of the levy, the log daily volume and the log mean value of taxed and untaxed transactions would have evolved in parallel. The weeks before July 15 provide the pre-policy period over which to test this directly. We report event-study versions of equation (1) that replace each post-period indicator with a series of leads and lags around its respective policy date, both to evaluate pre-trends and to trace the dynamic adjustment after each shock. The pre-period coefficients are small relative to the treatment response, and the movement among them is the month-end cycle (Appendix A.4). We return to them in Section 5 (Figure 6).

Spillovers. Equation (1) also assumes no spillover from the treated to the control series, the stable-unit-treatment-value assumption (SUTVA). It is unlikely to hold on a shared platform. The levy on person-to-person transfers reduces the value of holding a mobile-money balance, so users top up less, cash out, or shift wage receipts off the platform; these adjustments depress every series that requires a funded wallet, including the untaxed controls. Merchant payments, cash-in, and business-to-person payments in Figure 3 all fall alongside person-to-person transfers, though by smaller amounts. The coefficient β_1 therefore identifies the differential response of person-to-person transfers relative to a contaminated control, not the absolute response of person-to-person transfers to the levy. With both series shifted down by the same shock, the differential understates the absolute fall in taxed activity, and the main estimate in Table 2 is a lower bound on the absolute response of person-to-person transfers. The contamination may affect the magnitude of the estimate, but not the paper's conclusion: if anything, the true contraction in taxed activity is larger. We return to the welfare implications in Section 7.

Inference. Standard errors are clustered by series: seventeen series in Panel A of Table 2, eight of them taxed, and six in Panel B, three of them taxed. With so few clusters, cluster-robust standard errors can over-reject (Conley and Taber, 2011; MacKinnon and Webb, 2017), so we also report p-values from a wild cluster bootstrap. Because we observe essentially the entire population of transactions, the uncertainty that matters is not sampling noise but the autocorrelation of the daily series (Abadie et al., 2020). Appendix A.3 addresses it with placebo breaks at every pre-levy date; the mean-value outcome, which fails that test in the main specification, is the least secure of the three.

[9] The rate-cut window ends on October 13 to exclude a duplication artifact in the Airtel extract from approximately October 16 onward (Appendix A.1; notes to Table 2).

5 Results

5.1 The response to the levy introduction

Figure 3 plots the weekly volume of mobile money transactions by type from June 1 through October 13, 2021, normalized to 100 on the pre-policy reference week of July 8–14.[10] The corresponding estimates of equation (1) are in Table 2. Panel A, our main estimate, uses every untaxed series as the control; Panel B compares the taxed transfers against merchant payments alone (P2B in the tables), the untaxed series whose pre-levy path tracks the taxed series most closely.

Taxed activity falls sharply and immediately at the levy introduction: relative to the untaxed control, volume falls by approximately 23 percent and the mean transaction value by 13 percent, compounding to a 33 percent decline in total value. The untaxed series fall as well, though by less. The one visible movement before the levy, a rise of about 15 percent in business-to-person payments in the last week of June, is the end-of-month wage cycle, amplified by the close of the fiscal year on June 30; it is gone by the reference week, affects treated and control series alike, and drops out of the difference-in-differences by construction (Appendix A.4). The implied elasticity is approximately 10 with respect to the share of a transfer the sender keeps.

Taxed activity falls sharply at the introduction. Both taxed series drop at the July 15 policy date and stay down. Person-to-person transfers fall to a trough of 62 within three weeks and remain there through August; cash-out, taxed on the same schedule, falls less far, to a trough of 70.[11] Against merchant payments alone (Panel B of Table 2), the declines are 29 percent on volume and 33 percent on mean value; because the merchant-payment series' own decline is the smallest among the untaxed types, we read the Panel B figures as the upper range rather than the main estimate.

The two margins compound. The daily total value of taxed activity falls by approximately 33 percent relative to control, and by 52 percent against merchant payments. Figure 4 shows the same fall in daily total transacted value series by series: person-to-person transfers drop to approximately 47 percent of the reference week within three weeks and reach a trough of 42 percent in late August, and cash-out reaches a trough of 61 percent. Appendix Figure A3 plots daily mean transaction value, which falls on every type, taxed and untaxed, except merchant payments, whose mean value edges up.

Untaxed activity falls too. The statutorily untaxed series decline alongside the taxed ones. Cash-in drops to 85 and merchant payments to 91, consistent with the platform-level spillover discussed in Section 4. Business-to-person payments fall further, to 72, for the separate reason developed in Section 5.4: it carries

[10] The plotted series follow the same assignment as the estimation, except that business-to-person payments exclude transfers arriving from other networks and Vodacom's payouts from its own accounts (Section 5.4).

[11] The smaller decline is not a smaller price shock. Withdrawals and transfers carry identical levy amounts at every transaction size, and applying the July schedule to each type's size distribution over the sample gives an average effective rate of 2.4 percent per withdrawal against 2.1 percent per transfer. Only per shilling is the withdrawal rate lower, 1.5 against 1.9 percent, because withdrawal value is concentrated in large transactions on which the fixed levy is a small share.

the indirect incidence of the levy through the salary channel. No untaxed series in the figure rises above its pre-policy level between July 15 and August 31, ruling out within-platform substitution as the channel at the introduction. The figure-level numbers report absolute changes in each series, while the difference-in-differences coefficients report differentials relative to a control; the gap between the two is the spillover into the untaxed series, which we use formally in the welfare calculations of Section 7.

The response is immediate, and the movement before it is the month-end cycle. Figure 6 reports the event-study version of the main specification on the daily count series. The coefficients drop to -0.17 at week 0 and settle between -0.25 and -0.40 through the end of August. Before the levy they sit between -0.10 and $+0.01$, below zero in June and back at zero in the two weeks before collection. That movement is the month-end cycle of Appendix A.4: once the taxed series are given their own day-of-month profile, the pre-levy coefficients are jointly indistinguishable from zero (wild cluster bootstrap $p = 0.16$ on count and 0.86 on total value), and the main estimate barely moves (-0.38 on total value against -0.40).

The implied elasticity is about 10. For the representative 40,000 TZS transfer, the user cost rose by roughly 3.75 percentage points at the introduction, so the share of the amount kept fell by approximately 3.9 percent in log terms (Section 2.3). The main response of approximately 40 log points on total value then implies an elasticity of taxed activity with respect to the share kept of approximately 10 (10.3; the Panel B response of 73 log points implies 18.9). The reference transfer sits at the bottom of its bracket, where the levy is the largest share of the amount, so it gives the smallest share-kept elasticity the schedule allows: at the average levy, 2.1 percent of a transfer's value, the same response implies about 19.[12]

How large an elasticity is depends on the base more than on the taxpayers (Slemrod and Kopczuk, 2002). A transaction tax is escaped by moving the same activity to cash, a close substitute for the taxed instrument in most uses, or by forgoing the activity. Our data see the taxed base and not what its users did instead, so the estimate bundles avoidance and a real reduction in activity.[13] For the cost of the levy the split does not matter at the margin: a user who avoids the levy gives up at most the levy avoided, whether she pays in cash or does not transact, and the decline in the taxed base is the statistic Section 7 uses. The levy taxed the

[12] Studies of transaction taxes and of fees measure the price by the fee instead. Dividing the percentage fall in value by the percentage rise in the fee-inclusive price, which the levy raised by 82 percent on average over the taxed transactions of M-Pesa and Tigo at their recorded fees (Airtel's data record no fees), our elasticity is 0.40 (0.63 against merchant payments alone). On that convention, from a fee increase at Tigo in 2013, Economides and Jeziorski (2017) estimate elasticities with respect to the fee of -0.18 for the amount transferred and -0.40 for the amount withdrawn, the same order as ours. Closest to our setting in quasi-experimental design, Barczay et al. (2025) divide a 40 percent fall in taxed transaction value per user by a 19.45 percent rise in the fee-inclusive price of a taxable transaction and obtain an elasticity of approximately 2 for Cameroon's 0.2 percent levy, in a propensity-score-matched difference-in-differences design against Mali. Their 2.1 is about five times ours: the Tanzanian levy raised the price of a transfer by far more than the Cameroonian one did, and the base fell by less in proportion. Closest in country and policy, Fjeldstad et al. (2026) provide experimental evidence on the same Tanzanian levy: a lab-in-the-field experiment with 618 Dar es Salaam small-business owners finds that raising the salience of the levy reduces the share choosing mobile money over cash by 9–11 percentage points, operating through attention rather than information (96 percent of subjects already knew about the tax), the salience mechanism of Chetty et al. (2009) and Finkelstein (2009).

[13] One form of avoidance would have stayed on the platform: paying a business, which the levy did not tax, and having the business pay the recipient, which it did not tax either. For most transfers this route costs less than a taxed transfer even though it pays two provider fees, but we find no sign that it was used (Appendix A.10).

choice of payment instrument, one of the most substitutable margins in the economy, and the elasticity we estimate is a property of that choice of base.

5.2 The asymmetric response to the September cut

The response to the September levy cut is asymmetric. The levy is cut by nearly a third and taxed activity barely moves.

Taxed activity does not recover when the rate falls. Figure 7 summarizes both shocks for every series of Figure 3: the log change in daily activity after July 15, estimated on June 1 to August 31, and the additional log change after September 1, estimated on August 1 to October 13, the windows of Table 2, so that the gap between a taxed and an untaxed marker is close to the difference-in-differences of that table. At the introduction, every series falls: person-to-person transfers by 48 log points on count and 81 on total value, cash-out by 33 and 40, and the untaxed series by between 13 and 28 on count. At the cut, which lowered the levy at every bracket above the smallest by approximately 30 percent, every series rises by 5 to 14 log points on count, taxed and untaxed alike, and on total value the two series that rise most, merchant payments and business-to-person payments, are untaxed. The estimates of β_2 in Table 2 say the same thing: on the main specification the count of taxed transactions is unchanged relative to its high-rate level (-0.006; standard error 0.019), the value margins are flat, and against merchant payments the two value margins continue to fall, by roughly 4 to 6 percent. They are also smaller than an ordinary demand curve allows. The cut undid 30 percent of the levy's increase in the price of a transfer, so with a constant elasticity it should have restored 30 percent of the introduction's effect; on the main specification it restores none of the lost transactions or value, and we reject proportional recovery on both margins.[14] Six weeks after a cut of nearly a third in the levy, taxed activity has moved back toward its pre-policy level by no more than the series the cut never touched.

The untaxed comparison series does recover. Cash-in drops by approximately 15 points at the introduction and climbs back to roughly 93 by mid-October (Appendix Figure A10). Cash-in is untaxed, so its fall at the introduction cannot be a direct price effect; it is the platform-level spillover of Section 4, visible directly

[14] We stack the two windows of Table 2, so that the August days used by both estimates enter both coefficients, and test $\beta_2 = -0.305\,\beta_1$, where 0.305 is the ratio of the log net-of-tax price changes at the cut ($+0.0118$) and at the introduction (-0.0386); because the cut removed 30 percent of the levy in every bracket above the smallest, the ratio is close to 0.3 at every transaction size. On the main specification, with standard errors clustered by series (seventeen series), the share of the introduction's effect restored is -0.02 on count (95 percent confidence interval -0.19 to 0.12; $p = 0.001$) and -0.06 on total value (-0.16 to 0.12; $p = 0.007$), and a wild cluster bootstrap gives $p = 0.0002$ and 0.001. With the benchmark set instead by the fee-inclusive price of a transfer, which the cut moved back by 0.175 of its introduction change at the reference transfer and by 0.189 averaged over the schedule, the test still rejects on both margins (wild cluster bootstrap p between 0.006 and 0.010). The test does not reject on mean value ($p = 0.09$), whose response to the introduction is itself imprecisely estimated. Against merchant payments the shares restored are -0.05, -0.10 and -0.08 on count, mean and total value, and the test rejects on count and total value at the 5 percent level against the share-kept benchmark. Measured against its own pre-levy level, without netting out the untaxed series, the taxed aggregate restores 12 percent of lost transactions and 9 percent of lost value. The cut was passed through in full: within every levy bracket the modal provider fee is identical before and after September 1 at M-Pesa and Tigo (at Tigo net of the levy, which its data record inside the fee), and the levy falls exactly as gazetted; Airtel's data carry no fee field.

in the data: users who transfer less also top up less. The contrast across the September date is the asymmetry at its starkest. Cash-in, whose price never changed, recovers about half of its own loss; cash-out, whose levy was cut by nearly a third, recovers about a fifth of its own loss. After September every series rises by about the same amount, whether the levy on it was cut or not (Figure 7). That common rise is the platform slowly recovering from the shock of July, and it goes further for cash-in, which had lost little, than for cash-out, which had lost much.

Following senders, the contraction is close to proportional in transaction size. Figure 8 sorts a fixed cohort of senders once, on their pre-policy mean transfer, into five size groups, and tracks the weekly transaction count of each group relative to the reference week of July 8 to 14. The groups move together. Averaged over the thirteen weeks from July 18, the four upper quintiles sit between 49.6 and 51.0 percent of their reference week, and only the smallest quintile separates, at 57.4, upward. Part of that separation is composition: the schedule starts at 1,000 TZS, transfers below it were never levied and are concentrated in the smallest quintile, and at the introduction they fell by 0.19 log points, in line with the untaxed series and a third of the 0.58 of the brackets from 2,000 TZS up (Appendix A.19), so the group that sends them cuts back less. Nothing recovers toward the pre-policy baseline after the cut.[15] The evidence is consistent with a broad contraction that is close to proportional in transaction size, rather than with the selective exit of high-value users.[16] The rate structure makes the uniformity telling. At the median pre-levy transfer of each group, the levy at introduction takes 0.7 percent of the amount sent in the smallest quintile, 1.7, 2.2, and 3.3 percent in the next three, and 2.5 percent in the largest, while the cutback runs only from 43 to 50 percent of pre-levy activity: a group facing more than four times the rate of another cuts back by barely more. How much a sender cuts back does not track the rate the sender faces; the contraction is a response to the levy, not to its schedule.

Within the platform, the contraction is not substitution. The 3.8 million senders with at least three person-to-person transfers before the levy, followed across every type of transaction they initiate, cut all of them at the introduction: per sender and per day, transfers fall to 53 percent of their pre-levy rate, cash-out to 64, and the untaxed merchant payments to 83, and none recovers after the cut (Appendix Table A17). Compared with the same senders' merchant payments, the transfers of the 4.4 million senders who both sent money and paid merchants before the levy fall by 34 log points, as much as in the aggregate comparison, so the fall is not a change in who is on the platform; 36 percent of these senders stop sending transfers, against 25 percent who stop paying merchants. Substitution into cash would show up as fewer but larger withdrawals, and substitution into merchant payments as more merchant payments; among the fifth of senders who cut transfers the most, cash-out falls by 46 percent and merchant payments by 26 percent, and the mean withdrawal shrinks slightly. Nor do senders leave: 10 percent of the cohort initiate no transaction of any

[15] Sorting senders on a pre-policy window invites mean reversion. Appendix A.7 runs the same design on a window with no rate change: the mean reversion it finds runs the other way, and the level it produces with no policy change calibrates the post-levy level.

[16] Classifying transactions rather than senders produces an apparent size gradient, both at the introduction and in the missing recovery, with large transfers falling most and recovering least. Appendices A.5 and A.6 show that the gradient is mechanical and that users are not splitting transfers to chase lower effective rates.

type between July 15 and August 31 and 16 percent between September 1 and October 13, so more than eight in ten stay on the platform doing less of everything. The untaxed types fall for the very same senders, which corroborates the spillover used in Section 7. What the data cannot say is where the activity goes: cash, informal channels, another provider, and a genuine reduction in transacting are indistinguishable once a transaction no longer appears in the data. The Bank of Tanzania's national totals narrow this for one destination: the part of national transaction value that our data do not cover, which includes the three smaller providers, falls by about a quarter from June to August 2021 and by a third by September, much as the value at our three providers does, so the activity they lose does not appear to move to the others (Figure 9).

The levy does not push money out; it stops money coming in. If users were withdrawing balances to hold cash, cash-out would rise while cash-in fell. Both fall together, and by the same proportion: daily cash-out value falls from 81 to 55 TZS billion at the introduction and cash-in from 72 to 49, a third each, and neither recovers after the cut (Appendix Table A16). Income that would have been deposited stays in cash, and the money that stays on the platform moves less. The Bank of Tanzania's national aggregates show the same thing from the balance side (Figure 9). Each shilling held in mobile wallets moves about ten times a month in June 2021, seven in August, and eight in October, and in every month from August 2021 to June 2022 turnover stays below its level in every month from January 2019 to June 2021; the number of active accounts hardly moves (31.5 million in June, 31.3 million in September). In aggregate, users do not leave; they hold less and use it less.[17]

The one levy-free route through the platform recovers. The statute left one route untaxed on both legs. Merchant payments were exempt, and so in practice were merchants' own withdrawals. Paying a merchant electronically was therefore levy-free end to end, while a consumer's own routes to cash were taxed at rates up to 4.8 percent. Consistent with that wedge, merchant payments are the only series to recover after the September cut, running above their pre-levy baseline by October.

Lowering the levy does not bring activity back. Raising the levy and then lowering it are not opposite moves with opposite effects, and the asymmetry matters in two ways. The usual way of costing a tax assumes that activity returns when the rate is cut; here it does not, so the harm done at the high rate carries over into the period of the lower rate, and every shilling the government collects at the reduced rate costs users as much as a shilling collected before the cut (Sections 7 and 7.2). And a tax that is easy to impose turns out to be hard to take back: the government can lower the rate overnight, but it cannot restore the activity the tax has driven away. The next subsection sets the missing recovery against the conventional model and the Bank of Tanzania's later data; Appendix A.19 reports the tests behind the explanations the introduction rules out.

[17] Cash-in is not evidence of a spillover onto an untaxed party: deposits are made in order to transact, so its fall is the same behavior observed one step upstream.

5.3 The missing recovery is not a price response

The data reject the conventional model. We sketch the conventional model of a tax on transactions in Appendix A.16. A sender weighs each transfer on its own: the transfer is worth some amount over the best alternative, the schedule attaches a levy to it, and the sender makes it when its worth exceeds the levy, so daily volume is the share of potential transfers that clear the schedule in force that day. In this model only today's price matters, so a cut in the levy must raise volume: every transfer that was worth making at the higher levy is still worth making at the lower one, and some that were not become worth making. The September schedule is at or below the July schedule at every value and about 30 percent below it on every bracket above the first, so the model requires the reduced rates to bring back part of the lost volume, about 30 percent of it with a constant elasticity, whatever happened in between. It does not: the reduced schedule restores none of the lost transactions or value. Adjustment frictions extend the model but do not rescue it. Habit, inattention, or the cost of coming back can delay the return, but as long as what a transfer is worth to its sender is unchanged, the return must begin; the weeks from September 1 to October 13 show no movement toward recovery on any margin, and the fixed-cohort series in Figure 8 is flat to mid-October.

The base returns only when the levy on it is gone, and then slowly. In the Bank of Tanzania's annual series, transfer value in 2022, with two further cuts in force, stood at 68 percent of its 2020 level and at 47 percent of the level its pre-levy trend implied; after transfers were removed from the levy in July 2023, the gap to trend closed by 44 percent by 2024 and by 54 percent by 2025, while withdrawals, still taxed, stayed at their 2020 volume (Figure 10).[18] A base that comes back only once the levy is removed is what a response to the levy's presence produces and what a price response along a fixed schedule cannot.

5.4 Indirect incidence via the salary channel

The behavioral responses documented above are direct: the levy is statutory on person-to-person transfers and cash-out, and we estimate the response of those bases. A natural follow-up question is whether the levy depresses transaction types it does not statutorily tax. The mechanism we have in mind operates through the salary channel. Workers who receive their wages on the mobile money platform receive them through statutorily untaxed business-to-person payments (B2P in the tables), and pay the levy whenever they subsequently move that money via a person-to-person transfer or a cash-out. The levy therefore reduces the effective value of receiving wages on the platform, and may prompt workers to switch their employers' payment method to cash or bank transfer. About one in four Tanzanian wage-earners received wages via mobile money in 2021, so the salary channel is a real but bounded slice of the wage economy (Demirgüç-Kunt et al., 2022).

Payments from businesses into wallets fall by about a tenth, although the levy never touches them. At the introduction, the daily count of business-to-person payments at M-Pesa and Tigo falls by approximately 13 log points relative to the clean untaxed basket of cash-in and merchant payments, and the daily total value

[18] The trend is the pre-levy growth of total mobile money value, 19.5 percent a year; the 2020 to 2023 levels are printed identically in the Bank of Tanzania's 2023 and 2025 reports.

by approximately 11 log points; mean transaction value does not move (Panel A of Table 3). The estimate uses equation (1) with business-to-person payments as the treated series, compared within each provider. A transfer that arrives from another network is booked by the receiving provider as a business-to-person payment from its interoperability account, and the sending provider levies it, so we remove these transfers, together with Vodacom's payouts from its own accounts; they are about half of business-to-person payments by number and a third by value.[19] M-Pesa and Tigo carry 95 percent of the remaining payments. No pre-levy date produces a break in the count as large as the levy's in windows of 21 days on each side, and one in 78 does in windows of 14 days (Appendix A.3); the value effect is less precise. The implied elasticity is approximately 2.7, about a quarter of the main estimate of 10.3. On their own, without a control, these payments fall by about 28 log points on count, almost as much as cash-out, which is levied, and three fifths as much as transfers; in the same design as Table 2, they fall half as much as the taxed base on count and a quarter as much on value. After the September cut, they regain about half of the lost number, a gain the wild cluster bootstrap does not distinguish from zero ($p = 0.08$). The reading is the salary channel: as the levy makes wages received on the platform less valuable to workers, firms pay fewer of them there even though the payment itself is untaxed.

Workers with bank accounts do not move their wages to the bank. A wage paid straight into a bank account escaped the levy, so the fall could be workers with bank accounts moving their pay there. If it were, those workers would stop being paid on the platform sooner than workers without an account. They do not. Among M-Pesa wage recipients, those who had sent money from their wallet to a bank account between April 1 and May 14 (Appendix A.8), become more likely than the rest to keep being paid on the platform: the gap widens by 2.6 percentage points after the levy compared with the same comparison in the spring (standard error 0.35), and in September and October they are 6.1 points more likely to still be paid there. The fall is therefore not a move to bank payroll. What remains is wages paid in cash instead, or payments no longer made, and the data cannot tell the two apart (Appendix A.8).

The levy undid part of the formalization that mobile money had delivered. Over the past decade, a wage credited to a mobile wallet became data: digital, held by a regulated provider, and available for administrative use. The response of business-to-person payments measures how much of that data the levy erases. The wages that leave the platform do not move to bank payroll; they are paid in cash, which moves them into the informal economy, or the payments stop. Either way, the levy that exploits the digital base removes some of the flows around it from view, adding an informational cost on top of the fiscal cost.

5.5 Distributional incidence

The levy reduces activity by roughly the same proportion across all user groups we can distinguish. Senders of every size cut their transfers by a similar share, whatever rate they face (Figure 8), and senders with a

[19] At M-Pesa the arriving transfers are paid by the five accounts that the customer registry classes as mobile money interoperability; at Tigo, by the account whose payments match M-Pesa's outgoing transfers to Tigo day by day and amount by amount. Vodacom's own accounts pay amounts of a few hundred shillings and carry the duplication artifact of Appendix A.1.

bank account cut about as much as those without one: at M-Pesa both groups cut person-to-person transfers by 0.80 log points, and the banked by 4 points less once their larger volume and wealth are held fixed (Appendix A.9). A bank account therefore offers no better way out than the cash every user has. What leaves the platform therefore leaves the formal financial system rather than moving within it. Against this common pattern the differences are modest: the smallest senders cut back somewhat less than the rest, and urban senders somewhat more than rural ones.

Urban senders cut back somewhat more than rural ones. At Tigo, whose registry includes each sender's postal code, the response is -0.96 (standard error 0.03) for transactions whose sender's registered address is in Dar es Salaam and -0.84 (standard error 0.04) for the rest of the country, a difference of -0.12 (standard error 0.05; Appendix Table A14, Panel B). Bank accounts do not explain the gap. They are three times as common in cities (National Bureau of Statistics, 2023), but users with one do not cut back more than users without (Appendix A.9), and the data do not show why urban users cut back more. By gender the gap is similar in size, -0.95 (standard error 0.03) for women and -0.85 (standard error 0.03) for men, but gender is missing more often for senders who stopped transacting, so this comparison rests on a selected subsample (Panel A).

6 The levy and the network it taxed

Mobile money is a network good. An account is worth having in proportion to how many of a user's counterparties also hold one (Alvarez et al., 2025). When a user withdraws, the platform becomes less useful to everyone who transacted with them, a cost that the welfare analysis of Section 7, which treats each user's decision to transact less as private, does not count. The levy therefore did more than reduce volume. It thins the relationships the platform is made of, cutting their survival by a third relative to a no-policy benchmark; it stops new relationships from forming, at a rate that falls by sixty percent and stays there; and it stops the network from repairing itself around a lapsed relationship. The rest of this section documents these three facts on the complete transaction records of each provider. It then asks whether the levy's damage travels from one user to the next, by measuring how a sender's own activity moves when her counterparties go quiet, testing that reading against the alternative that similar users simply responded alike (Manski, 1993; Aral et al., 2009), and computing the share of the decline that ran through contacts going quiet.

Half the transacting relationships disappear and do not return. Figure 11 counts the distinct sender-receiver pairs that transact in each week, indexed to the week before the levy. The count falls by half within two weeks of the introduction, stays there through the September cut, and has not moved by mid-October. A count of active pairs can halve because the same relationships transact less often, which leaves them in place, or because relationships end, which does not. To tell the two apart, we take every pair with a person-to-person transfer between June 1 and July 14 and ask whether the same pair transacts again between July 15 and August 31. Many such relationships are one-off by nature, so the answer needs a benchmark, which we

take from pre-reform periods: of the pairs transacting between April 1 and May 14, about one in four (24.6 to 27.0 percent, depending on the provider) transact again between May 15 and July 1, and the other three in four go quiet with no policy change at all. Of the pairs transacting just before the levy, only about one in six (16.5 to 18.7 percent) transact again between July 15 and August 31 (Appendix Table A15; Appendix A.17 defines each measure in this section). The levy therefore lowers the chance that an existing relationship continues by roughly a third. So the halving is not only relationships transacting less often: a third fewer of them survive at all, and those that lapsed at the introduction do not resume at the reduced rate.

A relationship that ends is a loss to both of its parties, not only to the one who stopped: the other can reach one fewer person through the platform, a cost that the welfare calculation of Section 7, which counts only the loss to the user who transacts less, leaves out. The loss also outlives the rate that caused it. The September cut took back part of the increase in the price of a transfer, yet the pairs that had stopped do not resume, so the network the reduced rate applies to is the smaller one the introduction left behind, and the welfare cost at the reduced rate includes relationships lost at the higher one.

Three in ten senders stop transferring, and the rest keep a fifth fewer counterparties. Among senders with at least three pre-levy transfers, 31 percent make no person-to-person transfer at all between July 15 and August 31, and those who keep transferring keep fewer counterparties: the number of distinct receivers per active sender falls by 22 percent, averaged across the three providers, and by 21 to 24 percent at each (Appendix Table A5). The two margins together are the halving in Figure 11: seven in ten senders still active, each with a fifth fewer counterparties, is about half the pairs. Stopping transfers is not leaving the platform. Only 7 percent of these senders make no transaction of any kind, as sender or receiver, between July 15 and August 31, deposits, withdrawals, and merchant payments included, so about four in five of the senders who stopped transferring were still using the platform for something else (Appendix A.17). The levy took the taxed relationships and left the accounts in place, so the platform loses its transfer network faster than its users. The thinner network is borne by the users who stay: seven in ten keep paying the levy on a circle a fifth smaller than before.[20]

What remains reaches fewer people and is less tightly knit, but it holds together. The relationships that survive form a sparser network, not a broken one. Counting a sender who stopped as having no counterparties, the number of accounts she transacts with directly falls by 37 to 39 percent against the pre-reform benchmark. The number she can reach in two steps, through her counterparties and theirs, falls by more than half, because the counterparties she keeps have lost some of their own.[21] Her counterparties are also less often tied to one another. The share of them who transact with another of her counterparties, the support ratio of Chetty et al. (2022), falls by a fifth to a quarter, and it still falls among senders with the same number

[20] The loss is not confined to any part of the network either: receivers with dozens of senders lose the same share of them as receivers with one, once the tilt that busy accounts show in any window is netted out against the placebo (Appendix Table A6).

[21] Reach is measured for a fixed group of senders chosen for being active before the levy, so it also falls in the pre-reform benchmark, when 10 to 13 percent of the comparison group has no transfer, sent or received, in the following window. The range, 51 to 73 percent, covers the median, the mean, and the mean without the few accounts with 1,000 or more counterparties, each of which adds its whole circle to the reach of every sender linked to it (Appendix Table A19).

of counterparties, so smaller circles are not the whole story.[22] Yet most active accounts remain tied to the main body of the network through some chain of transfers: 73 to 82 percent after the levy, against 85 to 91 percent before. The levy thinned the network from within, leaving users connected through fewer and less overlapping relationships, rather than cutting it into separate pieces.

The network stops growing. New pairs formed per week fall by about sixty percent in the week of the introduction and remain there through the September rate cut (Appendix Figure A9). The deposit side mirrors the network: accounts that received person-to-person transfers from two to nine distinct counterparties before the levy cut their daily cash deposits by 0.42 to 0.49 log points, against 0.23 to 0.33 for accounts with no inbound person-to-person transfers at all, so money stopped entering the system where the network had been densest, not where users were isolated (Appendix Table A7).

New users still arrive, but few of them start sending. The fall in new relationships is not a fall in new users. At Tigo, whose registry records the day each account was activated, sign-ups run at about 14,900 a day in the six weeks before the levy and 14,700 in the seven weeks after, much as in the same weeks of 2020.[23] The change is in what new users do. The share who send a person-to-person transfer within two weeks of activating the account halves, from 4.2 to 2.1 percent, and stays there after the cut, while the share who receive one does not fall. A count of accounts would show a network still growing, while half as many of the accounts it adds start sending money to anyone.

The network stops repairing itself. A sharper test asks how a sender behaves when their own counterparties go quiet. For each sender with at least five pre-levy counterparties we track the share of those counterparties who have gone dormant, meaning no transfer to anyone in the preceding eight weeks,[24] and ask how the sender's own transfers to third parties move with that share, holding sender, period, and sender-volume-by-period effects fixed; we call this association contagion. Table 4 reports the results. Before the levy the association is positive: a ten percentage point rise in the dormant share comes with about 10 percent more third-party activity at Tigo, 8 percent at M-Pesa, and 7 percent at Airtel, the signature of users rerouting around relationships that lapse. After July 15 the association reverses, falling by 15 percentage points at Tigo, 11 at M-Pesa, and 9 at Airtel, with standard errors below half a percentage point, so a sender losing

[22] Senders are grouped by their number of counterparties in each window (two, three or four, five to nine, ten or more), so who is in each group changes between windows; within groups the support ratio falls by 9 to 28 percent against the pre-reform benchmark. A network-wide version, which weights each account by the number of pairs among its counterparties, shows no consistent change across providers, with or without the handful of accounts with more than 1,000 counterparties that dominate it.

[23] Tigo's registry passes two checks: 99 percent of the accounts in its transfer data are in it, and fewer than 0.1 percent of accounts registered between April and November 2021 transfer before their registration date. M-Pesa's registry fails the second check and Airtel's data include no registry, so the comparison is for Tigo only. In the same weeks of 2020, sign-ups run at about 10,800 and 10,400 a day (Appendix Table A20).

[24] Before the levy, a counterparty in a sender's circle who has not yet made any transaction is not a lapsed relationship but one that has not started, so her share of the circle enters the regression as its own variable and the quiet share counts only counterparties who were on the platform and went silent. After July 15 no counterparty is in that state, since every one in the circle had sent at least once by July 14. Senders whose circles held more such counterparties were more active, by about 19 percent per ten points at Tigo (Table 4).

counterparties now reduces activity with everyone else as well. Figure 12 shows the two cases. Figure 13 estimates the same association fortnight by fortnight. It is flat over the three pre-levy fortnights at Tigo and Airtel and rises by two points in the last of them at M-Pesa, changes sign in the first fortnight of the levy at Tigo and M-Pesa and in the second at Airtel, and stays there through the cut; the reversal is a break at July 15, not the end of a trend. The reversal is larger for senders with ten or more counterparties and holds for transfer counts, distinct receivers, and value. The September cut restores none of it: the post-cut change in contagion is under one percentage point at each provider (0.5 at Tigo, -0.6 at M-Pesa, and -0.0 at Airtel). In ordinary times the network repairs itself around a lapsed relationship; after the levy it does not, and the reduced rate does not revive the repair.

Dead relationships do not revive when the rate falls. If the cut made revival worthwhile where it was still possible, relationships that went silent should restart at higher rates in September, at least where both parties are still active with others. They do not. Among the 1.2 million pre-levy relationships at Tigo that fall silent between July 15 and August 3, about half have both parties still transacting with third parties. For those feasible pairs, the weekly rate of a first renewed transaction is already falling before the cut, from 3.2 percent four weeks before September 1 to 2.0 percent the week before, and it keeps falling afterward, to 1.5 percent four weeks later. M-Pesa shows the same path (2.8 to 1.7 to 1.2 percent). Pairs whose partner has stopped sending to anyone revive at lower rates on the same downward path. A lower price does not revive them, even where nothing but the price stands in the way. Against a straight-line decline, both groups step up slightly at the cut, by about a tenth of a percentage point, and the pairs that no price can revive step up more, so the step reflects the slowing of the decline rather than the cut.[25]

The reversal is not homophily. The natural objection to reading the reversal as contagion, one user's exit spreading to another, is that people transact with people like themselves, homophily, so a sender and a quiet counterparty may have withdrawn for the same reason: everyone faced the same tax, and people who resemble each other may have responded to it alike. One test separates the two readings, because they make opposite predictions about which counterparties matter. If homophily is causing the response, what matters is whether the sender and the counterparty are alike, so measured contagion should be strongest among the counterparties who most resemble her and fade as resemblance falls. If contagion is causing the response, what matters is that the counterparty stopped sending, not whether she resembles the sender, so measured contagion should not fade. We measure resemblance as a distance on the traits users sort on: at Tigo, where the registry records where users live, whether the two are registered in the same district; everywhere, how far apart their typical transfer sizes are. Appendix Table A8 reports the change in measured contagion at the levy by bin of each distance. It does not fade as homophily predicts. On district, counterparties elsewhere

[25] Pairs from the contagion-design rosters (senders with at least five observable pre-levy counterparties). A pair is dead if it has no transaction from July 15 to August 3 and feasible if both endpoints send to third parties in that window; the hazard is the share of pairs still dead at the start of a week that transact in it, for the four weeks either side of September 1. Pairs whose partner has stopped sending to anyone serve as the comparison: their hazard follows the same path, and the gap between the two series narrows across the cut rather than widening. Fitted with a linear week trend, the step at September 1 is 0.07 percentage points for feasible pairs at Tigo ($t = 2.7$) and 0.09 at M-Pesa ($t = 4.3$), against 0.10 and 0.12 for pairs whose partner has stopped sending to anyone.

carry over nine tenths of the measured contagion of counterparties in the sender's own district. On size it fades mildly at Tigo and M-Pesa and not at all at Airtel, and the counterparties whose transfers differ from the sender's by more than a factor of four still carry a reversal of 9 to 13 percent per ten points. Resemblance on what we observe accounts for at most about a third of measured contagion; the rest runs through counterparties who share nothing we can see with the sender, which is what contagion predicts and homophily cannot produce. The association therefore behaves as contagion: what moves a sender is that a counterparty went quiet, not who the counterparty is. A second test looks at the counterparties of a sender's counterparties, people she does not transact with. Contagion predicts that their going quiet reaches her through her circle; homophily predicts little, since she has less reason to resemble them than the people she pays. The data follow the first: holding the quiet share of her own circle fixed, the association between a sender's third-party transfers and the quiet share one step further out falls after the levy by 7 percentage points at Tigo, 4 at M-Pesa, and 4 at Airtel per ten points of that share (Appendix A.18).

We run three additional tests to bound what homophily could contribute and to rule out two mechanical alternatives. First, we bound how much of measured contagion could run through the traits we observe. We add to the main specification the typical transaction volume of each sender's counterparties and the typical size of their transfers, each interacted with the fortnight, so that contagion is identified only from counterparties going quiet within circles of the same kind. Volume alone removes about a fifth of measured contagion (19 percent at Tigo, 24 at M-Pesa, and 21 at Airtel), size alone a fifth, a quarter, and 3 percent, and the two together between a fifth and two fifths. By the coefficient-stability and Oster bounds of Appendix A.18, a trait we do not observe would have to select users one to four times as strongly as these two do to remove the rest. Second, we test for a trend already under way. If senders whose circles were thinning had been cutting back before the levy, measured contagion would drift before July 15; estimated fortnight by fortnight, it is flat over the three pre-levy fortnights at Tigo and Airtel, rises by two points in the last of them at M-Pesa, and reverses within two fortnights of the levy (Figure 13). Third, we test whether contagion runs the other way, from the sender's own cutback to the people she supported. Among counterparties who received less than a tenth of their pre-levy inflows from the sender, who cannot have gone quiet because she cut back, the association falls by 19 percentage points at Tigo, 13 at M-Pesa, and 6 at Airtel, three quarters of the change on the same senders' full circles at Tigo and M-Pesa and half at Airtel (Appendix A.18). We therefore interpret the reversal as causal, one user's withdrawal lowering the platform's value to the next. Applied to the share of each circle that has gone quiet, measured contagion gives the part of the decline that ran through contacts going quiet: by October, counting the contacts of a sender's contacts as well, about a fifth of the decline at Tigo and M-Pesa pooled (18 percent between September 9 and October 6; Appendix Table A9). We report the share by October because a contact counts as quiet only after eight weeks without a transfer, so contacts who stopped at the introduction register only from mid-September and the share measured earlier understates the effect. We report the lower end of the range, which uses transfers to people outside the sender's pre-levy circle and so cannot mechanically contain the quiet contacts. What carries it is contacts going quiet during the levy: senders whose circles were already thin on July 14 cut back no more than others (Appendix A.18).

7 Welfare, revenue, and the cost of public funds

7.1 The cost of the levy per shilling raised

A tax destroys value whenever it stops a transaction that was worth making without the tax but not with it: the person loses what the transaction was worth to her, somewhere between nothing and the tax, and the government collects nothing on it. The standard linear approximation values each stopped transaction at half the tax. The surplus lost is therefore $\frac{1}{2}\tau \cdot \Delta Q$, where τ is the levy as a share of transaction value and ΔQ is the taxed value the levy stopped, while the revenue is $\tau \cdot Q_1$, the levy on the value that remains; the ratio of the two, $\frac{1}{2}(\Delta Q/Q_0)/(1-\Delta Q/Q_0)$, depends only on how far the base is down. Taxed value stands 48.2 percent below its pre-levy level over the high-rate weeks, so the ratio is 0.47: the levy destroys 47 cents of surplus for every shilling it raises. The figure counts the whole fall in taxed activity as the levy's doing. Measured against the untaxed control instead, the fall is 33 percent (the main coefficient of -0.40) and the ratio is 0.25; because the control was itself depressed by the levy (Section 4), this figure is a lower bound. Two adjustments move the 47 cents in opposite directions, to 42 and to 74, and weighting the taxed series by value rather than equally would give 35; Appendix A.20 reports these. The evidence that the untaxed basket's own decline is the levy's spillover rather than a shock the two sides share, which is what counting the whole fall requires, is in Appendix A.12.

Cutting the rate lowers the cost per shilling only slightly. Under the reduced rate the taxed series are 45 percent below their pre-levy level, which corresponds to a cost of 41 cents, against 48 percent and 47 cents at the introduction rate. The cut gave up 30 percent of the rate and left the base nearly where the introduction had put it. Using aggregate data, we show that the two further cuts of 2022 did the same: transfer value in 2022 was 47 percent of the level its pre-levy trend implied (Section 5.3). Because the cost per shilling depends on how far the base is down and not on the rate, cutting the rate lowered revenue and barely lowered the cost of each shilling still raised.[26]

Deadweight loss as a share of revenue for a typical VAT in lower-income settings is estimated at well below 10 percent (Brockmeyer et al., 2024); the mobile money levy is several times to an order of magnitude above that benchmark.

Note that these figures understate the social cost of the tax. They count only surplus lost on taxed transactions, not the untaxed bases that contract because of the levy, and they treat each user's pre-levy activity as the right amount from society's point of view, when on a payment network it was probably too little, because part of the benefit of a user's activity goes to the people they transact with, who do not pay for it. In the closest calibrated study of person-to-person payments, Alvarez et al. (2025) find that the value a Costa Rican user gets from other users being on the platform is roughly six times the value of the platform on its own, enough that a planner would subsidize use by about 84 percent of the average user's benefit.

[26] With $\Delta Q/Q_0 = 0.451$ the ratio is $0.5 \cdot 0.451/0.549 \approx 0.41$. On the differential response the reduced-rate figure is 0.29, above the introduction's 0.25, because the untaxed control recovers more after September than the taxed series do (Section 5).

7.2 Revenue, the portfolio, and the alternatives

The levy's revenue was relatively small and fell fast because of the subsequent cuts to it. In its first two months of collections, August and September 2021, the line item "Transaction Levy on mobile money for sending and withdrawing" was about 2 percent of Tanzania Revenue Authority collections net of refunds; within three fiscal years, after the September 2021 cut, two further amendments in July and October 2022 that lowered the schedule, the removal of transfers from the levy in July 2023, and a base that stayed below its pre-levy trend, it was 0.8 percent (FY 2023/24). Over FY 2021/22 through FY 2023/24 the levy raised 288 billion TZS a year (Figure 5): about 5 percent of what the value-added tax raised net of refunds (6,045 billion), 10 percent of the payroll tax (PAYE, the income tax withheld from wages; 2,872 billion), and roughly seven times the property tax administered by the TRA (39 billion).[27] The seven-to-one ratio reflects what is currently collected from each instrument, not what could be collected at lower behavioral cost.

The levy also cost the government part of the revenue from the taxes already charged on provider fees, the 10 percent excise duty and the 18 percent VAT (Section 2.2). Fewer transactions meant fewer fees: at M-Pesa and Tigo, daily fee revenue on taxed transactions is 40 log points below its pre-levy level over the high-rate weeks and 38 after the cut, less than taxed value at the two providers (54 and 47 log points), because cash-out, which carries four fifths of the fees, fell least. The excise and VAT lost on these fees equal 16 percent of the levy collected over the high-rate weeks and 20 percent after the cut, or 12 and 24 percent when the fall in fees is measured against the fees on untaxed transactions.[28] Net of this loss, the levy raised an eighth to a quarter less than it collected, which would raise the marginal value of public funds computed below by 14 to 31 percent (Appendix A.21).

Tanzania's own retreat points to the main alternative. Two thirds of the levy's revenue came from withdrawals (68 percent, against 26 from person-to-person transfers and 7 from inter-operator transfers), and the Finance Act of 2023 removed the levy from transfers, leaving a withdrawal-only levy after two years of retreat. Other governments are nonetheless still adopting the instrument Tanzania began with: since 2025 Senegal, Malawi, Mali, Niger, and Gabon have legislated levies on the value of mobile money transactions, mostly at lower rates or on narrower bases than Tanzania's. The withdrawal-only design keeps most of the revenue while sparing the transfers that carry the network (Section 6). Matching the full revenue would require withdrawal rates about 1.6 times the enacted ones once withdrawals respond, with an elasticity of about 10 that we estimate imprecisely. Table 5 sets out what three other designs would require and who would bear them: a flat monthly charge per active account, a higher tax on provider fees, which would fall mostly on withdrawals because their fees are several times those on transfers, and stronger VAT enforcement.

[27] Recurrent property tax revenue in sub-Saharan Africa averages around 0.3 percent of GDP, roughly a third of the OECD average and well below the 2 to 3 percent collected in some advanced economies, so the base is underused rather than exhausted; cross-country averages from Norregaard (2013) and International Monetary Fund (2024); earlier estimates in Bahl and Martinez-Vazquez (2008) are broadly consistent.

[28] Fees are summed over every M-Pesa and Tigo transaction from June 1 to October 13, 2021; Airtel's data record no fees. From July 15 Tigo records the levy inside the fee, and we subtract it. The recorded fee includes the excise and VAT (Section 2.3), so the two taxes are $0.298/1.298$, or 23 percent, of it; if it excluded them, the shares would be 21 and 26 percent. Tigo stopped offering free on-network transfers on about June 25, 2021, which raised its fee revenue late in the pre-levy weeks, so the loss is slightly larger than measured. To the extent that users spent the fees they saved on goods that pay VAT, part of the VAT loss is recovered; the excise is not.

We do not rank them, because their behavioral costs depend on responses our data cannot measure.

The same cost can be stated as a marginal value of public funds (Hendren and Sprung-Keyser, 2020; Finkelstein and Hendren, 2020). The MVPF of a tax is what taxpayers would pay to avoid it per shilling of net revenue it raises; with a behavioral response that removes a share x of the revenue the old base would have yielded, it is $1/(1-x)$, which under the linear approximation of Section 7.1 equals one plus twice the deadweight loss per shilling. The levy's 47 cents therefore becomes an MVPF of 1.93 at the introduction rate and 1.82 at the reduced rate (Appendix A.21).

Table 6 sets these figures against the value-added tax and the payroll tax, for which we calibrate x from the statutory rates and from base elasticities in the literature on comparable settings, 0.5 for the VAT and 0.4 for the payroll tax. The implied MVPF is 1.12 for the VAT and 1.21 for the payroll tax. A shilling raised by the levy therefore costs users 50 to 60 percent more than a shilling of payroll tax, the more distortionary of the two, and 62 to 72 percent more than a shilling of VAT; even the lower bound is a quarter above the payroll tax. In the unit of Section 7.1, the VAT and the payroll tax destroy about 6 and 10 cents of surplus per shilling raised, against the levy's 47. The ranking would flip only if the payroll tax's elasticity were about 1.05 rather than 0.4, or about 0.77 to flip against the lower bound. The gap is what taxing a transaction rather than income or consumption does. A transfer is a means of moving purchasing power, not the purchasing power itself, and the same money moves by cash at no tax and little cost, so a levy on the transaction meets a substitute that a tax on wages or on consumption does not (Section 5). The size of the response is a property of that choice of base rather than of Tanzania, so the comparison carries over to any transaction levy on mobile money. A second difference, that the base is a network, is taken up at the end of this subsection.

Under equal welfare weights, pairing a tax with a spending program raises welfare only if the spending's MVPF exceeds the tax's, so revenue raised through the levy is worth raising only if it funds spending valued above 1.8 to 1.9 shillings per shilling spent. The levy was presented at introduction as financing rural health centres, classrooms, and roads,[29] and direct investment in the health and education of low-income children is the category in which Hendren and Sprung-Keyser (2020) find the highest MVPFs, often high enough to clear a threshold of this size. We therefore do not claim that the levy was poor policy; the case for the levy rests on the return to that capital, which we do not estimate.

Finally, we incorporate the cost of taxing a network into the MVPF formula. The levy's figures count each user's response as her own. On a network part of a user's response is caused by her counterparties' responses, and the loss it imposes on her is neither chosen nor counted. With a share s of the response arising this way, the burden on users per shilling of revenue is $1/(1-s)$ times the standard figure, the social multiplier of Glaeser et al. (2003), so the MVPF becomes $1/[(1-x)(1-s)]$ (Appendix A.23). With the share of the decline that spread from user to user at 18 percent by October (September 9 to October 6; Appendix Table A9), the MVPF rises from 1.93 to 2.4 at the introduction rate and from 1.82 to 2.2 at the reduced rate.

[29] Contemporaneous reporting of government statements describes the first four weeks of collections, approximately 48 billion TZS, as funding roughly 150 health centres and 560 classrooms, with subsequent collections directed to the rural roads agency. Tanzanian press reporting of government statements in August and September 2021 (*The Citizen*; *Daily News*; *The Chanzo*, 31 August 2021). These are reported figures rather than audited budget execution, and the reporting dates collections from 1 July while the levy took effect on 15 July, so the window is approximate. We use them to characterize the intended use of the revenue, not to measure it.

We apply the October share at both rates because the eight-week rule registers contacts who stopped at the introduction only from mid-September, so the share measured in the high-rate weeks understates it.

8 Conclusion

In this paper we estimate how a fully observable tax base responds when it is taxed and when the tax is eased, using two policy shocks seven weeks apart: the introduction of a notched levy on mobile money transfers and withdrawals in Tanzania on July 15, 2021, and a cut of about 30 percent on September 1. On the complete data from Tanzania's three largest providers, the introduction reduces the volume of taxed transactions by 23 percent and their mean value by 13 percent relative to untaxed transactions on the same platform, an elasticity of about 10 with respect to the share of a transfer the sender keeps. The cut brings none of it back, and neither do the two further cuts of 2022; only the removal of the levy on transfers in 2023 did, and slowly. The levy thins the payment network itself, so that part of the decline spreads from one user to another, and its reach runs beyond its statutory base: payments from businesses into mobile wallets, exempt by statute, fall by about a tenth at the two providers that carry nearly all of them.

The administrative case for taxing mobile money in low-income economies is strong: every transaction is captured electronically by a small number of regulated providers, the tax is withheld at source, and collection costs are near zero. The technology removes the observability and enforcement constraints that bind most low-income tax bases. Our estimates suggest that removing the administrative constraint is not enough: the base is highly responsive to being taxed, and the response does not undo itself when the rate falls. The levy destroys 47 cents of surplus for every shilling it raises (Section 7), several times the cost of a typical VAT in a lower-income setting, and since part of the decline spread from user to user, imposing losses on users who did not change their own behavior, even that figure is a lower bound.

Taxes of this kind are spreading fast in sub-Saharan Africa. Five governments have legislated levies on mobile money since 2025. Nothing in the administration forces the choice. We show that there are designs that would raise comparable revenue at a fraction of the behavioral cost, a flat access charge or a higher excise on provider fees, which are collected by the same operators through the same systems (Section 7). What a levy on transactions offers instead is a rate that looks small on a base that looks large: matching its revenue would take a tax of about two thirds of the value of provider fees or a visible monthly charge on every one of 18.6 million accounts (Table 5), while the levy is paid a few percent at a time by those who transact.

References

Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2020). Sampling-based versus design-based uncertainty in regression analysis. *Econometrica* 88(1), 265–296.

Altonji, J. G., T. E. Elder, and C. R. Taber (2005). Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools. *Journal of Political Economy* 113(1), 151–184.

Alvarez, F., D. Argente, F. Lippi, E. Méndez, and D. Van Patten (2025). Strategic complementarities in a dynamic model of technology adoption: P2P digital payments. NBER Working Paper 31280, December 2025 revision; conditionally accepted, *American Economic Review*.

Apeti, A. E. and E. D. Edoh (2023). Tax revenue and mobile money in developing countries. *Journal of Development Economics* 161, 103014.

Aral, S., L. Muchnik, and A. Sundararajan (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences* 106(51), 21544–21549.

Baca-Campodónico, J., L. de Mello, and A. Kirilenko (2006). The rates and revenue of bank transaction taxes. OECD Economics Department Working Paper 494, OECD Publishing, Paris.

Bachas, P., A. Jensen, and L. Gadenne (2024). Tax equity in low- and middle-income countries. *Journal of Economic Perspectives* 38(1), 55–80.

Bachas, P., O. Okunogbe, M. Shaukat, and D. Tortarolo (2026). *Raising Revenue Right: A Roadmap for Domestic Resource Mobilization*. Policy Research Report. Washington, DC: World Bank. https://doi.org/10.1596/978-1-4648-2034-2.

Bahl, R. and J. Martinez-Vazquez (2008). The property tax in developing countries: Current practice and prospects. In R. Bahl, J. Martinez-Vazquez, and J. Youngman (Eds.), *Making the Property Tax Work: Experiences in Developing and Transitional Countries*. Cambridge, MA: Lincoln Institute of Land Policy.

Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006). Who's who in networks. Wanted: The key player. *Econometrica* 74(5), 1403–1417.

Barczay, M., S. Hebous, F. Sawadogo, and J.-F. Wen (2025). Taxing mobile money: Theory and evidence. IMF Working Paper WP/25/255.

Batista, C. and P. C. Vicente (2025). Is mobile money changing rural Africa? Evidence from a field experiment. *Review of Economics and Statistics*, forthcoming.

Benzarti, Y., D. Carloni, J. Harju, and T. Kosonen (2020). What goes up may not come down: Asymmetric incidence of value-added taxes. *Journal of Political Economy* 128(12), 4438–4474.

Best, M. C., A. Brockmeyer, H. J. Kleven, J. Spinnewijn, and M. Waseem (2015). Production versus revenue efficiency with limited tax capacity: Theory and evidence from Pakistan. *Journal of Political Economy* 123(6), 1311–1355.

Björkegren, D. (2019). The adoption of network goods: Evidence from the spread of mobile phones in Rwanda. *Review of Economic Studies* 86(3), 1033–1060.

Brockmeyer, A., G. Mascagni, V. Nair, M. Waseem, and M. Almunia (2024). Does the value-added tax add value? Lessons using administrative data from a diverse set of countries. *Journal of Economic Perspectives* 38(1), 107–132.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *American Economic Review* 99(4), 1145–1177.

Chetty, R., M. O. Jackson, T. Kuchler, J. Stroebel, et al. (2022). Social capital I: Measurement and associations with economic mobility. *Nature* 608, 108–121.

Coelho, I., L. Ebrill, and V. Summers (2001). Bank debit taxes in Latin America: An analysis of recent trends. IMF Working Paper 01/67, International Monetary Fund, Washington, DC.

Conley, T. G. and C. R. Taber (2011). Inference with "difference in differences" with a small number of policy changes. *Review of Economics and Statistics* 93(1), 113–125.

Crouzet, N., A. Gupta, and F. Mezzanotti (2023). Shocks and technology adoption: Evidence from electronic payment systems. *Journal of Political Economy* 131(11), 3003–3065.

Demirgüç-Kunt, A., L. Klapper, D. Singer, and S. Ansar (2022). *The Global Findex Database 2021: Financial Inclusion, Digital Payments, and Resilience in the Age of COVID-19*. Washington, DC: World Bank.

Economides, N. and P. Jeziorski (2017). Mobile money in Tanzania. *Marketing Science* 36(6), 815–837.

Finkelstein, A. (2009). E-ZTax: Tax salience and tax rates. *Quarterly Journal of Economics* 124(3), 969–1010.

Finkelstein, A. and N. Hendren (2020). Welfare analysis meets causal inference. *Journal of Economic Perspectives* 34(4), 146–167.

Fjeldstad, O.-H., S. N. Ingholm, L. Katera, E. Løstegård, I. H. Sjursen, V. Somville, and J. Vietz (2026). Tax salience: Experimental evidence from Tanzania. *International Tax and Public Finance*, forthcoming. https://doi.org/10.1007/s10797-026-09957-5.

Garriga, P. and D. Tortarolo (2024). Firms as tax collectors. *Journal of Public Economics,* Elsevier, vol. 233(C).

Glaeser, E. L., B. I. Sacerdote, and J. A. Scheinkman (2003). The social multiplier. *Journal of the European Economic Association* 1(2–3), 345–353.

GSMA (2025). *The State of the Industry Report on Mobile Money 2025*. London: GSMA.

Hanson, A., E. P. Hargaden, and M. Harris (2025). Static and dynamic taxation of network goods. Working paper, November 2025.

Hearson, M., P. Mader, M. Abounabhan, M. Carreras, A. Diouf, A. Lees, H. Niesten, F. Santoro, and C. Wales (2024). Taxing mobile money in Africa: Risk and reward. ICTD Policy Brief 10, Institute of Development Studies, Brighton.

Hendren, N. and B. Sprung-Keyser (2020). A unified welfare analysis of government policies. *Quarterly Journal of Economics* 135(3), 1209–1318.

Higgins, S. (2024). Financial technology adoption: Network externalities of cashless payments in Mexico. *American Economic Review* 114(11), 3469–3512.

Innovations for Poverty Action (2023). *Transaction Cost Index: Year 1 Comparative Report*. New Haven, CT: IPA.

International Monetary Fund (2024). How to design and implement property tax reforms. IMF How-To Note 2024/006, Washington, DC.

Jack, W. and T. Suri (2014). Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution. *American Economic Review* 104(1), 183–223.

Klapper, L., D. Singer, L. Starita, and A. Norris (2025). *The Global Findex Database 2025: Connectivity and Financial Inclusion in the Digital Economy*. Washington, DC: World Bank.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics* 8, 435–464.

Kleven, H. J. and M. Waseem (2013). Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan. *Quarterly Journal of Economics* 128(2), 669–723.

MacKinnon, J. G. and M. D. Webb (2017). Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics* 32(2), 233–254.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *Review of Economic Studies* 60(3), 531–542.

Ministry of Finance and Planning, Tanzania (2021). Speech by the Minister for Finance and Planning, Hon. Dr. Mwigulu Lameck Nchemba Madelu (MP), presenting to the National Assembly the estimates of government revenue and expenditure for 2021/22. Dodoma, June 10, 2021.

National Bureau of Statistics, Tanzania (2023). *National Panel Survey, 2020–2021, Wave 5*. Ref. TZA_2020_NPS-R5_v02_M. Dar es Salaam: NBS. Distributed by the World Bank Living Standards Measurement Study, Integrated Surveys on Agriculture (LSMS-ISA).

Norregaard, J. (2013). Taxing immovable property: Revenue potential and implementation challenges. IMF Working Paper WP/13/129.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business and Economic Statistics* 37(2), 187–204.

Paul, R. W. and D. Sharma (2025). The impact of a mobile money levy on household welfare: Evidence from Tanzania. World Bank Policy Research Working Paper 11228.

Rivadeneyra, I., D. D. Suthers, and R. Juarez (2022). Mobile money networks with tax-incentives. *Humanities and Social Sciences Communications* 9, 62.

Slemrod, J. and W. Kopczuk (2002). The optimal elasticity of taxable income. *Journal of Public Economics* 84(1), 91–112.

Suri, T. (2017). Mobile money. *Annual Review of Economics* 9, 497–520.

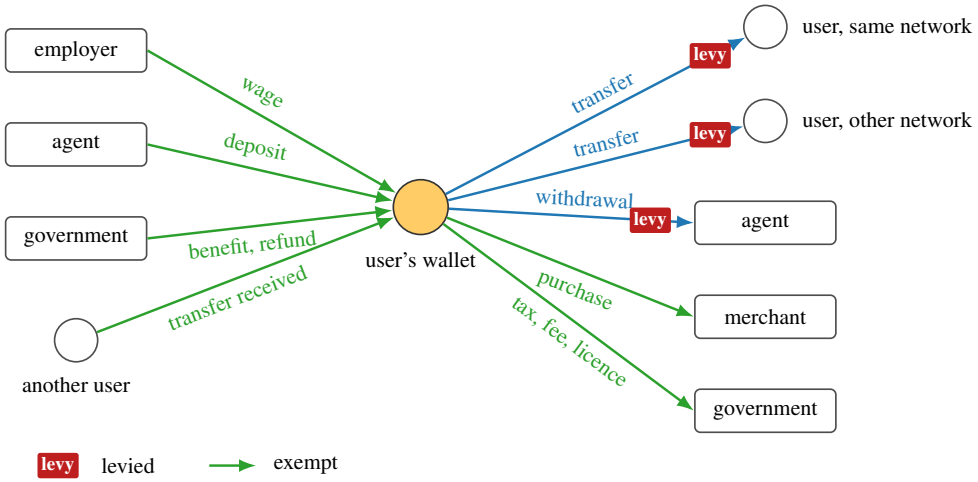
Tanzania Communications Regulatory Authority (2021). Quarterly communications statistics report, quarter ending December 2021. Dar es Salaam. URL: https://www.tcra.go.tz/tcra-tovuti/2021/mamlaka-website/documents/quarterly-statistics/Communications_Statistics_December_2021_1_1645554742_4eac16852a.pdf.

Wambura, B. (2021). Government considering a U-turn on mobile money transaction charges. *The Citizen* (Dar es Salaam), July 19, 2021.

World Bank (2023). *Tanzania Economic Update, Issue 19: The Efficiency and Effectiveness of Fiscal Policy in Tanzania*. Washington, DC: World Bank.

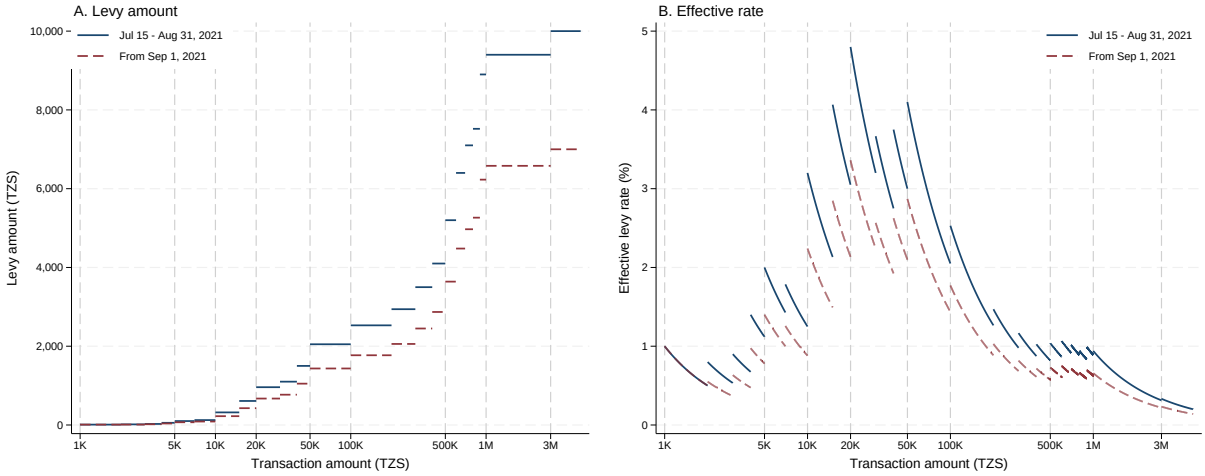
Figures

Figure 1: The levy's base



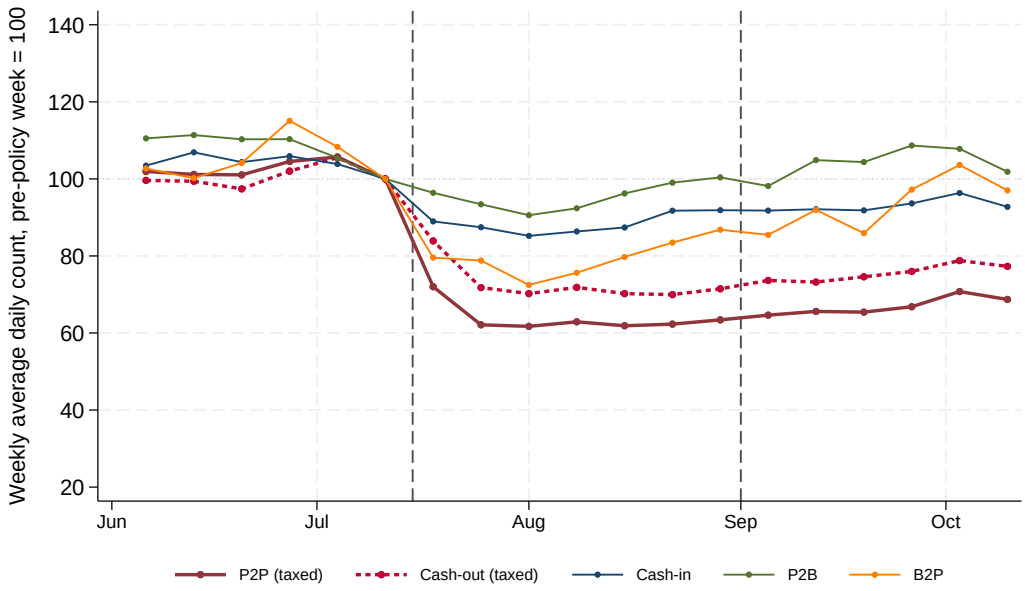
Notes: Flows through a user's mobile money wallet under the schedule in force from July 15, 2021. Person-to-person transfers, on the same network or to another network, and cash withdrawals at an agent were levied, at a fixed amount per transaction of 10 to 10,000 TZS in 23 brackets (Appendix Table A12), paid by the user sending or withdrawing. Deposits, wages and other business-to-person payments, payments from and to government, merchant payments and transfers received were exempt. Section 4 uses the exempt types as the control group.

Figure 2: The levy schedule under the two regimes



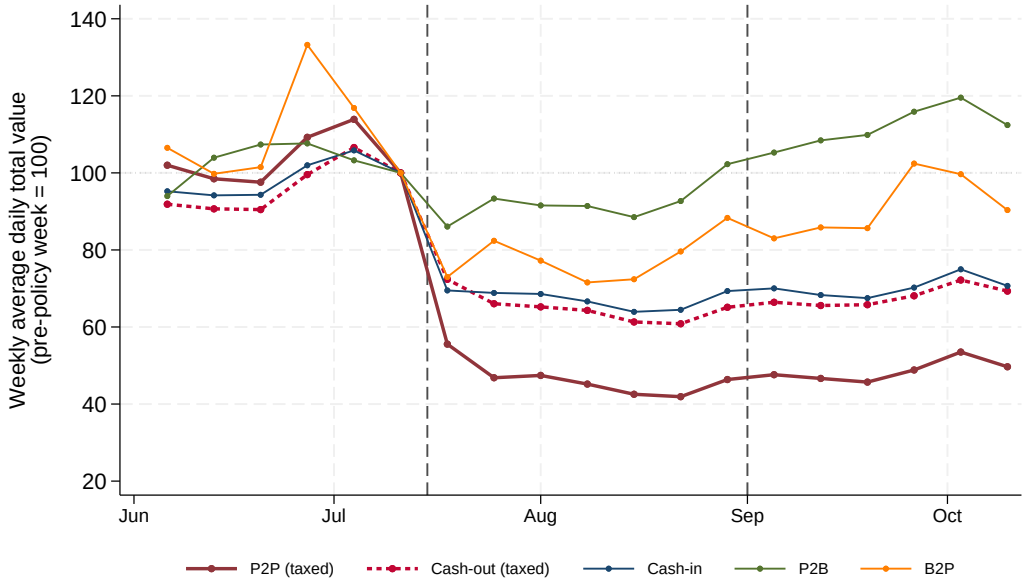
Notes: Panel A plots the fixed levy amount in each of the 23 statutory brackets (Table A12); panel B plots the same schedule as a share of the transaction value. Solid lines: Government Notice 496A, in force July 15 to August 31, 2021. Dashed lines: Government Notice 642A, in force from September 1, 2021. The horizontal axis is logarithmic; the open-ended top bracket is drawn to 5 million TZS.

Figure 3: Daily transaction counts by type around the two policy dates



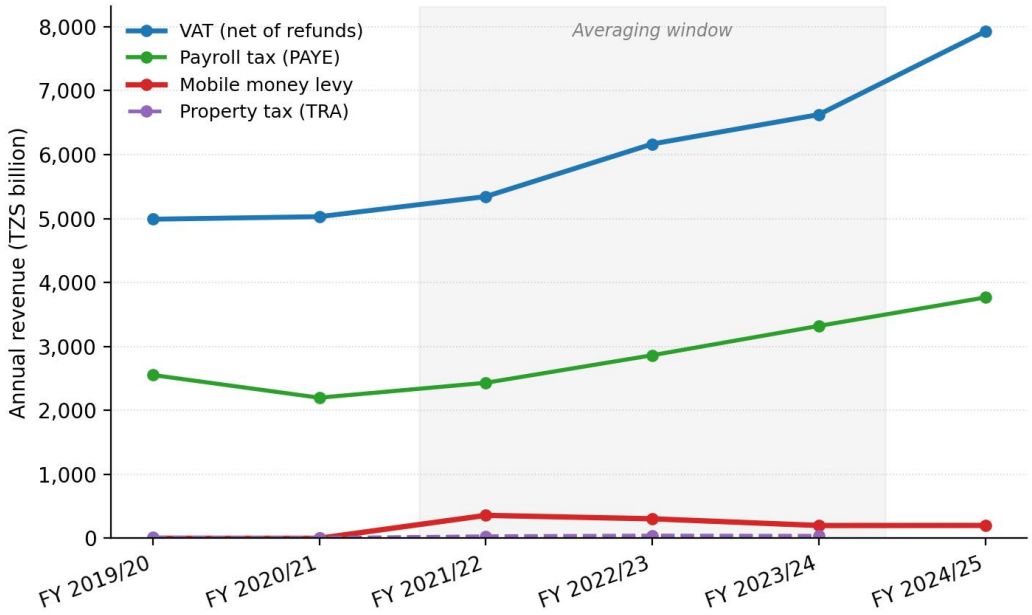
Notes: Daily count of transactions by type, normalized to 100 on the reference week of July 8–14, 2021; markers are non-overlapping seven-day averages, taxed series in heavier lines. Assignment as in the estimation (Appendix A.2): person-to-person transfers pool on-network and inter-operator transfers, the merchant-payment series keeps only genuine merchant payments. Three providers pooled, except cash-in (M-Pesa and Tigo only); payments to government omitted (Airtel only in this window); the window ends October 13, before the Airtel duplication artifact, and business-to-person payments omit three artifact days (Appendix A.1), transfers arriving from other networks, and Vodacom's payouts from its own accounts (Section 5.4). Dashed lines: July 15 introduction and September 1 cut.

Figure 4: Daily total transaction value by transaction type around the two policy dates



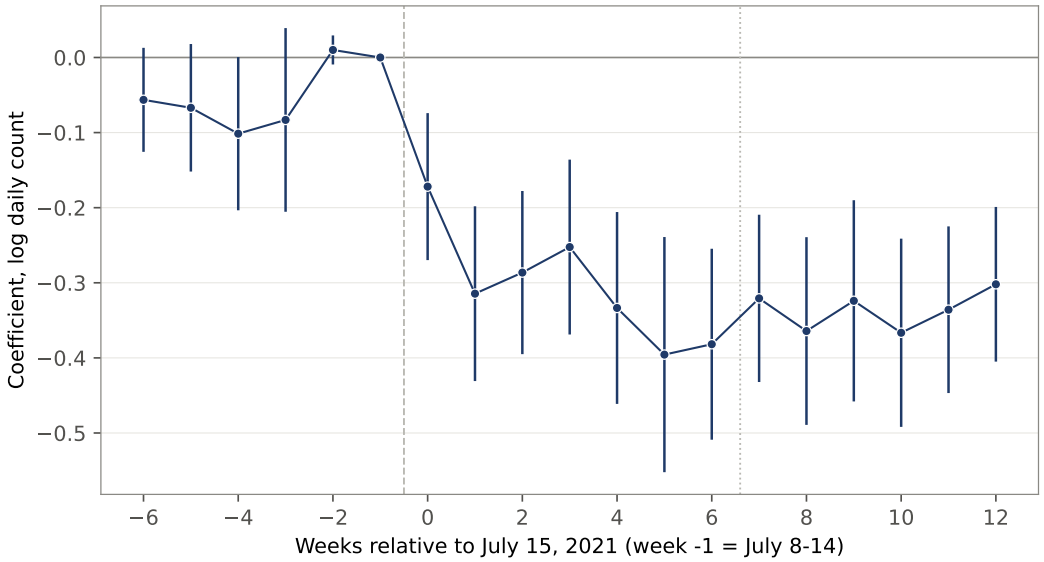
Notes: Weekly average of daily total transacted value, normalized to 100 on July 8–14, 2021; series, providers, and window as in Figure 3. No days are dropped: the M-Pesa duplication artifact splits bulk payments into several entries and leaves total value unchanged. Person-to-person transfers fall to about 47 percent of the reference week within three weeks and to 42 percent in late August, and cash-out to 61 percent, both steeper than the count-only falls in Figure 3. Dashed lines: July 15 introduction and September 1 cut.

Figure 5: Mobile money levy revenue against the main pillars of Tanzanian tax collection



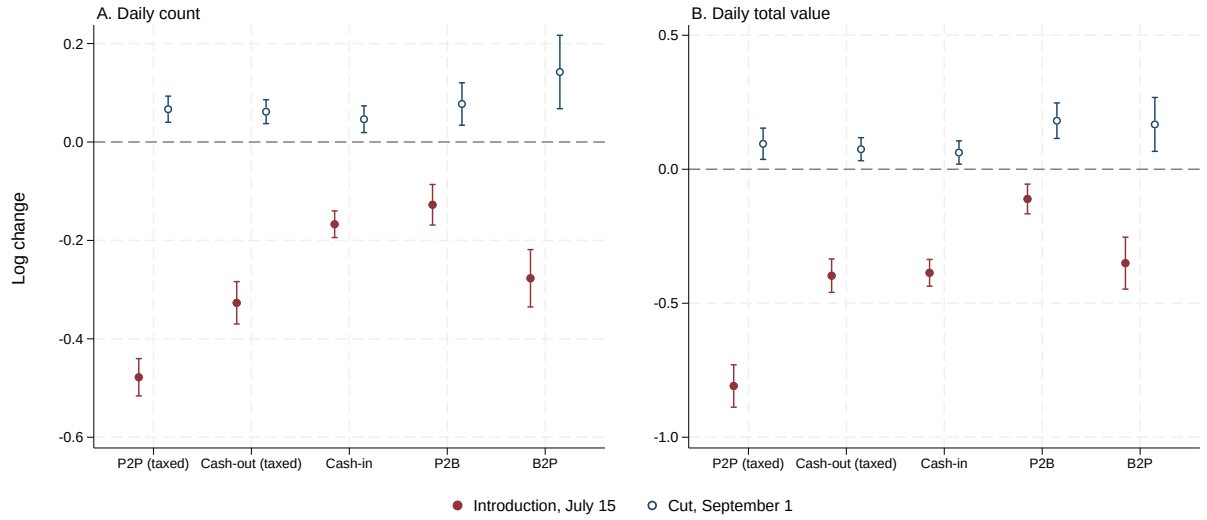
Notes: TRA-collected revenue by fiscal year, 2019/20 to 2024/25, TZS billion. The shaded band is the three-year window over which the text's comparisons are averaged. The property tax series excludes property tax collected by local government authorities, is not broken out in TRA workbooks before FY 2019/20, and ends when administration moved to local authorities at the end of FY 2023/24. Source: TRA quarterly tax revenue collections.

Figure 6: Event-study estimates around the levy introduction



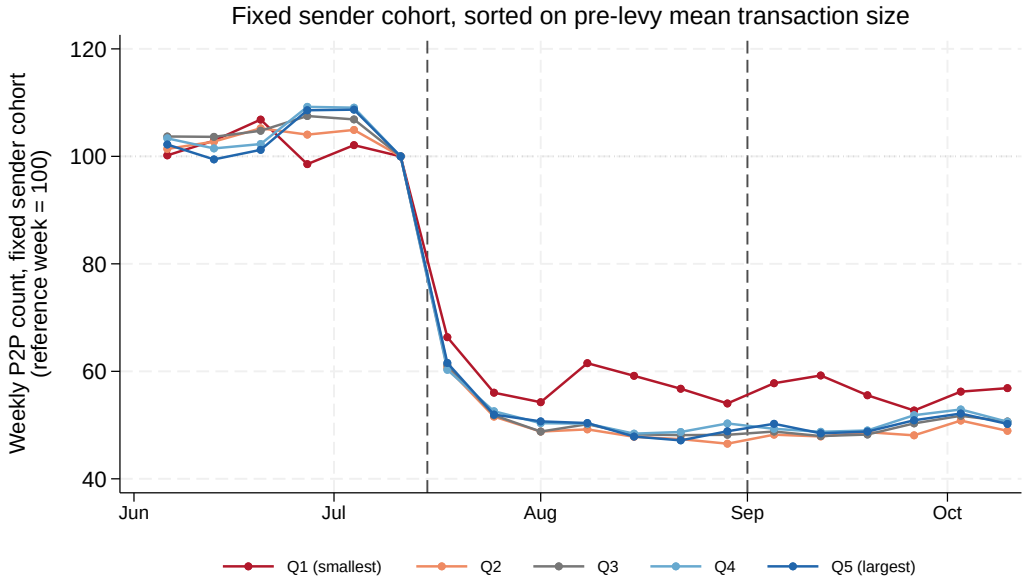
Notes: Weekly leads and lags of β in equation (1), main specification (Panel A of Table 2), on the daily count, one series per provider and transaction type with series and provider-by-day fixed effects, June 3 to October 13, 2021; week -1 (July 8–14) is the omitted base; bars are 95 percent confidence intervals from standard errors clustered by series. Dashed line: July 15 introduction; dotted line: September 1 cut. The coefficients drop to -0.17 at week 0 and lie between -0.25 and -0.40 through the end of August. Before the levy they lie between -0.10 and $+0.01$, a movement that disappears once the taxed series are given their own day-of-month profile (Appendix A.4).

Figure 7: The introduction and the cut, series by series



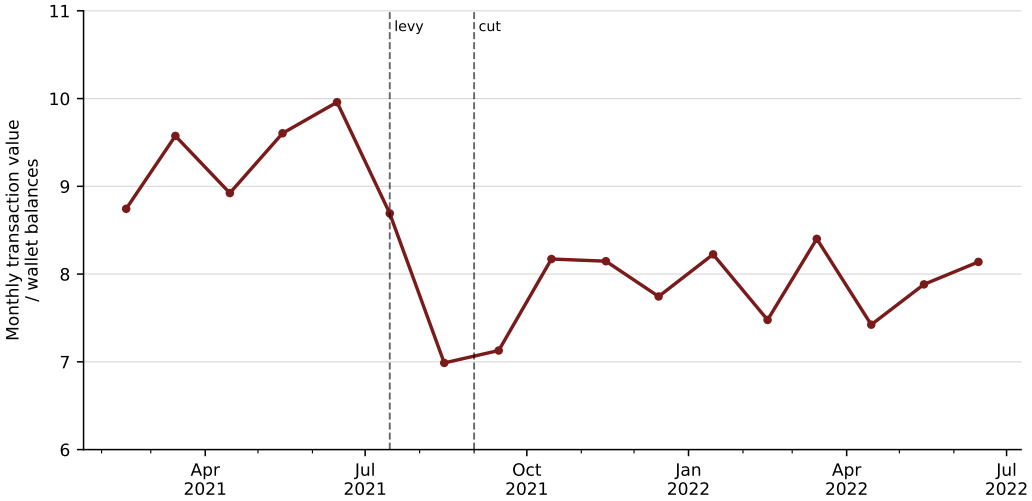
Notes: Single-series break coefficients from the complete transaction records, three providers pooled, series assigned as in Figure 3. Filled markers: coefficient on an indicator for July 15 onward in a regression of the log daily outcome on that indicator, June 1 to August 31, 2021. Hollow markers: coefficient on an indicator for September 1 onward, August 1 to October 13, 2021. These are the windows of Table 2, so the difference between a taxed and an untaxed marker is close to the difference-in-differences of that table, which compares series within provider. The figure is in log changes rather than elasticities because the untaxed series have no price change to divide by. For the taxed series the introduction implies an elasticity of about 10 for transfers and about 10 for withdrawals (Section 7); we do not convert the September coefficients, because the cut changed the log net-of-tax price by 0.012, an order of magnitude less than the introduction, so the ratio would be dominated by the coefficient's standard error. Panel A: daily count, with business-to-person payments omitting the three M-Pesa artifact days (Appendix A.1); Panel B: daily total transacted value, no days dropped. Bars are 95 percent intervals from Newey-West standard errors with a seven-day lag; like the other single-series standard errors in the paper, they understate uncertainty along time (Appendix A.3). Payments to government omitted (Airtel only in this window).

Figure 8: Weekly person-to-person transaction count for a fixed cohort of senders, by pre-policy size quintile



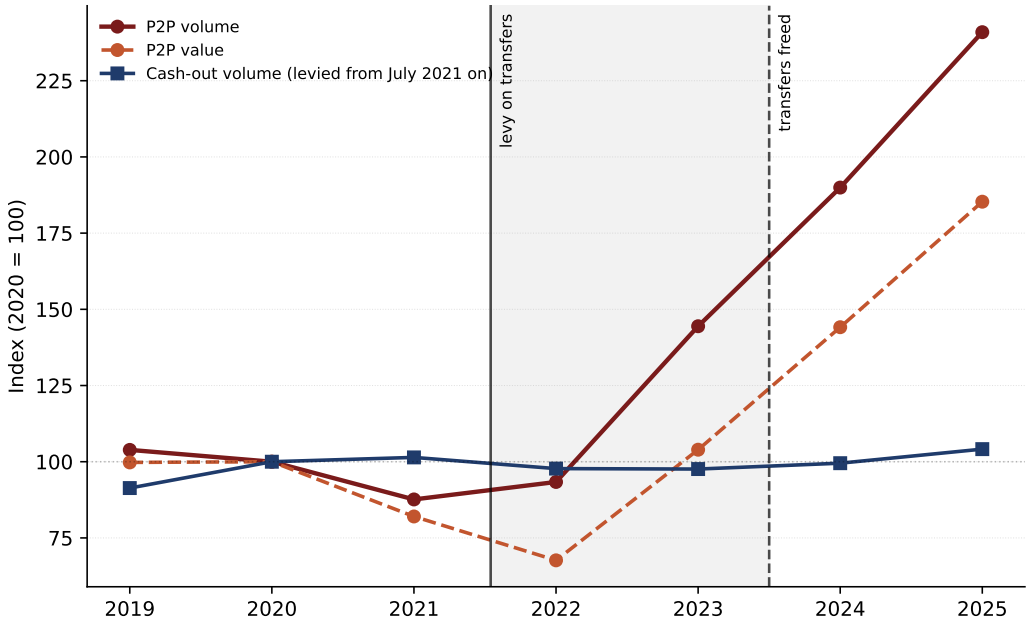
Notes: Senders with at least three person-to-person transactions in June 1 to July 14, 2021 are sorted once into quintiles by their pre-policy mean transfer and then followed (163,916 senders); each series is the cohort's weekly person-to-person count, normalized to 100 on July 8–14. Three providers, June 1 to October 13, non-overlapping seven-day bins. Because the cohort is fixed, a transfer that shrinks stays with its sender rather than moving between bins. Averaged over the thirteen weeks from July 18, quintiles two to five lie between 49.6 and 51.0 and the smallest quintile at 57.4. Account-sampled panel (Appendix A.15); the transaction-level counterpart is Appendix Figure A5.

Figure 9: Money turnover on the mobile money platform, 2021–2022



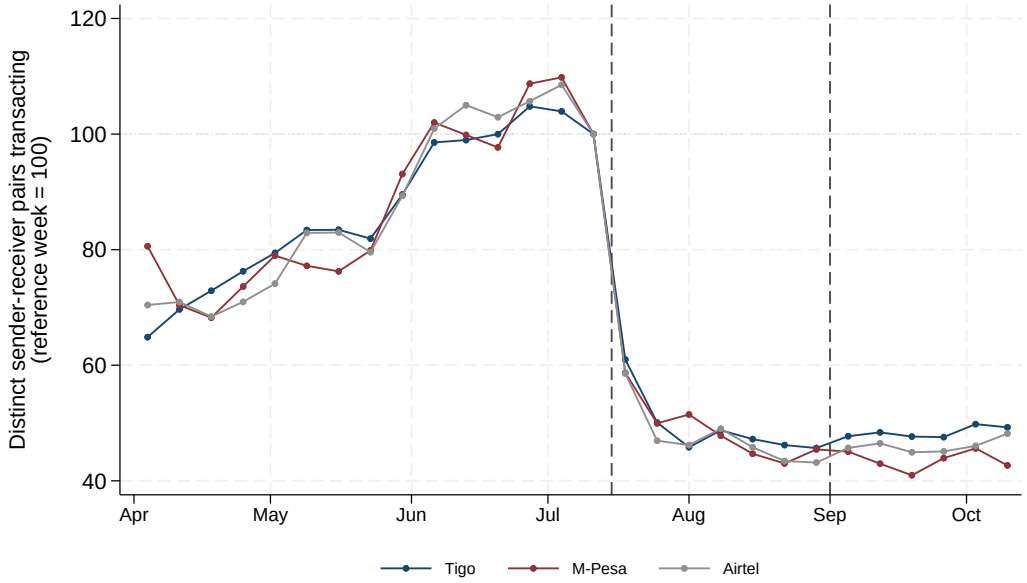
Notes: Bank of Tanzania monthly mobile payments statistics, national aggregates, February 2021 to June 2022; the window ends before the replacement schedule of July 2022. Turnover is the monthly value of mobile payments divided by the trust-account balance, the funds that back the money held in mobile wallets. Monthly points are plotted at mid-month, so July 2021, half of which precedes the levy, sits on the introduction line. Dashed lines: July 15, 2021 introduction and September 1, 2021 cut. Wallet balances are 1,173, 1,189, 1,156, and 1,045 TZS billion in June, July, August, and October 2021, and transaction value 11.7, 10.3, 8.1, and 8.5 trillion TZS. Turnover in every month from August 2021 to June 2022 is below every month from January 2019 to June 2021, including the months before February 2021 that the figure does not plot; over longer horizons the ratio drifts, and it ran between 6.2 and 9.2 in 2015 to 2017. Transaction value in our data is 60 to 63 percent of the national total in every month from May to September 2021. Indexed to June 2021, it stands at 66 in August and 69 in September, against 69 and 69 for the national total and 74 and 68 for the part of the national total our data do not cover.

Figure 10: Bank of Tanzania annual series: person-to-person transfers and cash withdrawals, 2019–2025



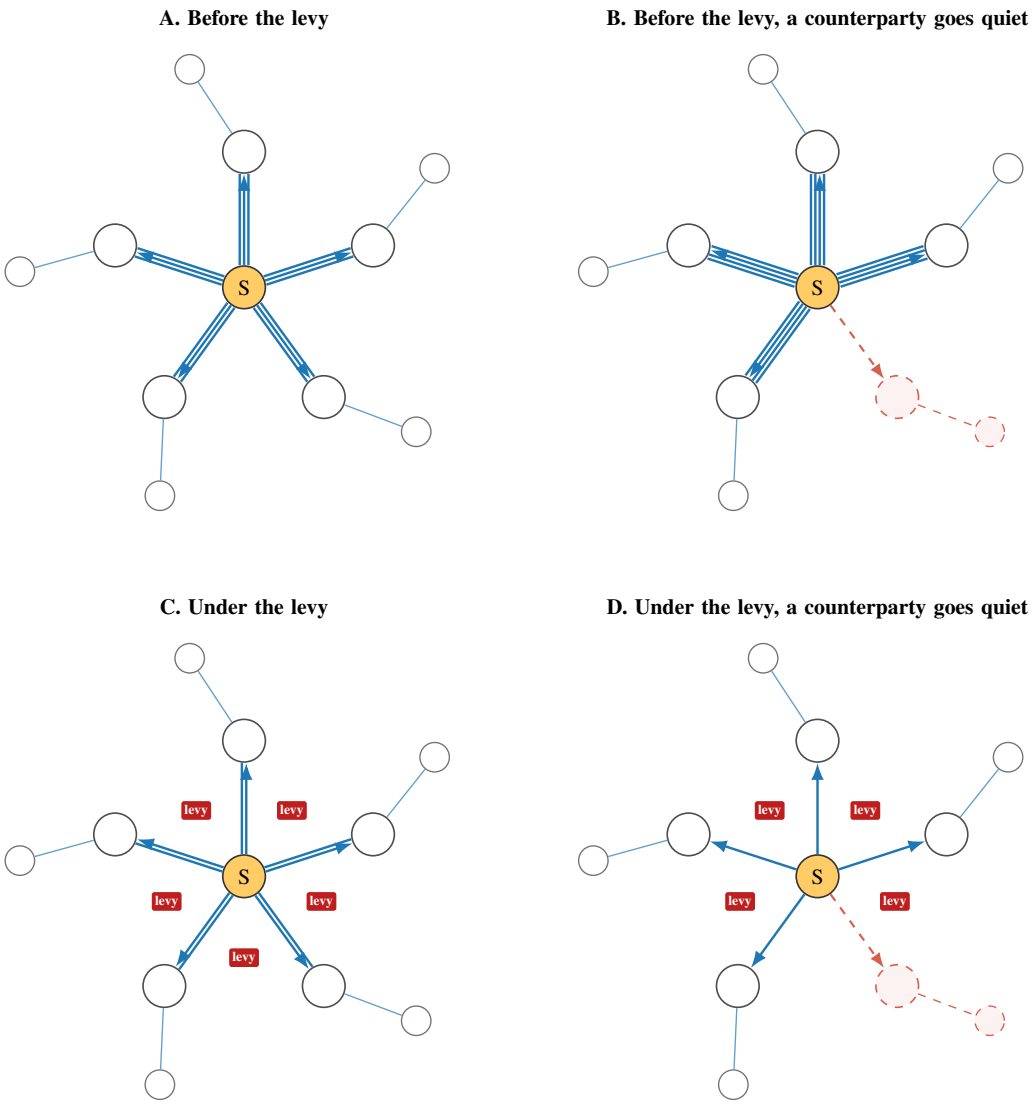
Notes: Annual person-to-person transfer volume and value and cash-withdrawal volume from the Bank of Tanzania's National Payment Systems annual reports, indexed to 2020 = 100. The shaded band marks the period in which transfers were levied (July 15, 2021 to June 30, 2023); withdrawals remain levied throughout. Person-to-person value in 2022, the first full year of the levy, stood at 68 percent of its 2020 level; by 2025 volume had reached 241 and value 185 percent of 2020, while cash-withdrawal volume stayed between 98 and 104 percent of 2020 in every year from 2021 to 2025 (91 percent in 2019).

Figure 11: Distinct sender–receiver pairs transacting, fixed cohorts by provider



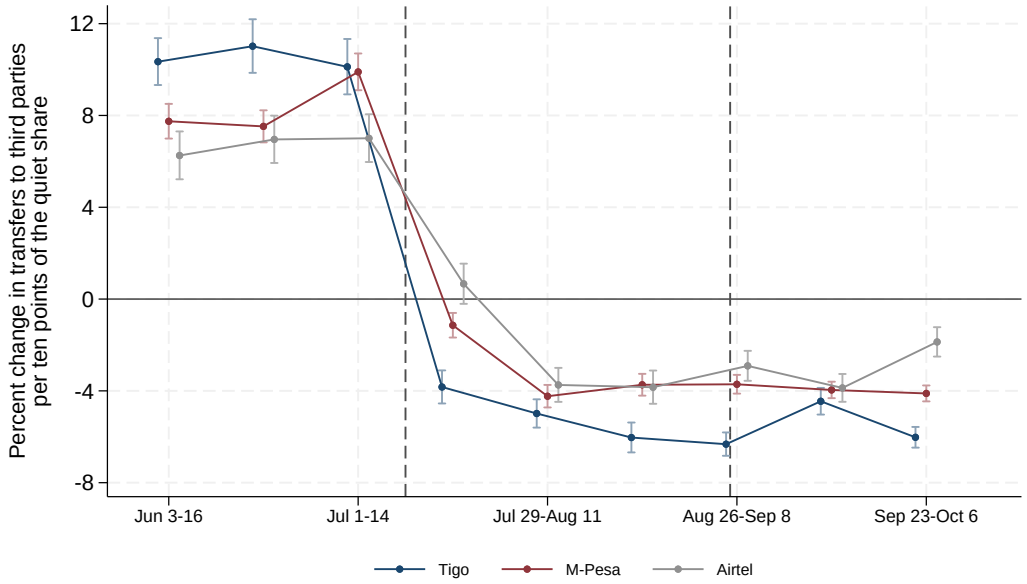
Notes: Distinct on-network person-to-person sender–receiver pairs transacting in each week, on each provider's complete data, for fixed cohorts of senders with at least three on-network person-to-person transfers between June 1 and July 14 (703,301 at Tigo, 1,156,402 at M-Pesa, 616,431 at Airtel), normalized to 100 on July 8–14, 2021. A pair counts once per week however many transfers it carries, so the series measures the extent of the network, not its volume. Inter-operator transfers are excluded at Tigo and M-Pesa; Airtel's transfer file is on-network by construction. Non-overlapping seven-day bins; window ends October 13. Definitions in Appendix A.17.

Figure 12: The rerouting reversal



Notes: Schematic of contagion in Table 4. In each panel S is a sender, the five circles around her are the counterparties she paid regularly before the levy, and the small circles are one further counterparty of each; the number of lines from S to a counterparty represents the volume S sends. A dashed counterparty has gone quiet, sending to nobody for eight weeks. Panels A and B: before the levy, a counterparty going quiet comes with S sending more to the others, the positive pre-levy association in the table. Panels C and D: under the levy every transfer S makes is levied and she sends less to everyone; a counterparty going quiet now comes with S sending less to the others, the negative post-introduction association. Volumes are illustrative.

Figure 13: Contagion fortnight by fortnight



Notes: The specification of Table 4, column one, with the quiet share of a sender's pre-levy counterparties interacted with each fortnight in place of the post-introduction and post-cut interactions. Each point is the percent change in the sender's transfers to receivers outside the pre-levy circle per ten percentage point rise in the quiet share, $100(e^{0.1b}-1)$, in that fortnight; bars are 95 percent intervals, standard errors clustered by sender. Fortnights are anchored on July 15, 2021 and run from June 3 to October 6; dashed lines mark July 15 and September 1. Senders with at least five pre-levy counterparties, complete transaction records of each provider. The pre-levy mean of the three points and the difference between the post-levy and pre-levy means lie within half a point of the corresponding cells of Table 4. A test that the three pre-levy coefficients are equal gives $p = 0.45$ at Tigo and 0.46 at Airtel and rejects at M-Pesa, where the July 1 to 14 coefficient is two points above the two before it.

Tables

Table 1: Summary statistics by transaction type, April–November 2021

Transaction type	Count (millions)	Mean value (TZS)	Median value (TZS)	Share of volume (%)
P2P (person-to-person)	170.25	37,113	9,000	13.14
P2B (person-to-business)	394.07	12,489	1,000	30.42
B2P (business-to-person)	124.87	55,637	5,000	9.64
P2G (person-to-government)	30.32	4,979	2,000	2.34
Cash-in (deposits)	324.76	43,271	9,000	25.07
Cash-out (withdrawals)	251.12	66,873	15,000	19.39
Total	1,295.39	37,969	5,000	100.00

Notes: Sample covers April–November 2021 for Vodacom (M-Pesa), Tigo (Tigo Pesa), and Airtel (Airtel Money). Counts are total transactions over the sample window, net of the Tigo delivery-duplication correction described in Appendix A.1. All columns, including the medians, are computed exactly on the complete transaction records. Share of volume is the share of total transaction count. The two categories B2B (business-to-business) and G2P (government-to-person) contain zero transactions in the available data extract and are omitted from the table. Types are the providers' own labels, with Tigo's inter-operator transfers counted as person-to-person transfers; M-Pesa's inter-operator transfers, which M-Pesa books as merchant payments, stay in that row here, unlike in every later exhibit (Appendix A.2).

Table 2: Difference-in-differences estimates of the response to the 2021 levy

	(1) Extensive margin log(daily count)	(2) Intensive margin log(mean value)	(3) Total value log(daily total)
Panel A. Main estimate: full untaxed basket, corrected assignment			
β_1 (levy introduction, July 15)	−0.255	−0.144	−0.399
	(0.065)	(0.069)	(0.111)
	[0.000]	[0.051]	[0.001]
β_2 (rate cut, Sept. 1)	−0.006	−0.016	−0.022
	(0.019)	(0.022)	(0.029)
Implied elasticity (introduction)			10.33
Series (taxed)	17 (8)	17 (8)	17 (8)
Observations (introduction / rate cut)	1,564 / 1,258	1,564 / 1,258	1,564 / 1,258
Panel B. Upper range: taxed transfers against clean P2B alone			
β_1 (levy introduction, July 15)	−0.336	−0.394	−0.730
	(0.073)	(0.098)	(0.080)
	[0.022]	[0.022]	[0.012]
β_2 (rate cut, Sept. 1)	−0.018	−0.040	−0.058
	(0.013)	(0.055)	(0.067)
Implied elasticity (introduction)			18.91
Series (taxed)	6 (3)	6 (3)	6 (3)
Observations (introduction / rate cut)	552 / 444	552 / 444	552 / 444
Series and provider-by-day FE	Yes	Yes	Yes

Notes: OLS estimates of equation (1) on the complete transaction records, one daily series per provider and transaction type, with series and provider-by-day fixed effects. Panel A: the treated side is person-to-person transfers (on-network and inter-operator), cash-out, and M-Pesa's inter-operator traffic booked under person-to-business payments; the control is every remaining untaxed series, with a levied share of 0.3 percent of rows after the reassignment (Appendix A.2); Tigo's payments to government, which have days with no transactions, are dropped. The seventeen series are, on the taxed side, transfers and cash-out at each provider, Tigo's inter-operator transfers, and M-Pesa's inter-operator traffic; on the untaxed side, merchant payments and business-to-person payments at each provider, cash-in at M-Pesa and Tigo (Airtel's extract has none), and payments to government at Airtel. Panel B: each provider's taxed transfers against its merchant payments alone. Standard errors clustered by series in parentheses; wild cluster bootstrap p-values for β_1 in brackets (Webb weights, 9,999 replications). Panel B gives the largest coefficients because the merchant-payment series' own decline is the smallest among the untaxed types, and we read it as the upper range. β_1 is estimated on June 1 to August 31, 2021 (June 1 to July 14 against July 15 to August 31); β_2 separately on August 1 to October 13 (August against September 1 to October 13), ending before the Airtel duplication artifact (Appendix A.1). Columns: log daily count, log daily mean value, log daily total value. The implied elasticity divides the column (3) coefficient by the log net-of-tax price change at the introduction, −0.0386 for a 40,000 TZS transfer (Section 2.3); none is reported at the cut, where the price change is too small to scale the coefficient (Figure 7).

Table 3: Indirect incidence on business-to-person payments (the salary channel)

	(1) Extensive margin log(daily count)	(2) Intensive margin log(mean value)	(3) Total value log(daily total)
Panel A. B2P vs. clean untaxed basket, M-Pesa and Tigo			
β_1 (levy introduction, July 15)	−0.128	+0.023	−0.105
	(0.094)	(0.067)	(0.151)
β_2 (rate cut, Sept. 1)	+0.062	−0.017	+0.045
	(0.017)	(0.034)	(0.051)
Indirect-incidence elasticity (introduction)			2.72
Series (treated)	6 (2)	6 (2)	6 (2)
Observations (introduction / rate cut)	552 / 444	552 / 444	552 / 444
Panel B. B2P vs. clean P2B, M-Pesa and Tigo			
β_1 (levy introduction, July 15)	−0.182	−0.065	−0.247
	(0.120)	(0.051)	(0.171)
β_2 (rate cut, Sept. 1)	+0.042	−0.049	−0.008
	(0.017)	(0.040)	(0.057)
Indirect-incidence elasticity (introduction)			6.40
Series (treated)	4 (2)	4 (2)	4 (2)
Observations (introduction / rate cut)	368 / 296	368 / 296	368 / 296
Series FE	Yes	Yes	Yes
Provider-by-day FE	Yes	Yes	Yes

Notes: Equation (1) with business-to-person payments as the treated series on the clean-control architecture of Table 2; the taxed series, including M-Pesa's off-net traffic booked under person-to-business payments, are excluded. Business-to-person payments exclude transfers arriving from other networks, which the receiving provider books as payments from its interoperability accounts and the sending provider levies, and Vodacom's payouts from its own accounts (Section 5.4). One daily series per provider and transaction type at M-Pesa and Tigo, which carry 95 percent of the remaining pre-levy business-to-person payments, with series and provider-by-day fixed effects; windows as in Table 2. Panel A: the control is each provider's cash-in and merchant payments (six series, two treated; M-Pesa books payments to government as merchant payments, and Tigo's, which have days with no transactions, are dropped). Panel B: each provider's merchant payments alone (four series); read as the upper range for the reason given in Table 2. Standard errors clustered by series in parentheses; with two treated series they are a guide only, and the placebo-date test of Appendix A.3 locates the July 15 break against every pre-levy date: none of 64 dates (±21-day windows) and one of 78 (±14) produces a count break as large as Panel A's, and 0 of 64 and 48 of 78 a value break as large. Wild cluster bootstrap p-values for β_2 (Webb weights, 9,999 replications): 0.08, 0.59, and 0.35 in Panel A and 0.17, 0.18, and 0.92 in Panel B. Days on which a provider's daily count for a series exceeds three times its within-window median are dropped; none qualifies once Vodacom's payouts, which carry the M-Pesa business-to-person duplication artifact (Appendix A.1), are removed. The elasticity divides the column (3) coefficient by the change in the share kept of a transfer, −0.0386, and is reported at the introduction only.

Table 4: Counterparty dormancy and own third-party activity

	Tigo		M-Pesa		Airtel	
	Transfers	Distinct receivers	Transfers	Distinct receivers	Transfers	Distinct receivers
Quiet share of counterparties	10.46	10.73	8.32	9.10	6.70	6.37
	(0.39)	(0.44)	(0.25)	(0.26)	(0.38)	(0.38)
Share of counterparties not yet on the platform	18.82	20.34	16.58	18.27	13.62	15.87
	(1.07)	(1.12)	(0.70)	(0.72)	(0.70)	(0.72)
Quiet share × post-introduction	−14.63	−14.50	−10.90	−11.30	−8.94	−7.64
	(0.35)	(0.38)	(0.24)	(0.25)	(0.40)	(0.40)
Quiet share × post-cut	0.47	0.37	−0.58	−0.74	−0.02	−0.08
	(0.21)	(0.21)	(0.17)	(0.16)	(0.29)	(0.28)
Observations	1,513,647	1,513,647	2,184,804	2,184,804	772,866	772,866
Senders	168,183	168,183	242,756	242,756	85,874	85,874

Notes: Poisson fixed-effects estimates on a sender-by-fortnight panel, June 3 to October 6, 2021 (the eight-week dormancy lookback is available from June). Cells report the percent change in a sender's on-network person-to-person activity with receivers outside their pre-levy counterparty set per ten-percentage-point rise in the share of those counterparties who are quiet (no transfer to anyone in the preceding eight weeks) or not yet on the platform (no transaction of any type before the fortnight); standard errors in parentheses, clustered by sender and expressed in the same percent units (delta method). Senders have at least five pre-levy counterparties who themselves sent at least once before the levy. Airtel's transfer file carries Airtel number prefixes on 99.99 percent of senders and receivers and is on-network by construction. Fixed effects: sender, fortnight, and pre-levy volume decile by fortnight. Post-introduction covers fortnights from July 15; post-cut covers fortnights entirely after September 1. Definitions in Appendix A.17.

Table 5: Alternatives to the enacted levy

Design	Revenue at enacted terms (share of levy revenue)	Matching the levy's revenue requires	Exposure to behavioral response	Who bears it
Enacted levy on transfers and withdrawals	100	—	Elasticity about 10; network damage on the transfer margin	Senders of small transfers who cannot substitute (Section 5)
Withdrawal-only levy at the enacted rates	68	Withdrawal rates 1.5 times higher before any response; 1.6 times (1.5 to 1.9) once withdrawals respond	Withdrawal elasticity about 10 (interval 0.4 to 19); multiples are lower bounds, since free transfers ease avoidance	Withdrawers; the transfer margin is spared
Flat access charge	—	1,883 TZS per active account per month over 18.6 million accounts; about 200 TZS more for each ten points of exit	Participation: exits of at least 4 percent, falling first on light users	Every active account equally: the median payer (700 to 950 TZS a month now) pays more, the three in ten who pay nothing start paying, the heaviest users pay less
Higher tax on provider fees	About 50 (the existing excise and VAT on fees)	Taxing fees at roughly their own value: a wedge of 63 percent of fees, 67 with the response	Increment of about 0.5 percentage points on a 40,000 TZS transfer and about 3 on a 40,000 TZS withdrawal, against 3.75 for the levy on both; incidence depends on pass-through	Users in proportion to fees paid, mostly withdrawers
VAT enforcement on the transaction data	—	An enforcement gain of 5 percent of VAT net of refunds	None on the platform; presumes a gain we cannot estimate, on data the levy thinned	VAT payers

Notes: Revenue shares are from the high-rate window (July 15 to August 31, 2021); a small negative entry from a provider-side sign convention accounts for the remainder. The withdrawal elasticity is from a withdrawal-specific version of the main design. The access charge matches high-rate revenue over the 18.6 million accounts active in the window; a fifth of active accounts already paid more than 1,883 TZS a month in levy and stayed, and the levy, with a similar average burden per payer, moved participation by 4 percent, a lower bound on exits under a charge that can be escaped only by leaving. The fee tax is the existing 10 percent excise and 18 percent VAT on provider fees, which take about 30 percent of the base fee and raised about half the levy's revenue over the same window; applying the main elasticity raises the required wedge from 63 to 67 percent of fees. The VAT figure is the levy's average annual take over FY 2021/22 to 2023/24 as a share of VAT net of refunds.

Table 6: Marginal value of public funds across Tanzanian tax instruments

Instrument	Rate τ	Base elasticity ε	x	MVPF
Panel A. Mobile money levy				
Introduction rate	0.0375	measured	0.482	1.93
Reduced rate	0.0263	measured	0.451	1.82
Introduction rate, differential response	0.0375	measured	0.329	1.49
Reduced rate, differential response	0.0263	measured	0.370	1.59
Panel B. Comparator instruments				
Value-added tax	0.18	0.5	0.110	1.12
Payroll tax (top bracket)	0.30	0.4	0.171	1.21

Notes: MVPF $= 1/(1-x)$ following Hendren and Sprung-Keyser (2020), with x the share of mechanical revenue lost to the behavioral response. For the levy, x is the fall in log daily total value of the taxed series from their pre-levy level in the main specification over June 1 to October 13, averaged with equal weight over the taxed provider-by-type series (introduction -0.658, standard error 0.086; reduced rate -0.600, standard error 0.089; clustered by series). The differential-response rows use the difference-in-differences coefficient over the same window (introduction -0.399, standard error 0.111; reduced rate -0.462, standard error 0.124), which nets out the untaxed control's own fall within each provider (0.25 log points under the high rate and 0.13 under the reduced rate, averaged with equal weights over the untaxed series). Levy rates are the levy component of the user cost for the reference 40,000 TZS transfer, 3.75 percent at introduction and 2.63 after the cut. For the comparators, $x = \varepsilon\tau/(1-\tau)$ with the rates and elasticities shown; elasticities from Bachas et al. (2024); Brockmeyer et al. (2024) (VAT) and Best et al. (2015); Kleven and Waseem (2013) (payroll tax); rates from the Tanzania Revenue Authority. With the large-reform correction of Finkelstein and Hendren (2020), $(1-x/2)/(1-x)$, the four levy rows read 1.47, 1.41, 1.25, and 1.29 and the comparators 1.06 and 1.10; the text reports the uncorrected figures for comparability with the library.

A Appendix

A.1 Data artifacts

The transaction-level extract exhibits a localized data artifact on a small number of days. On these days the daily count jumps to several multiples of its surrounding level, mean per-transaction value falls in inverse proportion, and daily total value is essentially unchanged. The data provider has confirmed the mechanism: transactions are duplicated on the 28th, 29th, and 30th of affected months, concentrated in end-of-month bulk-payment cycles. The empirical signature is confined to two slices: M-Pesa business-to-person (13 days, April through October 2021) and Airtel merchant payments (19 days, continuous from mid-April to early May 2021). The M-Pesa business-to-person artifact days are April 27, 28, and 29; May 29 and 31; June 24; July 23 and 30; and October 23, 24, 25, 27, and 28, 2021. The Airtel merchant-payment artifact days are April 16 through May 4, 2021. All other combinations (M-Pesa person-to-person transfers, merchant payments, cash-in, and cash-out; Tigo across all types) are clean of this signature. Most trimmed days fall outside the 28th–30th pattern (8 of the 13 M-Pesa days and 16 of the 19 Airtel days) and likely reflect related but less systematic occurrences of the same mechanism. The trimming rule is agnostic to calendar position and treats them identically.

Two further artifacts affect the raw deliveries rather than individual series. First, Tigo's April 2021 delivery duplicates every transaction wholesale: each daily file arrives twice (file-level duplication ratios 2.000 to 2.135), and the April 16 and 17 files arrive four times; the November delivery duplicates the files of November 28–30. Unlike the end-of-month artifact above, these duplicates are exact copies down to the timestamp, so they double count and value together and leave mean value unchanged. We remove them by deduplicating the affected raw files on every field including the timestamp, so that two genuine transactions identical in every field but time are preserved; rows identical down to the same second, about 0.1 percent of rows in unaffected files, cannot be distinguished from delivery duplicates and are collapsed with them. The correction removes 68.8 million rows, all in April and November. No difference-in-differences window contains an affected day; the bunching estimates and the relationship-survival placebo, which use April, are computed on the corrected data. Second, the Airtel extract duplicates transactions from approximately October 16, 2021 onward, inflating daily counts by a factor of 1.7 to 3.2 across every transaction type while leaving mean transaction value unchanged. Because the inflation factor differs between the treated and control series it does not difference out, and it would affect roughly half the days of a window extended to November 30, so every estimation window that would otherwise cross that date ends on October 13, which preserves all three providers.

A.2 Treatment assignment corrections

The main specification (Panel A of Table 2) uses the full set of untaxed series as the control, with two corrections to the treatment assignment that the providers' booking conventions require. Inter-operator person-to-person transfers enter as their own series rather than pooled with on-network person-to-person transfers, and M-Pesa books its inter-operator person-to-person traffic under person-to-business payments

through interoperability accounts; those flows carry the levy on 80 percent of rows, while genuine M-Pesa merchant payments are levied on fewer than 1 percent, so we reassign them to the taxed side. After the reassignment the untaxed basket has a levied share of 0.3 percent of rows. One convention is left in place: the receiving provider books transfers arriving from other networks as business-to-person payments from its interoperability accounts, and these rows, which the sending provider levies, stay in the control; row (7) of Appendix Table A4 drops business-to-person payments from the control. The treated side is person-to-person transfers (on-network and inter-operator) and cash-out plus the reassigned flows; fixed effects are at the series and provider-by-day level, with standard errors clustered by series.

A.3 Inference

Standard errors in Table 2 are clustered by series: seventeen series in Panel A, eight of them taxed, and six in Panel B, three of them taxed. Three features of the setting shape what those standard errors can mean, and we report two additional inference exercises alongside them.

First, the uncertainty here is not sampling uncertainty. We observe essentially the entire population of transactions on the three providers, and the day-level aggregates are computed over millions of transactions each. Decomposing the variance of the daily series, within-day sampling noise accounts for 0.006 percent of the residual variance of log count and 0.3 percent of log mean value: the Poisson bound on log count is 0.0013 against an across-day residual standard deviation of 0.171, and for log mean value the corresponding figures are 0.0048 and 0.089. Neither a count model nor transaction-level clustering would therefore change anything material. The relevant sample is the days of the estimation window, and the relevant dependence runs along time rather than across transaction types.

Second, with so few clusters, and few of them treated, cluster-robust standard errors can over-reject (Conley and Taber, 2011; MacKinnon and Webb, 2017). A wild cluster bootstrap with Webb weights (9,999 replications) gives p-values for the introduction coefficient of 0.0003 on count, 0.05 on mean value and 0.001 on total value in Panel A, and 0.02, 0.02 and 0.01 in Panel B. Pooling the providers into eight series leaves the point estimates unchanged and widens the intervals (Appendix Table A4, row 3).

Third, and more fundamentally, a standard error measures how precisely a level shift is estimated, not whether a shift of that size would have occurred anyway. We address this with in-time placebo randomization inference: the break is re-estimated, with the same series and provider-by-day fixed effects, at every pre-levy date admitting a window of the same length, and the true estimate is located in that distribution. With windows of ± 21 days, 64 placebo dates are available, none of which lies before April 1 or overlaps the levy. For log total value and log count, no placebo produces a break as large as the true one in either panel, so $p < 1/64$. For business-to-person payments (Panel A of Table 3), none of the 64 dates produces a count break as large as July 15's, and one of 78 does with windows of ± 14 days; in value, none of 64 and 48 of 78. The exception is the mean-value channel in the main specification: 17 of 64 pre-levy dates produce a break at least as large as the true one, giving $p = 0.27$ ($p = 0.18$ with ± 14-day windows); against merchant payments it passes ($p = 0.03$). We treat the mean-value result as the least secure of the three. Figure A4 shows the placebo distributions.

A.4 End-of-month cycles and the pre-introduction elevation

In the two weeks immediately before the introduction, the daily total transacted value of every series rises above its pre-policy baseline. For the seven-day bins centered on June 27 and July 4, the normalized total value of person-to-person transfers sits at 109 and 113, business-to-person payments at 133 and 117, and cash-out at 100 and 107, while the reference week of July 8–14 returns each series to its pre-spike level. The most plausible source is a recurring end-of-month cycle in mobile-money activity, amplified at June 30 by the close of Tanzania's fiscal year. The same two-bin structural pattern reappears at the end of October, when no policy change is taking place; computed on the M-Pesa and Tigo pool alone, since the Airtel extract is unusable after mid-October, person-to-person transfers, cash-out, and merchant payments rise by 8 to 15 normalized points into the week containing October 31 and fall back the following week. End-of-month wage disbursements arrive on the platform as business-to-person payments and are then spent through person-to-person transfers, cash withdrawals, and merchant payments, generating co-movement across all types; the fiscal-year-end at June 30 compounds the cycle for the late-June bins specifically, consistent with the spike in business-to-person payments of 33 normalized points at end-June. This pre-policy elevation does not bias our difference-in-differences estimates: the treated-minus-control differential between person-to-person transfers and merchant payments ranges from -10 to $+11$ normalized points across the pre-policy bins, with no systematic drift, so the cyclical pattern that affects both series drops out by construction. In the event study of the main specification (Figure 6), the pre-levy coefficients sit below zero in June and return to zero in the two weeks before collection; once the taxed series are given their own day-of-month profile they are jointly indistinguishable from zero (wild cluster bootstrap $p = 0.16$ on count and 0.86 on total value), and the introduction coefficient on total value moves from -0.40 to -0.38.

A.5 Transaction-size heterogeneity

Appendix Table A18 estimates equation (1) on subsamples defined by pre-policy deciles of transfer size. Transactions in the smallest decile (up to 1,000 TZS) show no response (0.00). The response grows across deciles 2 to 4, reaching -0.61 at decile 4 (roughly 3,000 to 5,500 TZS), holds between -0.50 and -0.62 across the middle deciles, and is largest in the top two, at -0.73 and -0.82, for transfers above roughly 39,000 TZS. Each transaction is classified by its own amount against fixed pre-policy cutoffs, so the gradient mixes fewer transfers with smaller ones: a transfer that shrinks in response to the levy leaves a higher decile and arrives in a lower one, depressing the first and inflating the second. The average transfer sent by the sender cohort of Appendix Table A17 shrinks to 78 percent of its pre-levy size, while the design of Figure 8, which fixes each sender's size group before the levy, shows the number of transfers falling by a similar share in every group. Appendix A.6 makes the same point for the cut.

The bracket schedule gives a second view of the same data. Because the levy is a fixed amount within each of the 23 brackets, the effective rate is hump-shaped in transaction size, peaking at 4.3 percent in the bracket that begins at 20,000 TZS, where the mean transfer is about 22,000 TZS, while the response is close to monotone in size: regressing the bracket-level change in transaction count on the bracket's effective

rate gives a coefficient whose sign depends on the weighting. Log transaction size, with provider fixed effects, accounts for 87 percent of the variation in the bracket-level response, and once size is controlled the effective-rate coefficient is indistinguishable from zero.[30] We do not report an elasticity estimated from cross-bracket variation.

A.6 Transaction-level size deciles across the cut

Table A1 classifies each taxed transfer by its own amount against fixed pre-policy decile cutoffs and compares daily counts and mean values between the pre-policy period (June 1 – July 14) and the post-cut period (September 1 – October 13).[31] Read at face value, the table shows a size gradient in the missing recovery. Small transactions recover most: the count in decile 1 stands at 89 percent of its pre-policy level, and decile 3 at 75 percent. Large transactions recover least: decile 7 stands at 59 percent of its pre-policy level and decile 10 at 46 percent. Appendix Figure A5 shows the same gradient in weekly dynamics by size quintile. The gradient is mechanical. Because each transaction is classified by its own amount, a transfer that shrinks in response to the levy does not stay in its bin and fall: it leaves a high bin and arrives in a low one, depressing counts at the top and inflating them at the bottom. The sender-cohort design of Section 5, which sorts senders once and follows them, removes this, and the gradient largely disappears (Figure 8). What the table does establish is that users are not splitting transfers into smaller pieces to chase lower effective rates: the mean transaction value within each of the bottom nine deciles is essentially flat across the policy episodes, and only the top decile, transactions above roughly 92,000 TZS, shows a within-decile decline, of about 8 percent.

A.7 Mean reversion in the sender cohort

Sorting senders on their activity in a pre-policy window and following them afterward invites mean reversion: a sender selected for being active in the sort window is, on average, less active outside it. Appendix Figure A6 runs the identical sort-and-follow exercise on a window with no rate change, sorting senders on April 1 to May 14, 2021 and following them to July 14. Mean reversion is plainly present and ordered by quintile, but in the opposite direction to the levy result: the largest quintile sits highest. The upward separation of the smallest quintile in Figure 8 therefore survives the placebo and is, if anything, understated by it. The placebo also calibrates the level. Selecting senders on activity in a window mechanically depresses their measured activity outside it, by 11 to 35 points in the placebo, so the post-levy level of about 50 in Figure 8 should not be read as the levy halving activity.

[30] The coefficient on the effective rate is -8.5 (standard error 3.4) unweighted and $+6.3$ (standard error 3.4) weighted by pre-policy volume; controlling for log size it is -0.3 (standard error 0.8). The estimation uses the complete transaction records, bracket-by-provider cells with at least 100 pre-levy transfers (69 cells) with provider fixed effects and standard errors clustered by bracket, and compares June 1 – July 14 with July 15 – August 31, 2021. Bracket membership is itself a function of transaction size, so the exercise mixes within-bracket exit with relocation across bracket boundaries; this is a further reason not to read the coefficient as an elasticity. Separating the two directly requires a bunching design at the bracket thresholds; Appendix A.13 implements one and finds that relocation across boundaries amounts to one to two percent of the missing transactions.

[31] The intermediate high-rate window (July 15 – August 31) is omitted because the question is whether the users who exited at the introduction returned when the rate fell, not the size of the introduction response itself.

Table A1: Daily count and mean value of person-to-person transactions by pre-policy size decile, pre vs. post-cut

	Daily count of P2P transactions			Mean transaction value (TZS)		
Decile	Pre	Sep–Oct	Ratio	Pre	Sep–Oct	Ratio
1	196,437	174,709	0.89	707	701	0.99
2	28,234	24,966	0.88	1,421	1,427	1.00
3	95,838	71,799	0.75	2,360	2,342	0.99
4	103,546	63,792	0.62	4,723	4,707	1.00
5	150,119	94,550	0.63	8,421	8,461	1.00
6	57,138	32,519	0.57	11,414	11,560	1.01
7	105,796	62,656	0.59	17,387	17,192	0.99
8	102,645	58,983	0.57	28,108	28,210	1.00
9	104,973	55,100	0.52	55,986	55,827	1.00
10	104,958	48,654	0.46	306,365	282,239	0.92

Notes: Taxed transfers only, on the clean assignment of Appendix A.2 (person-to-person transfers pooling on-network and inter-operator transfers, including M-Pesa's off-net traffic booked under person-to-business payments). Deciles are defined on the universe distribution of taxed-transfer values during the pre-policy window (June 1 – July 14, 2021) and applied to all transactions. Transfer values bunch at round amounts, so decile boundaries fall on mass points and the deciles hold unequal numbers of transactions. "Pre" columns report daily-count and mean-value averages over June 1 – July 14 (44 days); "Sep–Oct" columns report averages over September 1 – October 13 (43 days), the post-cut window of every other estimate, which ends before the Airtel duplication of mid-October (Appendix A.1). Computed on the complete transaction records, net of the Tigo delivery-duplication correction (Appendix A.1).

A.8 Banked wage recipients

The comparison in Section 5.4 uses the complete M-Pesa data. A wage recipient is an account with at least two payments from businesses in a six-week base window; it is still paid on the platform if it receives any such payment in the following six weeks. Payments from businesses exclude transfers arriving from other networks, which the five interoperability accounts pay, and Vodacom's payouts from its own accounts. The levy period has base window June 1 to July 14 and outcome window July 15 to August 27, 2021 (410,000 recipients); the spring comparison period has base window April 1 to May 14 and outcome window May 15 to June 27 (397,000). The eight documented M-Pesa business-to-person artifact days inside these windows are dropped (Appendix A.1). A recipient is banked if it made at least one transfer from its wallet to a bank account, identified from the provider's registry, between April 1 and May 14, before both outcome windows. The regression of continuation on banked status and its interaction with the levy period absorbs period-specific effects of the disbursement channel, the number of receipts, the mean receipt amount, and the pre-levy total wallet value, with standard errors clustered by channel. The channels are bulk-disbursement accounts rather than individual employers, so the comparison holds the channel fixed, not the employer. The interaction is 2.6 percentage points (standard error 0.35), against a minimum detectable effect of 1.0 point at 80 percent power. Under a broader definition of banked, any wallet-to-bank transfer between April 1 and July 14, it is 0.2 points ($t = 0.6$); under neither definition do the banked leave faster.

A.9 Bank accounts

At M-Pesa the customer registry carries a sector field for business and entity accounts, and 41 of them are classified as banks; a transfer from a wallet to one of these accounts is a wallet-to-bank transfer, so a customer who made one can be marked as banked. The data do not link a wallet to a bank account in any other way. We mark as banked the senders who made at least one wallet-to-bank transfer between April 1 and July 14, 2021. The levy cannot move this definition, and a user with a bank account who did not transfer to it counts as unbanked, which understates any difference between the groups. Of the 1,156,402 senders with at least three person-to-person transfers before the levy, 110,694 (9.6 percent) are banked. Among senders still active after the levy, person-to-person transfers fall by 0.80 log points for the unbanked and 0.79 for the banked. The banked are heavier and wealthier users, moving 3.9 times as much value before the levy. Within cells of pre-levy volume and transfer size they cut 9 log points less, and 4 points less (standard error 0.3) once pre-levy total value is also held fixed; on the same terms they are 2.2 points less likely to stop transacting (standard error 0.1).

Transfers from wallets to bank accounts carry the levy from July 15, 2021; Government Notice 642A of August 31, 2021 made the coverage explicit by defining a transfer to include one from "the user's mobile money account to bank account." Indexed to the June 1 to July 14 weekly mean, wallet-to-bank value averages 65 at the introduction rate and 62 after the cut, against 44 and 45 for on-network person-to-person transfers and 92 and 99 for exempt merchant payments.

A banked user could also leave the platform without moving any money through it, by being paid into the bank and paying from it. Six checks look for such a shift and find at most a small one. First, among on-network pairs active between June 1 and July 14, 13.9 percent of pairs of two unbanked accounts are still active between July 15 and August 31, against 19.3 percent of pairs of two banked accounts. Banked users transact more, however, and once the pair's own pre-levy activity and the sender's volume are held fixed, a pair in which both accounts are banked, the only kind with a bank at both ends, survives exactly as the separate effects of a banked sender and a banked recipient imply (interaction 0.07 percentage points, standard error 0.12). In daily transfer counts, a banked sender's transfers to banked recipients fall 3 log points more than those to unbanked recipients (standard error 3), against 9 points for unbanked senders.

Second, among the same 1.16 million senders, moving from no banked recipients to all banked recipients raises a banked sender's ratio of transfers per day after the levy to before by 0.039 (standard error 0.008), holding deciles of pre-levy volume and transfer size fixed; the effect on stopping altogether is -0.003 (standard error 0.006).

Third, business-to-person payments sent into M-Pesa wallets by the 14 bank accounts in the registry that make them, the route from a bank into a wallet that the levy did not reach until July 2022, fall after the levy: per day, July 15 to August 31 against June 1 to July 14, by 17 percent in number and 29 percent in value. They fall more than other business-to-person payments, which, net of transfers arriving from other networks and Vodacom's payouts (Section 5.4), fall by 10 percent in both, so this route lost ground rather than gained it; these payments also mix transfers made by bank customers with disbursements by the banks themselves, which the data do not separate.

Fourth, a banked user could withdraw cash at an ATM, which the levy did not reach until July 2022, without involving anyone else. Among the 2.19 million accounts with at least three cash-outs before the levy, of which 5.4 percent are banked, the banked cut their mobile cash-outs 3 log points less than others once volume, withdrawal size, and wealth are held fixed (standard error 0.3). By pre-levy withdrawal size, the banked cut 4.5 points more for withdrawals under 16,600 TZS (standard error 0.8), 1.6 points more between 16,600 and 47,700 TZS (standard error 0.5), and 6.6 points less above 47,700 TZS (standard error 0.3), where an untaxed ATM would save the most.

Fifth, a banked user who funds the wallet with cash could deposit it at the bank instead. Splitting users by the share of their pre-levy inflows that came from their own cash deposits, banked users funded mainly by their own cash cut back more than banked users funded mainly by others, by 2.4 log points on transfers, 4.0 on cash-outs, and 1.1 on cash deposits (standard errors 0.6, 0.5, and 0.5), but no more than unbanked users with the same funding mix. Inflows from others include payments from bank accounts into wallets, so the split is approximate.

Sixth, banks and ATMs are densest in Dar es Salaam. Using the district in the M-Pesa registry, available for about 70 percent of these users, the gap between banked and unbanked users is larger in Dar es Salaam than elsewhere, in the direction of banked users cutting back less, by 1.9 log points on transfers and 4.8 on cash-outs (standard errors 0.8 and 0.7); this is the opposite of what better access to banks would predict, although banked users in Dar es Salaam may also differ in kind from those elsewhere.

Two descriptive patterns suggest why banked users stay, without testing it. First, banks were not an everyday channel for spending. From January to June 2021, all bank channels together (mobile banking through bank applications and USSD, internet banking, card payments at points of sale, and ATM withdrawals) carried about 22 million transactions a month, against 275 million mobile money payments, and their typical transaction was five to ten times larger: 324,000 TZS through bank applications, 365,000 TZS by card, and 214,000 TZS at ATMs, against 38,000 TZS on mobile money. The country had 2,037 ATMs and 767,117 mobile money agents. The national instant payment system linking banks and mobile wallets recorded 38 transactions in August 2021 and reached scale only in 2022, and the value moving through banks' own mobile channels over July and August 2021 grew by no more than in the same months of 2019 and 2020.

Second, banked users move money between the bank and the wallet in both directions and spend from the wallet. Between April 1 and July 14, 2021, banked users sent 254 billion TZS from wallets to banks and received 191 billion TZS from banks into wallets, and half of them received at least one payment from a bank. A wallet-to-bank transfer is almost twice as likely as a random day to fall within three days of a business-to-person payment into the wallet (42 against 23 percent), which suggests that part of the income arriving in the wallet is set aside in the bank. After the levy, transfers from wallets to banks, which the levy taxed, fall by half for these users, partly because they were selected for making such transfers before the levy (across all users the fall is about a third), and transfers from banks into wallets, which the levy did not tax, fall by a fifth in number. Payments from banks into wallets mix transfers made by bank customers with disbursements by the banks themselves, and the timing comparison does not hold the user fixed. Read

together, the patterns fit banked users who keep part of their money in the bank but spend it through the wallet, where recipients, merchants, and agents are: the levy taxed the spending, and the bank offered no spending channel of comparable reach. The data do not test this reading.

A.10 Routing transfers through businesses

We find no sign that users avoided the levy by routing transfers through businesses, although for most transfers the route was cheaper than a taxed transfer. A sender could pay a business, which the levy did not tax, and have the business pay the recipient, which it did not tax either. The route pays two provider fees instead of one, but for most transfer sizes the levy exceeded both. At the median fees of June 2021, the route costs less than a taxed transfer at M-Pesa from 10,000 to 199,999 TZS and from 600,000 TZS up, ranges that cover 60 percent of transfers in June, and more below 10,000 TZS and between 200,000 and 599,999; at Tigo, where businesses pay named recipients without a fee, it costs less from 5,000 TZS up, a range that covers 76 percent of transfers. A recipient who wants cash pays the withdrawal levy on either route, so unlike paying a merchant (Section 5) this route is not free of the levy end to end. Only registered organizations with a disbursement account can pay individuals, 417 accounts at M-Pesa and 264 at Tigo between April and August 2021, so the route needs one of them to act as a go-between.

Used on any scale, the route would raise both payments to businesses and payments from businesses to individuals after July 15. Payments from businesses to individuals, net of transfers arriving from other networks (Section 5.4), fall instead: by 15 log points at M-Pesa and 34 at Tigo between June 1 to July 14 and July 15 to August 31, against changes of 0 and 3 log points over the same span in spring. Payments to businesses fall by 5 percent at M-Pesa, leaving out purchases of Vodacom's own products, and by 4 percent at Tigo. They rise after the September cut, by 1 and 27 percent, the recovery of merchant payments documented in Section 5, while payments from businesses stay below their pre-levy level.

Aggregate counts could hide routing on a small scale, so we also look for its signature. A routed transfer is a payment to a business followed, the same day or the next, by a payment of the same amount from a business to an individual. Businesses receive payments and make them under different account numbers, so we cannot follow money through a given business, and we match amounts across each provider's whole platform instead. For each day, we count the pairs formed by a payment to a business and a payment from a business of exactly the same amount, and compare this count with the pairs formed with payments from businesses a week earlier or a week later, which routing does not produce. Same-day pairs exceed this benchmark by less than 0.5 per 1,000 payments to businesses at M-Pesa and by about 2 per 1,000 at Tigo, before and after July 15 alike, and next-day pairs by less or not at all. The excess changes after July 15 by 113 pairs a day at M-Pesa (standard error 621) and 160 at Tigo (standard error 370), and after the September cut by -352 (standard error 598) and 101 (standard error 389). Amounts that are not a multiple of 1,000 TZS, which coincide by chance less often, give the same result. At the upper end of its 95 percent confidence interval, the change after July 15 amounts to about 1 percent of the on-network transfers lost each day, 116,000 at M-Pesa and 81,000 at Tigo. The test cannot detect a business that keeps a commission and pays out less, a payment made more than a day later, or a route that crosses providers.

A.11 Survey corroboration: TNPS Wave 5 pre/post comparison

The transaction-level findings of Sections 5 and 5.4 rest on administrative data alone; the Tanzania National Panel Survey (Wave 5) fieldwork window itself straddled the July 15, 2021 levy introduction, and so allows an independent cross-section check on the direction of the household-level response. Of the 4,709 households in Wave 5, 2,638 were interviewed on or before June 30, 2021 (pre-policy) and 1,621 on or after August 1, 2021 (post-policy); we drop the 450 households interviewed in July to keep the comparison clean, leaving 4,259 households merged with the financial-services module. Mobile money penetration is essentially flat across the policy window: 68.5 percent of pre-policy households report using any mobile money compared with 71.9 percent of post-policy households (a 3.4 percentage point difference, weighted, robust standard error 0.019, $t = 1.80$), and 1.6 points (standard error 1.8) once region and urban status are held fixed. The one exception is the top consumption decile, where mobile money penetration falls from 96.4 percent to 93.1 percent. The flatness in aggregate, together with the fall at the top of the distribution, is consistent with the admin-data finding (Section 5) that the volume response is concentrated within users rather than at the entry/exit margin: a household that sharply cuts its mobile money volume remains classified as a user in the survey.

Bank account ownership rises from 19.8 percent of pre-policy households to 25.0 percent of post-policy households (5.1 percentage points, robust standard error 0.017). Part of that gap is composition: fieldwork moved across regions during 2021, and the post-policy sample is more urban (50 percent against 44), which by itself raises measured bank use. Holding region and urban status fixed, the difference is 4.0 points (standard error 1.7); restricting to regions surveyed in both periods, 4.1; with region-by-urban cells, 3.4 (standard error 1.7). The rise is present in both urban (36.5 to 40.8 percent) and rural (12.4 to 15.8 percent) strata. The Findex secular trend for Tanzanian adult account ownership over 2017–2021 was approximately 1.25 percentage points per year; a three-to-four-point rise within the survey's six-month post-policy window is above that pace, though adult-level Findex and household-level TNPS aggregates are not directly comparable. It does not match the transaction data, in which users with a bank link leave the platform no faster than others (Section 5.4), and we do not read it as evidence that displaced users moved to banks. The cross-sectional design rules out causal interpretation: different households are interviewed in each period, and we cannot rule out secular trends or sampling drift in fieldwork. We read the results as corroboration of direction rather than as independent identification of the levy's effect.

A.12 Why the spillover is bounded rather than identified

Section 7 counts the whole fall in taxed activity as the levy's, on the argument that the 36-log-point decline of the untaxed basket is the levy's spillover, without estimating that spillover on its own. This appendix reports the two designs that tried to estimate it and why both fail in this setting. The failures are informative: they are what a shock common to the whole platform looks like from inside it.

The first design measured a user's exposure by the mechanical levy their pre-levy transactions would attract under the schedule. This exposure has no usable variation. The levy is a deterministic function of

transaction size; transaction sizes are massively bunched at round numbers (81 percent of pre-levy M-Pesa person-to-person transfers are divisible by 1,000 TZS); and the statutory bracket boundaries were legislated at those same round numbers, so the thousand-shilling interval beginning at the 20,000 TZS focal point holds about nine times more transactions than the one just below it. Any exposure measure built from size is therefore a measure of user type. Nor does the September cut supply differential variation: eighteen of the twenty-three levied brackets were cut by exactly 30 percent, three more by between 29.6 and 30.4 percent, one by 31 percent, and one not at all.

The second design measured exposure by pre-levy composition: the share of a user's customer-initiated value in taxed types, measured April 1 to May 26, with weekly untaxed activity as the outcome from June to mid-October, user and week-by-size-decile fixed effects, an unbiased five-percent sample of 497,901 senders at Tigo and Airtel (M-Pesa's merchant category mixes levied off-net transfers and cannot be split before the levy), and a leave-out counterparty version built from the full counterparty rosters of 180,588 network-cohort senders. The point estimates carry the signature of substitution, heavily exposed users cutting taxed activity more and shifting into untaxed types, but the design fails its own discipline. A placebo split inside the levy-free weeks yields half the main coefficient (0.22 against 0.50, both with t-statistics above 15), the event study shows the differential marching steadily upward through the entire pre-period, and the counterparty version repeats the pattern. Users measured with extreme taxed shares revert toward the mean mechanically, whether or not a levy arrives: composition-based exposure is confounded with composition dynamics by construction, and we did not search for a specification that hides this.

The reading that fits is the network one of Section 6. The levy's spillover onto untaxed activity travels through the platform: every user's active-counterparty share falls, and salary chains and agent economics deteriorate for everyone at once. A shock that is common within the platform has no within-platform control group. What the data support is the discontinuity at the legislated date and the absence of any shock in the payment channels off the platform. Summed across its series, the untaxed basket falls by 36 log points in value and 16 in count over the window of the main estimate (June 1 to August 31, 2021); with its series weighted equally, as in the difference-in-differences, by 25 and 17. It arrives discontinuously on July 15 itself (Figure 4), which a seasonal or macroeconomic explanation would have to reproduce by coincidence. It has no counterpart in the payment channels off the platform: in the Bank of Tanzania's series for July and August 2021 against the same months of 2020, ATM value runs 78 percent higher, point-of-sale 95 percent, internet banking 41 percent, electronic transfers 11 percent, and cheques 8 percent. And the substitution the levy pushed toward untaxed types runs the other way and masks part of the decline, so 36 log points understates the gross spillover. Only seasonality specific to mobile money, which no pre-2021 transaction data exist to exclude, could pull the whole-fall figure of Section 7 toward its lower bound.

A.13 Bunching at the bracket notches

The schedule of Table A12 is notched: a transfer of 19,999 TZS carries a levy of 610 TZS and a transfer of 20,000 carries 960. Notches of this kind ordinarily invite a bunching estimator in the tradition of Kleven and Waseem (2013), with excess mass accumulating just below each boundary and missing mass just above

it. Two features of this distribution defeat the standard polynomial counterfactual. Transaction amounts are concentrated on round numbers before any tax exists: 81 percent of pre-levy M-Pesa person-to-person transfers are divisible by 1,000 TZS, and the single 100-TZS cell at exactly 10,000 holds 6 percent of all pre-levy person-to-person volume. And the boundaries were legislated at those same round numbers, so every notch sits on a focal amount whose mass no smooth function can reproduce.

We therefore difference the bunching against the pre-levy distribution itself. For each of the twenty-two boundaries we form histograms of the person-to-person line, pooled across the three providers, in a window of twenty percent on either side of the boundary, for the pre-introduction period (April 1 to July 14) and the post-introduction period (July 15 to August 31, 2021). The counterfactual for each post-period bin is its own pre-period count, scaled by a smooth reshaping function that absorbs the overall contraction and the shift of the distribution toward smaller amounts; we estimate it by Poisson regression of post-period counts on a fourth-degree polynomial in distance to the notch (second- and third-degree at the 2,000 and 3,000 TZS boundaries, which have few bins to fit), with the log pre-period count as an offset, using only bins outside every boundary's excluded region. Round-number mass appears identically in the observed and counterfactual histograms and cancels bin by bin, and the predicted counts are positive by construction. We report the excess mass within five percent below each boundary and, separately, at and above it, both normalized by the counterfactual mass of the region. Standard errors come from resampling days with replacement within each period. As a placebo, we apply the identical estimator to two levy-free subperiods, April 1 to May 22 against May 23 to July 14.

Table A2 reports three results. First, bunching is present, and precisely estimated, at the boundaries where the levy jump is large relative to the amounts involved: normalized excess mass below the boundary is 0.40 (standard error 0.02) at 20,000 TZS, 0.52 (0.02) at 30,000, 2.01 (0.08) at 40,000, 1.09 (0.03) at 50,000, and 1.63 (0.04) at 5,000, against placebo values of at most 0.07 in absolute value at each of these boundaries. Second, the relocated mass is trivial against the missing mass. The excess below summed over all twenty-two boundaries is approximately 178,000 transactions over the 48 post-introduction days, or 310,000 over the nine boundaries where the estimate is positive and statistically significant, against roughly 18 million person-to-person transactions missing from the same window at the pre-period daily rate: relocation across boundaries accounts for one to two percent of the response. Third, there is no missing mass at and above the boundaries where bunching is sharpest, because the focal amounts themselves are the stickiest part of the distribution. The daily volume of transfers of exactly 20,000 fell to 57 percent of its pre-levy rate, other round amounts in the surrounding window fell to 60 percent, and non-round amounts fell to 36 percent, so the estimator finds excess rather than missing mass at and just above these boundaries. The excess below 20,000 is equivalent to about two percent of the pre-levy volume at exactly 20,000, although moving one shilling lower would have saved 350 TZS.

These estimates are upper bounds on notch avoidance, because the excess below a boundary mixes genuine relocation with the greater persistence of round amounts just documented. Genuine relocation is nonetheless visible at the bin level: arrivals concentrate at amounts that carried little mass before the levy, and the 100-TZS cell holding 19,900 to 19,999 stands at nearly two and a half times its pre-levy share

of person-to-person volume. Users overwhelmingly either pay the levy at the amount they always sent or leave the taxed base; almost none re-optimize the amount to a cheaper bracket. This is the boundary-level counterpart of the finding in Section 5 that transaction size, not the effective rate, explains the bracket-level response, and it is the direct evidence behind treating the notch structure as second order in the welfare calculation of Section 7. It also settles a design question: a schedule with kinks in place of notches, a levy proportional to the amount, would have removed the incentive to relocate across a boundary, but relocation is one to two percent of the response, so the shape of the schedule is not what drives the behavioral cost. Users respond to being taxed, not to where the steps fall, and a proportional levy of the same yield would have met the same response.

Table A2: Difference-in-bunching at the bracket notches

		Excess mass below boundary				At/above boundary		Placebo, below	
Notch	Δ levy	Transactions		Normalized		Normalized		Normalized	
2,000	6	−11,342	(862)	−0.144	(0.009)	−0.128	(0.019)	−0.048	(0.014)
3,000	11	2,202	(950)	0.028	(0.012)	−0.115	(0.007)	0.053	(0.011)
4,000	29	−1,034	(864)	−0.021	(0.017)	−0.211	(0.011)	−0.047	(0.013)
5,000	44	123,066	(1,944)	1.634	(0.039)	1.992	(0.039)	0.008	(0.013)
7,000	25	−47,762	(1,225)	−0.470	(0.006)	−0.373	(0.006)	0.019	(0.010)
10,000	195	22,479	(6,357)	0.054	(0.016)	−0.204	(0.007)	−0.002	(0.014)
15,000	290	−24,367	(1,015)	−0.313	(0.008)	0.167	(0.013)	−0.049	(0.013)
20,000	350	38,592	(1,327)	0.396	(0.022)	0.651	(0.029)	−0.034	(0.013)
30,000	140	21,607	(493)	0.522	(0.023)	0.622	(0.022)	−0.038	(0.009)
40,000	400	51,503	(648)	2.010	(0.075)	1.485	(0.058)	−0.007	(0.011)
50,000	550	44,252	(642)	1.094	(0.031)	0.919	(0.025)	−0.063	(0.012)
100,000	480	5,935	(551)	0.086	(0.009)	−0.358	(0.006)	−0.018	(0.009)
200,000	410	−9,354	(509)	−0.219	(0.011)	−0.403	(0.009)	−0.016	(0.010)
300,000	560	146	(239)	0.008	(0.012)	−0.116	(0.013)	−0.003	(0.012)
400,000	600	−804	(206)	−0.053	(0.013)	−0.176	(0.010)	0.015	(0.016)
500,000	1,100	−4,226	(340)	−0.248	(0.015)	−0.211	(0.015)	0.015	(0.025)
600,000	1,200	58	(330)	0.007	(0.040)	−0.113	(0.039)	0.059	(0.032)
700,000	700	−793	(215)	−0.109	(0.026)	−0.284	(0.027)	0.013	(0.026)
800,000	420	476	(123)	0.098	(0.028)	−0.002	(0.023)	−0.031	(0.023)
900,000	1,380	−1,051	(797)	−0.199	(0.113)	−0.306	(0.159)	−0.063	(0.120)
1,000,000	500	−30,291	(13,102)	−0.532	(0.090)	0.107	(0.122)	−0.196	(0.113)
3,000,000	600	−953	(210)	−0.251	(0.044)	0.027	(0.062)	−0.005	(0.058)

Notes: Difference-in-bunching estimates at the 22 interior boundaries of the schedule (Table A12), person-to-person line pooled across the three providers. Pre period April 1 to July 14, post period July 15 to August 31, 2021; the placebo compares April 1 to May 22 against May 23 to July 14. Δ levy is the statutory increase in the levy at the boundary under the July schedule, in TZS. Excess mass is observed minus counterfactual transactions in the bins within five percent below the boundary; the normalized columns divide the excess below, and at and above, the boundary by the counterfactual mass of that region. The counterfactual rescales the pre-period histogram by a smooth function estimated by Poisson regression of post-period bin counts on a fourth-degree polynomial in distance to the notch (second and third degree at the 2,000 and 3,000 boundaries) with the log pre-period count as offset, fitted outside every boundary's excluded region within a twenty-percent window. Bins are 100 TZS up to 50,000, 1,000 TZS from 100,000 to 900,000, and 10,000 TZS at 1 and 3 million; the 1 and 3 million boundaries are inclusive, so their threshold bin belongs to the region below; the 2,000 estimate rests on six fitted bins. Standard errors from resampling days within each period, 200 replications.

A.14 Lower-bound decile incidence from TNPS remittance flows

This subsection computes a lower-bound estimate of the levy's first-order incidence by consumption decile, using TNPS Wave 5 remittance inflows as a calibration target for s_h, the share of household consumption flowing through mobile money. Under a first-order welfare approximation, the burden of the levy as a fraction of consumption is $\Delta V_h/C_h \approx -\tau \cdot s_h$, where s_h is averaged over all households in decile h, households that do not use mobile money counting as zero, so that mobile-money penetration α_h (Figure A8) is already inside it; τ is the effective rate at the reference 40,000 TZS transaction (Table A13; 3.75 percent at introduction, 2.63 percent after the September cut), and s_h is estimated here.

The TNPS remittance module (`hh_sec_q2`) reports, for each source, the TSH amount received in cash over the past twelve months and up to three channels through which transfers arrived. Channels include four mobile money codes (M-Pesa, Tigo Pesa, Airtel Money, and EzyPesa), a bank-account code, friends/relatives (cash hand-delivery), and other. We compute s_h as the household sum of mobile-money-channel inflows divided by household annual consumption (the LSMS consumption aggregate per adult equivalent, multiplied by the adult-equivalent count and annualized). Because the amount variable is reported at the source level while channel is reported as up to three codes per source, we cannot apportion a single amount across multiple listed channels, and instead report two attribution bounds: Option A attributes the amount only to the primary channel; Option C attributes the amount whenever any of the three listed channels is mobile money. Estimates are weighted using TNPS sampling weights and winsorized at the 99th percentile to prevent a small number of low-consumption, high-remittance households from dominating the means.

Table A3: Lower-bound first-order incidence of the levy by consumption decile (TNPS Wave 5)

Consumption decile	α_h	s_h (A–C, %)	Burden at intro rate (A–C, %)
1	0.42	0.45–0.51	0.016–0.018
2	0.50	0.75–0.89	0.027–0.033
3	0.57	0.40–0.42	0.015–0.016
4	0.57	0.50–0.61	0.019–0.023
5	0.67	0.54–0.67	0.020–0.025
6	0.76	0.24–0.27	0.009–0.010
7	0.76	0.58–0.66	0.022–0.025
8	0.81	0.82–0.84	0.031–0.032
9	0.83	0.61–0.62	0.023
10	0.96	1.27–1.42	0.048–0.053

Notes: α_h is the share of decile-h households using any mobile money (TNPS Wave 5, pre-policy). s_h is the share of household annual consumption flowing through mobile-money remittance inflows, calibrated from `hh_sec_q2`. A–C reports the lower (Option A: primary channel only) and upper (Option C: any of three listed channels) attribution bound for the channel–amount allocation. Burden is the first-order incidence $100 \times \tau_{\text{intro}} \cdot s_h$ at $\tau_{\text{intro}} = 0.0375$, the levy on a 40,000 TZS transfer, with s_h taken over households using mobile money and the others counted as zero; α_h is reported for reference. All means weighted by TNPS sampling weights and winsorized at the 99th percentile.

The estimated s_h rises with consumption: 0.5 percent in the bottom decile, climbing to 1.3–1.4 percent

in the top decile. The implied first-order burden at the introduction rate ranges from about 0.017 percent of consumption in the bottom decile to 0.05 percent in the top, an approximately three-fold gradient. Per household, this corresponds to roughly 300 TZS per year in the bottom decile and 3,400 to 4,100 TZS in the top. Aggregated across roughly 13 million Tanzanian households at the mean burden of about 1,000 to 1,200 TZS, the implied annual burden from the remittance-channel slice of mobile money is about 14 to 16 billion TZS. Statutory revenue from the levy was approximately 288 billion TZS per year, so remittance flows account for roughly 5 percent of the incidence implied by the formula; the remaining 95 percent flows through own-funded mobile-money transactions (person-to-person sends, bill payments, cash-outs) that TNPS does not measure at the channel level.

Two qualifications matter for interpretation. First, this is the lower-bound *statutory* incidence in the absence of behavioral response: the top-decile burden is what households would pay if they kept their pre-policy transaction pattern at the new rate. Section 5 finds that senders of every size cut back by a similar share, which suggests that the behavioral response lowers the burden in similar proportion at every decile rather than shifting it across the distribution; the survey cannot be linked to the transaction data, so we cannot check this at the household level. Second, the calculation holds τ constant across deciles. Linking deciles to typical transaction sizes would lower τ for the top decile (whose typical transactions face approximately 1 percent rather than the 3.75 percent reference) and would flatten the rise in the burden with consumption shown in the table.

Redistribution lowers the bar of Section 7.2, the spending MVPF the levy's revenue has to clear, but by an amount our data cannot size. The usual argument for tolerating a distortionary tax discounts its welfare cost by the weight placed on beneficiaries relative to payers, and presumes the payers are the better-off party. Here they partly are. Senders of every size cut back by a similar share (Section 5), so the levy is paid across the whole range of users, and mobile money use rises steeply with household consumption; the survey-based lower bound puts the levy's burden at about 0.02 percent of consumption in the poorest tenth of households and 0.05 percent in the richest (above). The spending the levy financed targeted poor rural households, so the discount is positive. How large it is depends on the welfare weights and on the part of the burden the survey does not measure, about 95 percent of it.

A.15 Specification, provider, and sample robustness

Table A4 checks the main introduction estimate against the choices behind it. Row (1) is the main estimate: one series per provider and transaction type, compared within provider and day. Replacing the provider-by-day effects with calendar-day effects (row 2) or pooling the series across providers (row 3) moves the total-value coefficient by less than 0.01; pooling widens the standard error from 0.111 to 0.136. Leaving M-Pesa's levied off-net traffic in the control, where the raw type labels put it (row 4), gives -0.391 (0.123); moving it to the taxed side, where the levy incidence puts it, raises the coefficient by 0.008.

Rows (5) and (6) re-estimate the main specification on Tigo alone and without Airtel. The total-value coefficients are -0.394 (standard error 0.221) and -0.344 (0.123), elasticities 10.2 and 8.9, within a standard error of the main figure; the precision falls because each variant has fewer series. Row (7) drops business-to-

person payments from the untaxed side, since part of them are transfers arriving from other networks, which the sending network levies (Section 5.4). The total-value coefficient becomes −0.449 (0.119), an elasticity of 11.6, so the main estimate, which keeps them in the control, is the smaller of the two.

Panel B checks the account-sampled panel on which the cohort figure rests (Figure 8). The panel keeps every transaction of the senders whose account identifier ends in one of five percent of the possible two-digit endings. For each provider and type it compares the pre-levy to high-rate log change in daily value computed on the sampled senders with the same change computed on everyone else. For the series the paper's estimates use, the two agree within 0.04 log points: person-to-person transfers −0.95 against −0.92 at Tigo, −0.80 against −0.82 at M-Pesa, and −0.94 against −0.93 at Airtel; cash-out within 0.02 at all three providers; cash-in within 0.01. Business-to-person payments and payments to government are omitted because their senders are businesses and government agencies, which a sample drawn on senders barely contains; no result in the paper uses the sampled panel for those types.

Table A4: Specification, provider, and sample robustness

Panel A. Introduction coefficient β_1 by specification (provider-by-type series unless stated)	Log count	Log mean value	Log total value	Elasticity
(1) Provider-by-day fixed effects (main)	−0.255	−0.144	−0.399	10.3
	(0.065)	(0.069)	(0.111)	
(2) Calendar-day fixed effects	−0.255	−0.152	−0.407	10.5
	(0.061)	(0.067)	(0.102)	
(3) Series pooled across providers	−0.266	−0.135	−0.401	10.4
	(0.086)	(0.080)	(0.136)	
(4) M-Pesa off-net traffic left in the control	−0.268	−0.123	−0.391	10.1
	(0.068)	(0.078)	(0.123)	
(5) Tigo only	−0.273	−0.121	−0.394	10.2
	(0.116)	(0.126)	(0.221)	
(6) Without Airtel	−0.250	−0.094	−0.344	8.9
	(0.071)	(0.070)	(0.123)	
(7) Without business-to-person payments	−0.297	−0.152	−0.449	11.6
	(0.072)	(0.084)	(0.119)	

Panel B. Pre-levy to high-rate log change in daily value, sampled senders vs. all others	P2P	Cash-out	Cash-in	P2B
Tigo, sampled	−0.954	−0.339	−0.379	−0.049
Tigo, others	−0.921	−0.356	−0.378	−0.013
M-Pesa, sampled	−0.798	−0.411	−0.389	−0.403
M-Pesa, others	−0.818	−0.432	−0.394	−0.374
Airtel, sampled	−0.936	−0.391	—	−0.135
Airtel, others	−0.934	−0.393	—	−0.154

Notes: Panel A: equation (1) on the complete data, one series per provider and transaction type with series and provider-by-day fixed effects unless stated, June 1 to August 31, 2021, standard errors clustered by series in parentheses; the elasticity divides the log total value coefficient by the log net-of-tax price change of −0.0386. Row (1) is Panel A of Table 2; row (3) pools each transaction type across providers, with calendar-day fixed effects. Panel B: log change in mean daily value between June 1 to July 14 and July 15 to August 31, 2021, by provider and type, for senders in the account-sampled panel and for all other senders; the Airtel extract has no cash-in series.

A.16 The conventional benchmark

The conventional model of Section 5.3 treats each transfer on its own. Formally, a transfer of value v carries a surplus s over the best alternative, drawn from a distribution $F(s \mid v)$; the schedule assigns a levy $L(v)$; the transfer is made if $s \geq L(v)$; and daily volume is proportional to

$$Q(L) \;=\; \mathbb{E}_v[\,1 - F(L(v) \mid v)\,].$$

Two properties follow. First, Q depends on the schedule in force on the day of the transfer and on nothing else, so a schedule that is weakly lower at every value must raise volume. Second, adjustment frictions that delay the return do not change this: any specification in which F is fixed predicts monotone convergence toward the higher post-cut volume.

A.17 How the network is measured

This appendix defines each object in Section 6 and states how it is computed. Every measure is computed on the complete transaction records of each provider; none uses the sampled panel.

The graph. A node is an account. For a window of days w, the edge set $E(w)$ is the set of ordered sender-receiver pairs (i, j) with at least one on-network person-to-person transfer from i to j in w. A pair counts once per window however many transfers it carries, and (i, j) and (j, i) are distinct pairs. On-network means that both accounts are with the same provider. Tigo's records book transfers to other operators as a separate type, which we exclude; M-Pesa's person-to-person type is on-network by construction, its transfers to other operators being booked under merchant payments; Airtel's transfer file carries Airtel number prefixes on 99.99 percent of senders and receivers and is on-network by construction. Rows whose sender or receiver identifier is the placeholder zero are dropped, since the placeholder would otherwise appear as one account with millions of counterparties. The records are collapsed to one row per sender, receiver, and day, carrying the number and value of transfers, and every measure below is computed from that file.

Windows. The pre-levy window is June 1 to July 14, 2021 (44 days), the high-rate window July 15 to August 31 (48 days), and the post-cut window September 1 to October 13 (43 days). The placebo shifts the first two back to a stretch with no rate change: a base window of April 1 to May 14 (44 days) and a follow-up window of May 15 to July 1 (48 days). Weekly series use non-overlapping seven-day bins anchored on the reference week of July 8 to 14 and run from April 1 to October 13.

Cohorts. The descriptive measures are computed for a fixed cohort of senders with at least three person-to-person transfers between June 1 and July 14: 703,301 at Tigo, 1,156,402 at M-Pesa, and 616,431 at Airtel. The placebo applies the same rule to its own base window. Fixing the cohort keeps a sender who stops transacting in the denominator rather than letting them leave the sample. The contagion design uses a narrower cohort, senders with at least five counterparties, defined below.

Pairs transacting (Figure 11). For each seven-day bin b, the series is $|E(b)|$ over the cohort's pairs, indexed to 100 on the reference week. A bin with fewer than six days of data is dropped.

Survival (Appendix Table A15). With E^{pre} the cohort's pairs between June 1 and July 14 and E^{high} its pairs between July 15 and August 31, survival is

$$s = \frac{|E^{\text{pre}} \cap E^{\text{high}}|}{|E^{\text{pre}}|},$$

the share of pre-levy pairs with at least one transfer between July 15 and August 31. The placebo survival s^{pl} is the same ratio for the placebo cohort, base window, and follow-up window, and the relative decline in the table is $s/s^{\text{pl}} - 1$. The table also reports the share of pre-levy pairs with no transfer through October 13. That column has no placebo counterpart, since a follow-up of the same length from May 14 would cross July 15, and the text does not use it.

Counterparties per sender (Appendix Table A5). For each cohort sender and window, the number of distinct receivers in the sender's pairs in that window. The mean and quantiles are taken over the cohort senders with at least one transfer in the window, so a sender with none, whose count is zero, is left out of them and shows up instead in the share of the cohort active in the window. The July 15 to August 31 window is four days longer than the June 1 to July 14 window, which works against the fall in the mean. The share of the cohort with no transaction of any type between July 15 and August 31, in either role and across every transaction type in the account-by-window file, is 5.2 percent at Tigo, 6.4 at M-Pesa, and 9.3 at Airtel, 6.7 percent pooled; among cohort senders with no person-to-person transfer in that window, 78 percent have some other transaction in it (82, 80, and 72 percent by provider).

New pairs (Appendix Figure A9). A pair is new in the bin of its first transfer on or after April 1, the first day of the data. The series counts the cohort's pairs that are new in each bin, indexed to 100 on June 10 to 16. A pair active in March appears as new when it first transacts in April, so the early bins overstate formation and the figure starts in May.

Receiver degree and retention (Appendix Table A6). For each receiver j with at least one sender in the base window, the in-degree d_j is the number of distinct senders in j's pairs in that window, and retention is j's in-degree in the follow-up window divided by d_j. Receivers are grouped by d_j into one, two to four, five to nine, ten to 49, and 50 or more senders, and the table reports mean retention by group for the levy windows (June 1 to July 14 followed by July 15 to August 31) and for the placebo windows (April 1 to May 14 followed by May 15 to July 1). This measure uses all senders rather than the cohort, so that a receiver's degree counts every sender. The ratio of the two retentions is the levy's toll net of the tilt that busy receivers show in any window, and it is this ratio that the text describes as close to proportional across groups.

Deposits by inbound degree (Appendix Table A7). For each account with any cash-in, the inbound degree is the number of distinct senders from whom it received an on-network person-to-person transfer between June 1 and July 14, zero for an account that received none. Accounts are grouped as zero, one, two to four, five to nine, ten to 49, and 50 or more senders, and the table reports the change in the log of each group's total cash-in value from the June 1 to July 14 window to the July 15 to August 31 window, together with the group's share of the total fall. Airtel's extract carries no cash-in records. The two windows differ in length by four days, which shifts every group's change by the same amount and leaves the comparison across groups unaffected.

Counterparty circles and contagion (Table 4). Sender i's circle R_i is the set of receivers with at least one on-network transfer from i between June 1 and July 14 who themselves sent at least one transfer before July 15, so that their later silence is observable. The design keeps senders with $|R_i| \geq 5$ and follows them over fortnights $k = -3, \ldots, 5$ anchored on July 15, from June 3 to October 6; earlier fortnights are unusable because the eight-week lookback would leave the data, which begin on April 1. A counterparty j is dormant at k if j sent to nobody in fortnights $k-4$ to $k-1$. Dormancy is a state rather than an exit, and a later transfer ends it. Before the levy, a dormant counterparty may instead be one not yet on the platform: dormant, with no transaction of any type before k and no on-network transfer received before k, and with its first transfer in or after k. After July 15 no counterparty meets that definition. The quiet share S^q_{ik} is the share of R_i dormant and not new at k, and S^n_{ik} the share not yet on the platform. The outcome y_{ik} is the number of transfers from i in fortnight k to receivers outside R_i, or the number of such receivers, with zeros kept; it cannot mechanically contain the counterparties whose silence defines the treatment. We estimate the Poisson regression

$$\mathbb{E}[y_{ik}] \;=\; \exp\Big(\alpha_i + \lambda_k + \gamma_{d(i)k} + \beta\, S^q_{ik} + \beta^n S^n_{ik} + \delta\, S^q_{ik}\,\mathbf{1}[k \geq 0] + \delta^c\, S^q_{ik}\,\mathbf{1}[k \geq 4]\Big),$$

with sender effects α_i, fortnight effects λ_k, and effects $\gamma_{d(i)k}$ for the sender's pre-levy volume decile by fortnight, the decile being of i's transfers to R_i between June 1 and July 14; standard errors are clustered by sender. Fortnights $k \geq 4$ begin on September 9 and lie entirely after the cut. Table 4 reports $100\,(e^{0.1b} - 1)$ for each coefficient b, the percent change in the outcome per ten percentage points of the share: β is contagion before the levy, δ its change at the introduction, and δ^c its further change at the cut. Figure 13 replaces the two interactions with the quiet share interacted with each fortnight and plots the nine coefficients. The bounds on the share of the decline that spread, and the looser definitions behind them, are in Appendix A.18.

Dead and feasible pairs. From the same circles, a pair is dead if it has no transfer from July 15 to August 3 and feasible if both of its accounts send to third parties in that window, so that nothing but the relationship itself stands in the way of a transfer. Dead pairs are at risk from August 4. For each week from four weeks before September 1 to four weeks after, the revival rate is the share of pairs still dead at the start of the week that transact in it. Pairs whose counterparty sends to nobody in the window are the comparison group.

Reach, support ratio, and the largest connected part (Appendix Table A19). For each window the graph is made undirected: two accounts are linked if either sent the other at least one on-network transfer in the window. For each cohort sender, one-step reach is the number of accounts linked to her and two-step reach the number linked to her or to one of those accounts, herself excluded; a sender with no link counts zero. The support ratio of a sender with at least two links is the share of the accounts linked to her that are linked to at least one other of them, the support ratio of Chetty et al. (2022) computed around each sender. The largest connected part is the largest set of accounts any two of which are joined by a chain of links, and the table reports its share of the accounts with at least one link in the window. The placebo applies the same measures to the placebo cohort in the windows April 1 to May 14 and May 15 to July 1, and the net change is the levy change relative to the placebo change. Because each cohort is chosen for activity in its first window, 10 to 13 percent of the placebo cohort has no link in the next window, which is why reach falls in the placebo windows; the mean is less affected by these zeros than the median. One row drops accounts with 1,000 or more links in the window (at most 38 per provider and window), since such an account adds its whole circle to the two-step reach of every sender linked to it.

New users (Appendix Table A20). Tigo's registry records the date each account was activated; we count subscriber accounts by that date, in the levy windows and in the same weeks of 2020, 364 days earlier so that weekdays match. The sending and receiving shares are for subscribers registered between June 1 and 30, July 15 and August 17, and September 1 and 29, so that the 14 days after registration fall within one window, and count on-network person-to-person transfers on the day of registration and the 13 days after. Two checks support reading the date as the day the account was opened: 99.0 to 99.8 percent of the accounts in Tigo's transfer data are in the registry, and 0.07 percent of those registered between April and November 2021 transfer before their registry date. M-Pesa's registry fails the second check, with 9.3 percent transferring before the recorded date, and covers only 40 to 54 percent of the accounts whose first transfer falls after May; Airtel's data include no registry.

A.18 The share that spreads: bounds and checks

Table A9 reports the part of the decline that ran through contacts going quiet, from the measured contagion of Table 4 applied to the share of each circle that has gone quiet. Before the levy a quiet contact went with more sending to others, not less (Table 4), so measuring the drag against a circle with no quiet contacts understates the change in what quiet contacts do to their senders rather than overstating it. The floor uses the post-levy association on transfers to receivers outside the pre-levy roster, the outcome that cannot mechanically contain the quiet contacts; the ceiling uses the change in the association on all transfers. The second step adds the same calculation for the contacts of a sender's contacts, and the joint rows estimate both steps together. Holding the sender's own quiet share fixed, the association between her transfers outside the roster and the quiet share among her contacts' contacts falls after the levy by 0.70 at Tigo, 0.38 at M-Pesa, and 0.40 at Airtel (standard errors 0.05, 0.04, and 0.05), about 7, 4, and 4 percentage points per ten points of that share. The main definitions are an eight-week quiet window, quiet meaning no transfer to anyone, contacts

Table A5: Distinct counterparties per sender, by window and provider

Provider	Window	Share of cohort active	Mean	Median	90th percentile
Tigo (703,301 senders)	Pre-levy (Jun 1–Jul 14)	1.00	5.28	3	10
	High rate (Jul 15–Aug 31)	0.72	4.03	2	8
	Post-cut (Sep 1–Oct 13)	0.64	4.03	2	8
M-Pesa (1,156,402 senders)	Pre-levy (Jun 1–Jul 14)	1.00	4.88	3	9
	High rate (Jul 15–Aug 31)	0.68	3.78	2	7
	Post-cut (Sep 1–Oct 13)	0.58	3.64	2	7
Airtel (616,431 senders)	Pre-levy (Jun 1–Jul 14)	1.00	4.52	3	7
	High rate (Jul 15–Aug 31)	0.66	3.56	2	6
	Post-cut (Sep 1–Oct 13)	0.58	3.63	2	6

Notes: Number of distinct receivers in a sender's on-network person-to-person pairs in the window, for the fixed cohort of senders with at least three transfers between June 1 and July 14 (Appendix A.17). Share of cohort active is the share of cohort senders with at least one transfer in the window; the mean and quantiles are over those senders, so the count of an inactive sender, zero, is excluded from them. Complete transaction records of each provider.

Table A6: Retention of a receiver's senders by base-window in-degree, levy windows against placebo windows

Senders in base window	Receivers (levy windows)	Mean retention		Ratio
		Levy windows	Placebo windows	
Tigo				
1	1,502,592	0.444	0.753	0.59
2–4	761,953	0.448	0.742	0.60
5–9	115,795	0.471	0.818	0.58
10–49	34,531	0.492	0.854	0.58
50 or more	2,388	0.513	0.861	0.60
M-Pesa				
1	2,345,823	0.466	0.798	0.58
2–4	1,255,543	0.420	0.694	0.61
5–9	151,677	0.436	0.734	0.59
10–49	39,867	0.497	0.841	0.59
50 or more	4,736	0.486	0.875	0.56
Airtel				
1	1,621,068	0.376	0.657	0.57
2–4	667,841	0.378	0.620	0.61
5–9	58,792	0.430	0.672	0.64
10–49	13,484	0.533	0.831	0.64
50 or more	2,514	0.557	0.950	0.59

Notes: Receivers with at least one sender in the base window, grouped by the number of distinct senders in their on-network person-to-person pairs in that window; all senders, not the cohort. Retention is the receiver's number of distinct senders in the follow-up window divided by the number in the base window, averaged within the group. Levy windows: June 1 to July 14 followed by July 15 to August 31; placebo windows: April 1 to May 14 followed by May 15 to July 1. Ratio is levy-window retention over placebo retention. Complete transaction records of each provider (Appendix A.17).

Table A7: Change in cash-in value by number of pre-levy person-to-person senders to the account

	Tigo			M-Pesa		
Senders to the account	Accounts	Change in log value	Share of fall	Accounts	Change in log value	Share of fall
0	3,430,883	−0.246	0.254	4,552,595	−0.147	0.126
1	1,154,908	−0.311	0.203	1,831,578	−0.323	0.212
2–4	725,979	−0.332	0.310	1,180,892	−0.393	0.370
5–9	114,992	−0.335	0.149	150,284	−0.399	0.169
10–49	34,218	−0.231	0.078	39,676	−0.332	0.109
50 or more	2,340	−0.187	0.007	4,722	−0.275	0.013

Notes: Accounts with any cash-in between June 1 and July 14 or between July 15 and August 31, grouped by the number of distinct senders from whom the account received an on-network person-to-person transfer between June 1 and July 14, zero for accounts that received none; all senders, not the cohort. Change in log value is the log of the group's total cash-in value between July 15 and August 31 over its total between June 1 and July 14; share of fall is the group's fall in value as a share of the fall across all groups. The July 15 to August 31 window is four days longer than the June 1 to July 14 window, which raises every group's change by the same 0.09 and does not affect the comparison across groups; the text reports the changes per day. Airtel's extract carries no cash-in records. Complete transaction records (Appendix A.17).

not yet on the platform identified from every transaction type and counted separately, and sender-volume-decile-by-fortnight controls. The loosest definitions use a four-week window, count a contact as quiet when it sends to no one other than the sender, and drop the volume controls; the community controls add the estimated position of each sender in the pre-levy graph, interacted with fortnight. The loosest bounds differ across providers, running to 42 and 56 percent at Airtel, and we do not use them in the text.

How strong an unobserved trait would have to be. The similarity tests hold fixed the traits we observe. Two standard calculations ask how strong a trait we do not observe would have to be to produce the reversal on its own, each benchmarked on the two strongest observed ones together, the typical activity of a sender's counterparties and the typical size of their transfers, which explain 17 percent of the variation in the quiet share left after the fixed effects and under 0.3 percent of the variation in the sender's own transfers. First, coefficient stability: adding the two controls removes 29 percent of the reversal at Tigo, 38 at M-Pesa, and 23 at Airtel in the Poisson design of Table 4 (activity alone removes 19, 24, and 21 percent, size alone 20, 28, and 3); by the ratio of Altonji et al. (2005), selection on unobservables would have to be 1.5 (M-Pesa) to 2.6 (Airtel) times the observed selection to explain the reversal away, and 2.7 to 3.9 times at Tigo and M-Pesa against activity alone. Second, the bound of Oster (2019) on the within-sender fit: the bias-adjusted coefficient crosses zero only if unobservables are 1.2 (Airtel) to 3.9 (Tigo) times as important as observables, against the conventional threshold of one. Estimated positions in the pre-levy graph, which absorb similarity visible in network structure, explain 1.6 percent of the leftover variation in the quiet share and move the pooled reversal from −12.5 to −12.1 percent per ten points. The reversal is therefore robust to selection as strong as any we observe.

Resemblance and measured contagion. Table A8 reports the test of Section 6 that splits each sender's circle by how much each counterparty resembles her, with resemblance measured as a distance rather than

a shared quantile, on the two traits on which users sort: transfer size everywhere and district at Tigo. The change in measured contagion at the levy fades mildly with distance in size at Tigo and M-Pesa and not at Airtel, is nearly the same for counterparties in the sender's district and elsewhere, and among the counterparties least like her, in another district and far from her in transfer size, is three quarters of the change among those most like her. The homophily share of measured contagion on observed traits is therefore at most about a third, in line with the share the controls remove.

Table A8: The change in measured contagion at the levy, by resemblance between sender and counterparty

	Change in measured contagion at the levy, by bin			Equal across bins
Panel A. Distance in transfer size				
	Within a factor of two	Two to four	Beyond four	p
Tigo	−15.8 (0.58)	−15.1 (0.61)	−13.1 (0.44)	< 0.001
M-Pesa	−13.9 (0.42)	−12.3 (0.44)	−8.8 (0.33)	< 0.001
Airtel	−8.9 (0.65)	−9.9 (0.72)	−8.5 (0.53)	0.26
Panel B. District (Tigo)				
	Same district	Other district		p
Tigo	−15.2 (0.50)	−14.0 (0.40)		0.03
Panel C. District and size together (Tigo)				
	Same district, size within two	Rest	Other district, size beyond four	p, first vs third
Tigo	−16.5 (0.92)	−14.7 (0.43)	−12.3 (0.51)	< 0.001

Notes: The specification of Table 4, column one, with the quiet share of a sender's pre-levy counterparties split into bins by how much the counterparty resembles the sender, each bin interacted with the post-introduction indicator; cells are the post-introduction change in the association for that bin, $100(e^{0.1b} - 1)$ percent per ten percentage points of the bin's quiet share, with standard errors clustered by sender in parentheses, in the same percent units (delta method), and the last column the p-value of a test that the bins' changes are equal. Panel A: distance is the absolute difference between the logarithms of the sender's and the counterparty's mean value per transfer over April 1 to July 14, 2021; the bins hold 37, 28, and 35 percent of counterparties at Tigo, 36, 28, and 36 at M-Pesa, and 35, 27, and 38 at Airtel. Panel B: the sender's and the counterparty's registered three-digit district codes are the same (34 percent of counterparties) or differ, from Tigo's registry. Panel C: the first bin is counterparties in the sender's district with a transfer size within a factor of two of hers (13 percent), the third is counterparties in another district with a transfer size more than four times larger or smaller (23 percent), and the second is the rest. Senders with at least five pre-levy counterparties; sender, fortnight, and volume-decile-by-fortnight fixed effects; complete transaction records. Homophily predicts a change that is largest in the first bin and fades across the row; contagion predicts a flat row. Users sort on both traits: pairs fall in the same quintile of transfer size 1.3 to 1.4 times as often as random pairing would make them, and same-district pairs at Tigo are five times as common; on transaction volume the ratio is 0.99 to 1.08, so that dimension is not used.

The reversal is not the sender's own cutback. A sender who cuts back for her own reasons pays her counterparties less, and a counterparty who lived off those payments may go quiet as a result, so contagion could run from the sender to her circle rather than from the circle to her. To separate the two we measure each counterparty's dependence on the sender, the sender's share of the counterparty's on-network inflows by value from April 1 to July 14, and re-estimate the main specification on the counterparties who received less than a tenth of their inflows from her, who cannot have gone quiet because she cut back, for senders with at least five such counterparties (61,759 at Tigo, 63,652 at M-Pesa, 15,019 at Airtel). Among these counterparties the post-levy change in the association is −19.0 percent per ten points at Tigo ($t = -21$), −12.9 at M-Pesa ($t = -21$) and −6.0 at Airtel ($t = -5.5$), against −24.2, −17.6 and −12.5 on the same

senders' full circles; the pre-levy rerouting is present in every band and the cut changes nothing. The reversal is not concentrated where the reverse channel would be strongest: among counterparties who received more than half of their inflows from the sender it is -9.8 at Tigo, -7.9 at M-Pesa and -7.0 at Airtel. Measured by counts of transfers rather than by value the pattern is the same at Tigo and M-Pesa (-17.2 and -8.2 among independent counterparties), while at Airtel, on 8,897 senders, the independent-counterparty association does not move (-0.8, $t = -0.4$). The gap between a circle's independent counterparties and the whole circle bounds what the reverse channel could contribute, and part of it is only that quietness measured on part of a circle is a noisier measure of the whole.

A circle's state before the levy does not predict its sender's response. With the quiet share of each sender's roster fixed at a pre-levy date and weekly transfers followed from April to October, senders with thin circles and senders with full circles move together before July 15, separate by less than one log point per ten percentage points of the quiet share at the levy (0.002, standard error 0.002, pooled), and move together again at the cut, on an April roster with quietness measured through May or through July 14. The association in Table 4 is therefore with contacts going quiet during the levy period. One version of the design, which fixes the quiet share of the June roster in early June, separates the groups at the levy by 6.6 log points per ten points on transfers to third parties; under the roster's definition a June counterparty quiet through early June is one that had only just begun sending, so that version measures the fragility of newly formed relationships rather than attachment, and the two groups were already moving apart before the levy.

Table A9: Share of the decline that spread from user to user, by definition and provider

	Tigo		M-Pesa		Airtel		Tigo and M-Pesa	
	High rate	After cut	High rate	After cut	High rate	After cut	High rate	After cut
Main definitions, first step (floor to ceiling)	6.4–10.9	14.3–27.4	5.0–7.8	12.7–20.5	4.8–11.2	11.1–22.6	5.8–8.6	13.6–22.0
Second step alone (floor to ceiling)	3.4–3.8	7.2–10.1	1.7–2.6	3.5–5.6	2.0–9.3	2.2–18.2	2.6–3.3	5.3–8.2
Both steps jointly (floor to ceiling)							8.1–11.2	18.4–28.2
First step, partner-type control (floor)	5.2	12.6	3.4	10.8	2.7	9.6	4.4	11.8
First step, community controls (floor to ceiling)							5.5–7.4	12.7–19.1
Loosest definitions, both steps (floor to ceiling)	17.9–21.9	29.4–34.0	13.5–23.8	22.8–36.9	11.0–42.1	15.7–56.1	15.8–23.7	26.1–36.6
Loosest definitions with community controls							14.9–20.1	24.3–31.0

Notes: Percent of the decline in a sender's on-network transfers attributable to counterparties going quiet, senders with at least five observable pre-levy counterparties. High rate: fortnights from July 15 to September 8; after the cut: September 9 to October 6. Floors and ceilings as described in the text; the pooled columns estimate one coefficient across Tigo and M-Pesa with provider effects.

A.19 What the cut did not do

This section reports the tests behind the explanations of the missing recovery that the introduction rules out. Each was written down before it was run.

Providers did not capture the cut. On the complete M-Pesa and Tigo data for June 1 to October 13, the modal provider fee within each of the 23 levy brackets is the same in the weeks before and after September 1, for transfers and for cash-out. It covers at least 95 percent of transactions in every bracket except two

groups: M-Pesa transfers of 1,000 to 1,999 TZS (86 to 89 percent), and Tigo cash-outs of 100,000 TZS or more, where the share at the mode falls after the levy, to about half in the top bracket, because more cash-outs are charged less than the mode; the mean fee in those brackets falls slightly rather than rises. The levy per transaction falls exactly as gazetted, for example from 320 to 224 TZS in the 10,000 to 14,999 bracket. Airtel's data carry no fee field. The user cost fell by the statutory 30 percent of the levy.

Users with nothing to unlearn respond no more. We compare users whose first transaction of any type falls in a given week, who have no habit to break and no relationship to revive, with users active since April, over blocks of four weekly entry cohorts, users entering between July 15 and August 11 and users entering between September 2 and 29. Newcomers' transfers in their first four weeks fall by 0.16 log points at Tigo and 0.17 at M-Pesa from the July block to the September block, while incumbents' rise by 0.06 and 0.04; the difference is 0.22 and 0.21 log points (standard errors 0.02). A lower rate did not make new users transact more.

No tipping within circles. A network that tipped into disuse, one in which each user waits for the others to return, would leave three marks, and none appears. Quitting is smooth in the share of a user's contacts already gone: relative to the pre-levy profile, transfers to third parties fall by 0.17, 0.31, 0.44, 0.51, 0.67, 0.65, and 0.95 log points across bins of the quiet share of width 0.1, with no step, and a split-sample search for a break finds one in the wrong direction. Circles do not leave or return together: within circles of at least five members, the correlation of retention across members, given each member's activity, is 0.04 to 0.08 at Tigo and 0.01 to 0.03 at M-Pesa, against 0.36 or more needed for an all-in or all-out pattern. And returning depends on contacts only at the floor: among senders who sent nothing between July 29 and September 8, the probability of a first renewed transfer between September 9 and October 6 rises with the share of their contacts still active only below about 15 percent, which covers 5 percent of quiet senders, and is flat above it. No chain of exits runs along links: within a circle, the next member to leave is linked to an earlier leaver no more often than a random order among that window's leavers would produce, during the levy or in the spring.

Places hit hardest recover no less than usual. If a low-use state were held in place locally, districts whose taxed activity fell furthest in July should recover least at the cut, beyond what local shocks usually leave behind. Regressing each district's post-cut change on its July fall, against the same regression across levy-free windows in the spring, gives nothing at Tigo (169 districts; the levy-era slope differs from the placebo slope by -0.15, standard error 0.15, unweighted, and by $+0.13$, standard error 0.11, weighted) and a mixed result at M-Pesa (175 districts; $+0.34$, standard error 0.16, unweighted; -0.14, standard error 0.10, weighted). Almost every district rose a little after the cut, with no split into recovered and stuck places.

The smallest levied bracket is not a clean test. Transfers of 1,000 to 1,999 TZS carried a 10 TZS levy in both regimes and were never cut. Their daily count fell by 0.20 log points at the introduction, the same as

the never-levied transfers below 1,000 TZS and far less than the 0.58 of the cut brackets, and nothing moved between the groups at the cut. Read at face value this says the response tracks the amount at the bottom of the schedule; but the comparison inherits the composition problem of Appendix A.5, since transfers shrink after the levy and small brackets are filled from above, and cash-out shows the mechanics in reverse, with the smallest bracket falling more than the cut brackets. We do not use it for either reading.

A.20 The deadweight-loss figure: weighting and robustness

Section 7 values the surplus lost on the transactions the levy stopped at half the levy, $\frac{1}{2}\tau \cdot \Delta Q$, and divides it by the revenue collected on the value that remains, $\tau \cdot Q_1 = \tau \cdot (1 - \Delta Q/Q_0) \cdot Q_0$, so the ratio is $\frac{1}{2}(\Delta Q/Q_0)/(1 - \Delta Q/Q_0)$; with $\Delta Q/Q_0 = 0.482$, it is $0.5 \cdot 0.482/0.518 \approx 0.47$. The fall is the treated arm of the main regression: the post-July-15 change in log daily total value of each taxed series (0.66 log points, standard error 0.09, clustered by series), averaged with equal weight, which is the weighting of the difference-in-differences coefficient itself. Weighting the taxed series by value instead gives a fall of 0.53 log points, 41 percent, and a ratio of 0.35. On the Panel B comparison, $\Delta Q/Q_0 = 0.52$ and the ratio is 0.54.

Three features of the setting could move the figure. Replacing linear demand with constant-elasticity demand, holding the measured decline fixed, lowers the ratio from 0.47 to 0.42. Charging the levy against the pre-existing distortion rather than an undistorted baseline raises it to between 0.67 and 0.74: each transaction the levy stops also loses the 10 percent excise and the 18 percent VAT on its provider fee, which we measure at 12 to 16 percent of the levy collected at the introduction rate (Section 7.2); the rest of the fee is a market price rather than a wedge. The notched schedule raises the local welfare cost at bracket boundaries, but the bracket-level response is explained by transaction size rather than by the effective rate (Section 5), and the difference-in-bunching estimates of Appendix A.13 put the transactions relocated across a boundary at one to two percent of those missing from the taxed base, so the notch term is second order. The two adjustments that move the figure pull in opposite directions; the defensible range around 47 percent runs from 42 to 74.

A.21 The MVPF: inputs, calibration, and conventions

Following Hendren and Sprung-Keyser (2020), Section 7 computes the MVPF as the ratio of mechanical to net revenue. If the behavioral response removes a share x of the revenue the tax would have raised on the old base, each shilling actually raised costs $1/(1-x)$; for a tax that moves the base from Q_0 to $Q_0(1-x)$, that share is the proportional decline in the base itself. Two features of our setting make x better identified than in most applications: we measure it directly from the fall in the taxed base rather than extrapolating it from an estimated elasticity, and we observe it separately under each of the two rate regimes.

The levy's MVPF is 1.93 at the introduction rate and 1.82 at the reduced rate, within a standard error of each other. The inputs are the same falls in taxed value as above, measured over June 1 to October 13: the taxed series stand 48.2 percent below their pre-levy level at the introduction and 45.1 percent under the reduced rate (0.658 and 0.600 log points, standard errors 0.086 and 0.089, clustered by series), so $x = 0.482$ and 0.451. These are the whole fall in taxed activity, counted as the levy's for the reasons given above. The

differential response against the untaxed control, 0.399 log points at the introduction (standard error 0.111) and 0.462 under the reduced rate (standard error 0.124), gives $x = 0.329$ and 0.370 and an MVPF of 1.49 and 1.59, the floor. The floor orders the two regimes the other way: the untaxed control recovers more after September than the taxed series do, so the differential widens while the taxed series' own fall narrows.

For the value-added tax and the payroll tax, for each comparator the share of revenue lost to the behavioral response is the elasticity of the base times the rate over one minus the rate, $x \approx \varepsilon\tau/(1-\tau)$, the fiscal-externality term Hendren and Sprung-Keyser (2020) use, where τ is the statutory rate and ε the elasticity of the base with respect to $(1-\tau)$. Tax rates are taken from the Tanzania Revenue Authority; base elasticities are calibrated from the empirical literature on tax bases in comparable settings, 0.5 for VAT (Bachas et al., 2024; Brockmeyer et al., 2024) and 0.4 for the payroll tax (Best et al., 2015; Kleven and Waseem, 2013). We do not compute an MVPF for the property tax. Its statutory base is close to inelastic, so the identity returns a value indistinguishable from 1, and the compliance and administrative costs that would move it are not something we can quantify from our data.

The comparator MVPFs rely on base elasticities calibrated from comparable settings, not estimated from Tanzanian data, so their precision is lower than that of the MVPF for the levy itself, which inherits the precision of our difference-in-differences estimates. The qualitative ranking is robust across plausible perturbations of the calibration: the payroll tax's MVPF would have to reach 1.82 for the ranking to flip, which requires its calibrated elasticity to rise from 0.4 to approximately 1.05, and 1.49 to flip against the floor, which requires approximately 0.77.

Finkelstein and Hendren (2020) apply a large-reform correction to their tariff example, valuing half the behavioral revenue loss at the margin and giving MVPF $= (1-x/2)/(1-x)$. That correction is appropriate for a reform of this size and would place the levy at 1.47 and 1.41 rather than 1.93 and 1.82. We report the uncorrected figures because the library values we benchmark against are computed without the correction, and the comparison requires the same convention on both sides. The pre-existing 10 percent excise and 18 percent VAT on provider fees generate a further fiscal externality when the base contracts. Measured from the fees in the data, the excise and VAT lost equal a share f of 12 to 16 percent of the levy collected at the introduction rate and 20 to 24 percent at the reduced rate (Section 7.2). Counting it multiplies each MVPF by $1/(1-f)$, raising 1.93 to between 2.2 and 2.3 and 1.82 to between 2.3 and 2.4; we omit it from the reported figures for the same reason.

A.22 The levy's MVPF against the MVPF library

These numbers can be read against the MVPF library compiled by Hendren and Sprung-Keyser (2020) for U.S. policies. The comparable set is their five top marginal rate reforms, which are the revenue-raising instruments in the library and are computed on the same basis as our figures: 1.16 for the 2013 reform, 1.37 for 2001, 1.85 for 1993 (95 percent confidence interval 1.19 to 4.07), 44.27 for 1986, and infinity for the 1981 cut. Finkelstein and Hendren (2020) add 1.2 to 1.5 for the 2018 U.S. tariffs on Chinese imports, the lower figure coming from the specification that incorporates spillovers and terms-of-trade effects. Our figures, 1.93 at the introduction rate and 1.82 after the cut, exceed the 2013 and 2001 reforms and the tariff

range and straddle the 1993 estimate of 1.85.

The two large values are not informative upper comparators. Both reforms cut a top rate of 50 percent or more, where the behavioral response offsets nearly all of the mechanical revenue change and the net cost to government approaches zero, so the ratio diverges. The 1986 estimate implies a net cost of 0.02 dollars per dollar of mechanical revenue, with a confidence interval from 2.37 to infinity; for 1981 Hendren and Sprung-Keyser (2020) report a fiscal externality above 1, meaning the pre-reform rate was on the wrong side of the Laffer curve, and they cannot reject an MVPF of either 1 or infinity. The MVPF is an unstable statistic in that region because it is a ratio of two quantities both approaching zero. Our own fiscal externality of 0.45 to 0.48 sits well away from that boundary, which is why the levy's MVPF is a stable number rather than a near-singular ratio.

Two differences in construction qualify the comparison. The library's tax-reform values apply a Pareto adjustment specific to reforms of a top bracket, which has no analogue in a levy applied across the whole transaction distribution, and they infer the behavioral response from estimated elasticities of taxable income rather than from a measured base decline. The remainder of the library consists of spending programs, for which the MVPF measures value delivered per dollar of government outlay rather than burden per dollar raised; those values sit on the other side of the budget and are not comparators for a revenue instrument. The structural parallel with tariffs remains direct, both being taxes on a relatively elastic quantity base, though Finkelstein and Hendren (2020) note that the tariff range they compute is similar to the cost of raising revenue through the U.S. income tax. The levy sits above that range.

A.23 The cost of public funds on a network

This appendix derives the factor used in Section 7.2. User i chooses transfers $x_i \geq 0$ with utility $U_i = v_i(x_i) + \delta\, x_i \bar{x}_i - (\tau_0 + t)\, x_i$, where v_i is concave, $\bar{x}_i = \sum_j w_{ij} x_j$ is the weighted activity of i's counterparties with weights $w_{ij} \geq 0$, t is the levy and τ_0 the pre-levy user cost. The network term is the linear complementarity of Ballester et al. (2006): what a transfer is worth to a sender rises with how active her counterparties are. We assume that counterparties matter to a user only through her own transfers, that levy revenue is a transfer and τ_0 a resource cost, and that the equilibrium is interior.

The first-order condition is $v_i'(x_i) + \delta \bar{x}_i = \tau_0 + t$. Differentiating, with $c_i = -v_i''(x_i) > 0$, $c_i\, dx_i/dt = -1 + \delta\, d\bar{x}_i/dt$. Holding counterparties fixed the response would be $-1/c_i$, the private response; the remainder, $n_i = (\delta/c_i)\, d\bar{x}_i/dt$, is the part caused by counterparties, and $s_i = n_i/(dx_i/dt)$ is its share. The identities $\delta\, d\bar{x}_i/dt = c_i n_i$ and $c_i = -1/[(1-s_i)\, dx_i/dt]$ give $\delta x_i\, d\bar{x}_i/dt = -x_i s_i/(1-s_i)$. By the envelope theorem user i's loss from a marginal levy is

$$-\frac{dU_i}{dt} \;=\; x_i - \delta x_i \frac{d\bar{x}_i}{dt} \;=\; \frac{x_i}{1-s_i},$$

the levy paid plus the loss from counterparties transacting less, which the user does not choose. Summing over users, the burden is $\sum_i x_i/(1-s_i)$ against a mechanical burden of $\sum_i x_i$, and net revenue is $Q(1-x)$ with

$Q = \sum_i x_i$ and x the share of mechanical revenue lost to the response, so

$$\text{MVPF}_{\text{network}} \;=\; \bar{m}\,\text{MVPF}_{\text{standard}}, \qquad \bar{m} \;=\; \frac{\sum_i x_i/(1-s_i)}{\sum_i x_i}.$$

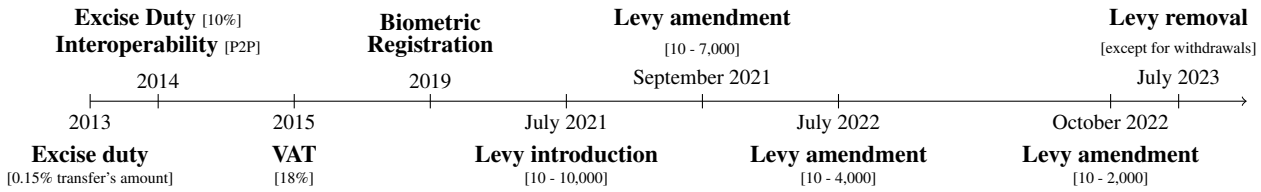
With a common share, $\bar{m} = 1/(1-s)$, the social multiplier of Glaeser et al. (2003), and the MVPF is $1/[(1-x)(1-s)]$ in the convention of Section 7.2. The same formula holds when users choose only whether to transact, with s the share of the participation response due to others' participation. A numerical check on a 400-user network with heavy-tailed degrees reproduces the formula exactly with heterogeneous s_i; the common-share approximation $1/(1-s)$ understates the exact factor by at most 0.01.

The share s is the share of the decline that spread from user to user in Table A9, both steps, at the lower end of its range and in the October window: 18 percent between September 9 and October 6, applied at both rates for the reason given in Section 7.2. It is measured against each sender's own fall in on-network transfers, the same whole-fall convention as x. The exact factor weights $1/(1-s_i)$ by activity, so with a share measured on the aggregate decline $1/(1-s)$ is an approximation; the bounds come from senders with at least five pre-levy counterparties, and applying them platform-wide assumes a similar share elsewhere.

A.24 Taxes and charges on mobile money transactions

Pre-existing taxes. The levy was applied on top of two taxes already in force on mobile money charges. The excise duty was introduced in July 2013 as a percentage of the transaction value (0.15 percent of the amount transferred, for transactions above 30,000 TZS). One year later, in 2014, it was replaced by a 10 percent duty on the charges of service providers. The standard 18 percent value-added tax (VAT) has applied to these charges since July 2015. Figure A1 shows the timeline of tax changes on mobile money in Tanzania.

Figure A1: Timeline of major regulatory and tax changes on mobile money



Notes: Based on Official Documents.

The levy. The levy was introduced on July 15, 2021, motivated by a widening fiscal deficit, partly driven by the economic impact of the COVID-19 pandemic. It consists of transaction-specific fixed amounts that vary with the transaction value. The initial schedule had 23 brackets, each with clearly defined lower and upper bounds. By construction, the average tax rate decreases within each bracket as the transaction value increases, with discrete jumps in tax liability at the thresholds.[32] This is a *notched* scheme, which typically

[32]The average tax rate is the ratio between the amount paid as a result of the levy and the transfer's value.

induces intensive-margin responses such as bunching around thresholds (Kleven, 2016).

Scope. The levy targeted two transaction types: person-to-person transfers and cash-out (withdrawals). All other types, including merchant payments, business payments, government transactions, and cash-in, were never subject to it. Since July 2023 the levy no longer applies to person-to-person transfers and is levied only on withdrawals. Table A10 summarizes the treatment by transaction type and its evolution.

Collection. The levy was collected at the source. When a customer initiated a transfer, the mobile money provider (for example, Vodacom) withheld the levy together with its own service fee. The levy was computed on the value of the transaction sent.

Schedule. The government partially reduced the levy on September 1, 2021, seven weeks after its introduction, and amended it again in July and October 2022. Table A12 reports the thresholds and levy amounts before and after the September 2021 reduction. The introduction followed by a partial reduction generates an increase and a subsequent decrease in transaction costs, albeit of different magnitudes, which allows us to study asymmetric responses to changes in the price of mobile money.

Table A10: Levy's regulatory scheme

	Period				
	July 2021	September 2021	July 2022	October 2022	July 2023
Number of brackets	23	23	22	22	22
Targeted (affected by the levy)					
P2P	Yes	Yes	Yes	Yes	No
P2B	No	No	No	No	No
B2P	No	No	No	No	No
P2G	No	No	No	No	No
G2P	No	No	No	No	No
Cash-in (deposits)	No	No	No	No	No
Cash-out (withdrawals)	Yes	Yes	Yes	Yes	Yes
Legal documents					
Identification	GN 496A	GN 642A	GN 478V	GN 596C	FA 2023
Published on	30Jun2021	31Aug2021	5Jul2022	5Oct2022	30Jun2023

Source: the Government Notices (GN) and the Finance Act (FA) 2023 identified in the table; each identifier links to the published text.

User cost. To normalize and compare the effects of the two policy interventions, we use the *user cost* of a mobile money transaction. It has four components: (i) the provider fee, (ii) the excise duty, (iii) the VAT, and (iv) the levy. The excise duty and the VAT are proportional to the provider fee, whereas the levy is a fixed amount that depends on the transaction type and value. The total user cost (tax-inclusive), which the

customer pays on top of the amount sent, is

$$\text{Customer deduction}_{t,j,p} = \left[Charge_{t,j,p} \cdot (1+0.10)\right] \cdot (1+0.18) + Levy_{t,j}, \tag{2}$$

where t denotes the transaction type, j the transaction amount, and p the mobile money provider. The provider charge is set by each provider and varies with the amount and with whether the transaction involves users of different providers (typically, sending money to a user of another provider is more expensive). The levy depends only on t and j.

Taxes on service charges are typically embedded in the displayed prices: mobile money fees are usually shown inclusive of the excise duty and VAT. Taxes levied directly on transaction values, such as the levy, are instead presented to users as an additional, separate cost. Hearson et al. (2024) document this practice across several African countries. Figure A2 reproduces the schedules published by Tigo and Airtel at the introduction of the levy, which list the provider fee and the levy in separate columns.

Illustration. Consider a person-to-person transfer of 40,000 TZS (approximately USD 17) between users of the same provider. Before the levy, the total user cost was 410 TZS, an average effective cost of 1.03 percent.[33] After the introduction of the levy (1,500 TZS), the cost increased nearly fivefold, to 1,910 TZS (4.78 percent). Taxes account for most of this increase: approximately 83 percent of the 1,910 TZS (1,594 TZS) is tax, namely 1,500 TZS of levy plus 94 TZS of excise duty and VAT on the fee. After the September 1 reduction, the levy fell to 1,050 TZS and the total cost to 1,460 TZS (3.65 percent), with the provider fee unchanged (Section 5.2).

Table A11: User cost of a 40,000 TZS person-to-person transfer within the same provider

	Before levy	Jul 15 – Aug 31	After Sep 1
Provider fee (TZS)	315.87	315.87	315.87
Excise + VAT on fee (TZS)	94.13	94.13	94.13
Levy (TZS)	0	1,500	1,050
Total user cost (TZS)	**410**	**1,910**	**1,460**
Share of the amount (%)	1.03	4.78	3.65

[33] $410 = 315.87$ [provider fee] $\times (1+0.10)$ [excise] $\times (1+0.18)$ [VAT].

Figure A2: Provider fee schedules and levy at the introduction (July 2021)



tigo pesa

Ni Zaidi ya Pesa

Ada za wateja. Piga *150*01#

Kiwango (Tsh)		Kutuma pesa kwenda Tigo Pesa			Kutuma pesa kwenda mitandao washirika			Kutoa pesa kwa Wakala/ATM		
Kuanzia	Mwisho	Ada ya Tigo Pesa	Tozo ya Serikali	Jumla ya Makato	Ada ya Tigo Pesa	Tozo ya Serikali	Jumla ya Makato	Ada ya Tigo Pesa	Tozo ya Serikali	Jumla ya Makato
100	999	15	-	15	15	-	15	-	-	-
1,000	1,999	30	10	40	50	10	60	300	10	310
2,000	2,999	30	16	46	50	16	66	400	16	416
3,000	3,999	50	27	77	100	27	127	600	27	627
4,000	4,999	60	56	116	100	56	156	650	56	706
5,000	6,999	130	100	230	200	100	300	950	100	1,050
7,000	9,999	150	125	275	200	125	325	1,000	125	1,125
10,000	14,999	360	320	680	550	320	870	1,450	320	1,770
15,000	19,999	360	610	970	550	610	1,160	1,450	610	2,060
20,000	29,999	380	960	1,340	600	960	1,560	1,850	960	2,810
30,000	39,999	400	1,100	1,500	680	1,100	1,780	1,850	1,100	2,950
40,000	49,999	410	1,500	1,910	750	1,500	2,250	2,350	1,500	3,850
50,000	99,999	720	2,050	2,770	1,250	2,050	3,300	2,700	2,050	4,750
100,000	199,999	1,000	2,530	3,530	1,600	2,530	4,130	3,650	2,530	6,180
200,000	299,999	1,200	2,940	4,140	1,900	2,940	4,840	5,300	2,940	8,240
300,000	399,999	1,500	3,500	5,000	2,300	3,500	5,800	6,500	3,500	10,000
400,000	499,999	1,500	4,100	5,600	2,500	4,100	6,600	7,000	4,100	11,100
500,000	599,999	2,200	5,200	7,400	3,200	5,200	8,400	7,500	5,200	12,700
600,000	699,999	3,300	6,400	9,700	4,300	6,400	10,700	8,000	6,400	14,400
700,000	799,999	3,300	7,100	10,400	4,300	7,100	11,400	8,000	7,100	15,100
800,000	899,999	3,500	7,520	11,020	4,300	7,520	11,820	8,000	7,520	15,520
900,000	1,000,000	3,500	8,900	12,400	6,000	8,900	14,900	8,000	8,900	16,900
1,000,001	3,000,000	5,000	9,400	14,400	6,000	9,400	15,400	8,000	9,400	17,400
3,000,001	10,000,000	5,000	10,000	15,000	6,000	10,000	16,000	10,000	10,000	20,000

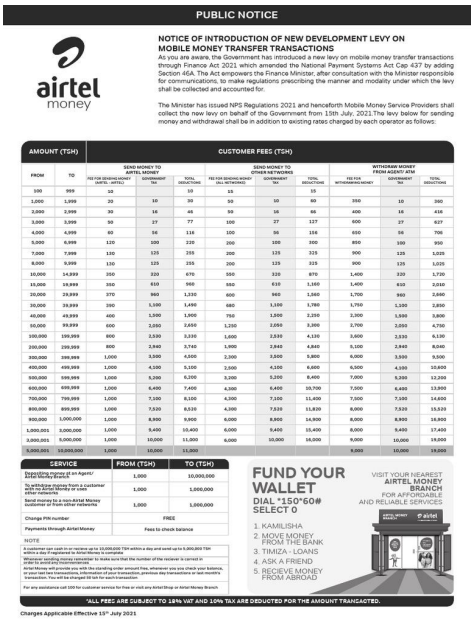
tigo pesa

www.tigo.co.tz

Kuwasaliana na huduma kwa wateja,
Piga **100** bure au tembelea duka lolote la Tigo!
Ada hizi ni kuanzia July 2021

tigo

(a) Tigo Pesa



PUBLIC NOTICE

airtel money

NOTICE OF INTRODUCTION OF NEW DEVELOPMENT LEVY ON MOBILE MONEY TRANSFER TRANSACTIONS

As you are aware, the Government has introduced a new levy on mobile money transfer transactions through Finance Act 2021 which amended the National Payment Systems Act Cap 437 by adding Section 46A. The Act empowers the Finance Minister, after consultation with the Minister responsible for communications, to make regulations prescribing the manner and modality under which the levy shall be collected and accounted for.

The Minister has issued NPS Regulations 2021 and henceforth Mobile Money Service Providers shall collect the new levy on behalf of the Government from 15th July, 2021.The levy below for sending money and withdrawal shall be in addition to existing rates charged by each operator as follows:

AMOUNT (TSH)		CUSTOMER FEES (TSH)								
		SEND MONEY TO AIRTEL MONEY			SEND MONEY TO OTHER NETWORKS			WITHDRAW MONEY FROM AGENT/ ATM		
FROM	TO	FEE FOR SENDING MONEY (AIRTEL - AIRTEL)	GOVERNMENT TAX	TOTAL DEDUCTIONS	FEE FOR SENDING MONEY (ALL NETWORKS)	GOVERNMENT TAX	TOTAL DEDUCTIONS	FEE FOR WITHDRAWING MONEY	GOVERNMENT TAX	TOTAL DEDUCTIONS
100	999	10		10	15		15			
1,000	1,999	20	10	30	50	10	60	350	10	360
2,000	2,999	30	16	46	50	16	66	400	16	416
3,000	3,999	50	27	77	100	27	127	600	27	627
4,000	4,999	60	56	116	100	56	156	650	56	706
5,000	6,999	120	100	220	200	100	300	850	100	950
7,000	7,999	130	125	255	200	125	325	900	125	1,025
8,000	9,999	130	125	255	200	125	325	900	125	1,025
10,000	14,999	350	320	670	550	320	870	1,400	320	1,720
15,000	19,999	350	610	960	550	610	1,160	1,400	610	2,010
20,000	29,999	370	960	1,330	600	960	1,560	1,700	960	2,690
30,000	39,999	390	1,100	1,490	680	1,100	1,780	1,750	1,100	2,850
40,000	49,999	400	1,500	1,900	750	1,500	2,250	2,300	1,500	3,800
50,000	99,999	600	2,050	2,650	1,250	2,050	3,300	2,700	2,050	4,750
100,000	199,999	800	2,530	3,330	1,600	2,530	4,130	3,600	2,530	6,130
200,000	299,999	800	2,940	3,740	1,900	2,940	4,840	5,100	2,940	8,040
300,000	399,999	1,000	3,500	4,500	2,300	3,500	5,800	6,000	3,500	9,500
400,000	499,999	1,000	4,100	5,100	2,500	4,100	6,600	6,500	4,100	10,600
500,000	599,999	1,000	5,200	6,200	3,200	5,200	8,400	7,000	5,200	12,200
600,000	699,999	1,000	6,400	7,400	4,300	6,400	10,700	7,500	6,400	13,900
700,000	799,999	1,000	7,100	8,100	4,300	7,100	11,400	7,500	7,100	14,600
800,000	899,999	1,000	7,520	8,520	4,300	7,520	11,820	8,000	7,520	15,520
900,000	1,000,000	1,000	8,900	9,900	6,000	8,900	14,900	8,000	8,900	16,900
1,000,001	3,000,000	1,000	9,400	10,400	6,000	9,400	15,400	8,000	9,400	17,400
3,000,001	5,000,000	1,000	10,000	11,000	6,000	10,000	16,000	9,000	10,000	19,000
5,000,001	10,000,000	1,000	10,000	11,000				9,000	10,000	19,000

SERVICE	FROM (TSH)	TO (TSH)
Depositing money at an Agent/ Airtel Money Branch	1,000	10,000,000
To withdraw money from a customer registered on Airtel or own other networks	1,000	1,000,000
Send money to a non-Airtel Money customer or from other networks	1,000	1,000,000
Change PIN number	FREE	
Payments through Airtel Money	Fees to check balance	

NOTE

A customer can cash in or recieve up to 10,000,000 TSH within a day and send up to 5,000,000 TSH within a day as applied by Airtel Money to comply.

Whenever sending money remember to make sure that the number of the reciever is correct in order to avoid any inconveniences.

Airtel Money will provide you with the standing order amount free, whenever you you check your balance, however you will be charged 10 tsh for each transaction. Limitless day transactions or last month's transaction. You will be charged 10 tsh for each transaction.

For any assistance call 100 for customer service for free or visit any Airtel Shop or Airtel Money Branch

FUND YOUR WALLET

DIAL *150*60#
SELECT 0

1. KAMILISHA
2. MOVE MONEY FROM THE BANK
3. TIMIZA - LOANS
4. ASK A FRIEND
5. RECIEVE MONEY FROM ABROAD

VISIT YOUR NEAREST **AIRTEL MONEY BRANCH** FOR AFFORDABLE AND RELIABLE SERVICES

*ALL FEES ARE SUBJECT TO 18% VAT AND 10% T&C ARE DEDUCTED FOR THE AMOUNT TRANSACTED.

Charges Applicable Effective 15th July 2021

(b) Airtel Money

Notes: Tigo Pesa fee schedule and Airtel Money public notice, effective July 2021. Both list the provider fee and the government levy in separate columns, and their sum as the total deduction. In the Tigo schedule (in Swahili), *Ada ya Tigo Pesa* is the provider fee, *Tozo ya Serikali* the government levy, and *Jumla ya Makato* the total deductions.

Table A12: Mobile money levy schedule, July–September 2021

	Bracket bounds (TZS)		Levy amount (TZS)
Bracket	Lower	Upper	July / September 2021
1	1,000	1,999	10 / 10
2	2,000	2,999	16 / 11
3	3,000	3,999	27 / 19
4	4,000	4,999	56 / 39
5	5,000	6,999	100 / 70
6	7,000	9,999	125 / 88
7	10,000	14,999	320 / 224
8	15,000	19,999	610 / 427
9	20,000	29,999	960 / 672
10	30,000	39,999	1,100 / 770
11	40,000	49,999	1,500 / 1,050
12	50,000	99,999	2,050 / 1,435
13	100,000	199,999	2,530 / 1,771
14	200,000	299,999	2,940 / 2,058
15	300,000	399,999	3,500 / 2,450
16	400,000	499,999	4,100 / 2,870
17	500,000	599,999	5,200 / 3,640
18	600,000	699,999	6,400 / 4,480
19	700,000	799,999	7,100 / 4,970
20	800,000	899,999	7,520 / 5,264
21	900,000	1,000,000	8,900 / 6,230
22	1,000,001	3,000,000	9,400 / 6,580
23	3,000,001	—	10,000 / 7,000

Notes: Fixed levy amounts applied to person-to-person transfers and cash withdrawals. Bracket boundaries unchanged across the two regimes; levy amounts reduced by approximately 30 percent on September 1, 2021. Source: Government Notice 496A (June 30, 2021) and Government Notice 642A (August 31, 2021).

Table A13: Effective levy rate at representative transaction sizes

Transaction value (TZS)	July 2021	September 2021
5,000	2.00%	1.40%
20,000	4.80%	3.36%
40,000	3.75%	2.63%
100,000	2.53%	1.77%
500,000	1.04%	0.73%
3,000,001+	0.33%	0.23%

Notes: Levy as a share of transaction value at the lower bound of the bracket containing each transaction size. Excludes the provider's service fee, the 10 percent excise duty, and the 18 percent VAT, which apply separately to the fee.

Table A14: Difference-in-differences response by sender gender and district (Tigo subsample)

Stratum	$\hat{\beta}_1$ on log(total value)	Standard error	Observations
Panel A. Sender gender			
Female	−0.953	(0.031)	184
Male	−0.853	(0.032)	184
Panel B. Sender district			
Urban (Dar es Salaam + districts)	−0.965	(0.031)	184
Rural (rest of country)	−0.843	(0.036)	184

Notes: Subsample difference-in-differences estimates of equation (1), taxed transfers against merchant payments, June 1 to August 31, 2021, on Tigo's complete data with sender attributes from Tigo's user file. Panel A: transactions whose sender's gender is known; the user file covers 99 percent of cohort accounts but gives gender for 57 percent, and less often for senders who were inactive after the cut (44 percent) than for those still active (61 percent), so Panel A is estimated on a selected subsample. Panel B: urban means the sender's registered address carries a Dar es Salaam postcode (prefixes 11 to 17, the region's five districts; postcodes are present for 98.5 percent of registry accounts), rural all other postcode-bearing addresses. Outcome: log daily total value; HC1 standard errors in parentheses; observations are 92 days of two series. The urban–rural difference of −0.122 has standard error 0.047.

Table A15: Relationship survival across the levy, by provider

	Pre-levy pairs (Jun 1–Jul 14)	Share transacting again Jul 15–Aug 31	May 15–Jul 1 (placebo)	Relative decline (%)	Never again (%)
Tigo	3,716,041	0.187	0.270	−30.9	75.6
M-Pesa	5,639,870	0.165	0.246	−33.0	78.7
Airtel	2,783,494	0.170	0.251	−32.4	78.0

Notes: A pair is a sender–receiver combination with at least one on-network person-to-person transfer in the window (Airtel's transfer file is on-network by construction), restricted to senders with at least three transfers in the respective base window. Pre-levy pairs transact between June 1 and July 14, 2021 and are followed from July 15 to August 31 (forty-eight days). The placebo repeats the calculation for pairs transacting April 1 to May 14, followed from May 15 to July 1 (forty-eight days). Relative decline is the July 15 to August 31 share relative to the placebo share. Never again means no further transaction through October 13, 2021. Computed on the complete transaction records of each provider. Definitions in Appendix A.17.

Table A16: Daily transacted value by transaction type across the policy windows

	Daily value (TZS billion)			Δlog value		Δlog count
Transaction type	Pre-levy	High rate	Post-cut	Introduction	Post-cut	Introduction
Cash-out (taxed)	81.4	54.8	57.2	−0.40	−0.35	−0.32
Cash-in	71.9	48.8	51.2	−0.39	−0.34	−0.17
P2P (taxed)	45.6	20.4	21.4	−0.81	−0.76	−0.48
B2P	24.0	16.9	19.9	−0.35	−0.19	−0.27
P2B	14.4	12.9	15.6	−0.11	+0.08	−0.13
P2G	0.6	0.5	0.5	−0.13	−0.06	+0.03
All types	237.9	154.2	165.8	−0.43	−0.36	

Notes: Mean daily transacted value and its log change, on the complete transaction records of the three providers with the clean assignment of Appendix A.2 (person-to-person transfers pool on-network and inter-operator transfers, including M-Pesa's off-net traffic booked under person-to-business payments). Windows: pre-levy June 1 to July 14, high rate July 15 to August 31, post-cut September 1 to October 13, 2021. Log changes are relative to June 1 to July 14. Cash-in is M-Pesa and Tigo only; payments to government are Airtel only in this window. Business-to-person payments exclude transfers arriving from other networks and Vodacom's payouts from its own accounts, as in Section 5.4. Every type fell at the introduction and none absorbed the loss; only merchant payments stood above their pre-levy level after the cut.

Table A17: Activity by transaction type for a fixed cohort of senders, complete data

	Daily count		Daily value		Mean size	
	High rate	Reduced rate	High rate	Reduced rate	High rate	Reduced rate
P2P (taxed)	52.6	50.5	40.8	40.7	77.6	80.6
Cash-out (taxed)	63.5	60.1	60.8	58.6	95.9	97.5
P2B	82.8	82.3	86.4	95.6	104.4	116.2
P2G	93.1	88.2	82.0	81.1	88.1	92.0

Notes: Per-day activity per cohort sender by transaction type, indexed to the pre-levy window (June 1 to July 14, 2021) = 100; high rate is July 15 to August 31, reduced rate is September 1 to October 13. The cohort is the 3,821,187 senders with at least three person-to-person transfers between June 1 and July 14, inter-operator transfers included, three providers, complete data (the network exhibits count on-network transfers only and follow 2,476,134 senders); a sender with no activity of a type in a window counts as zero. Only the types the customer initiates are observed for a sender; off-net transfers are counted as person-to-person transfers (Appendix A.2). The share of the cohort initiating no transaction of any type is 9.9 percent between July 15 and August 31, 16.1 percent between September 1 and October 13, and 6.1 percent in both. Among the fifth of senders with the largest own fall in per-day transfers, cash-out falls by 46 percent, merchant payments by 26 percent, and mean withdrawal size by 3 percent.

Table A18: Difference-in-differences response by pre-policy person-to-person transaction-size decile

Decile	$\hat{\beta}_1$ on log(total value)	Standard error	Observations
1	0.003	(0.104)	552
2	−0.212	(0.067)	552
3	−0.298	(0.073)	552
4	−0.606	(0.039)	552
5	−0.509	(0.028)	552
6	−0.618	(0.034)	552
7	−0.500	(0.048)	552
8	−0.540	(0.046)	552
9	−0.728	(0.088)	552
10	−0.825	(0.117)	552

Notes: Subsample difference-in-differences estimates of equation (1), taxed transfers against merchant payments (the assignment of Appendix A.2), narrow window (June 1 – August 31, 2021), one series per provider and line with series and provider-by-day fixed effects (six series), restricted to transactions in each pre-policy decile of taxed-transfer amount; both the treated and the control series are restricted to the decile. Deciles are defined on the universe distribution of taxed-transfer values during the pre-policy window (June 1 – July 14, 2021). Outcome is log(daily total transaction value). Standard errors clustered by series in parentheses. Computed on the complete transaction records.

A.25 Additional tables

A.26 Additional figures

This subsection collects the supplementary and robustness figures referenced in the text.

Table A19: Reach, cohesion, and connectivity of the transfer network, by window and provider

	Pre-levy (Jun 1–Jul 14)	High rate (Jul 15–Aug 31)	Post-cut (Sep 1–Oct 13)	Placebo change (%)	Net of placebo (%)
Accounts transacted with directly (mean)					
Tigo	7.81	4.42	3.97	−7.2	−38.9
M-Pesa	7.35	4.01	3.28	−10.8	−38.9
Airtel	6.49	3.53	3.11	−13.5	−37.2
Accounts within two steps (median)					
Tigo	51	13	10	−17.0	−69.3
M-Pesa	53	10	6	−29.4	−73.3
Airtel	26	6	4	−34.6	−64.7
Accounts within two steps (mean)					
Tigo	269.4	107.9	89.9	+10.2	−63.7
M-Pesa	289.3	119.8	84.5	−3.5	−57.1
Airtel	152.3	49.1	38.6	−17.3	−61.0
Accounts within two steps (mean, without accounts with 1,000+ links)					
Tigo	100.7	46.1	38.1	−5.7	−51.4
M-Pesa	141.2	51.7	37.9	−14.7	−57.1
Airtel	118.1	48.5	37.8	−15.4	−51.5
Support ratio (mean)					
Tigo	0.168	0.125	0.113	+0.1	−25.6
M-Pesa	0.138	0.101	0.087	−0.9	−25.9
Airtel	0.157	0.119	0.113	−4.9	−20.6
Share of active accounts in the largest connected part					
Tigo	0.877	0.792	0.783	+0.8	−10.4
M-Pesa	0.911	0.824	0.804	+1.3	−10.7
Airtel	0.853	0.732	0.725	+1.2	−15.1

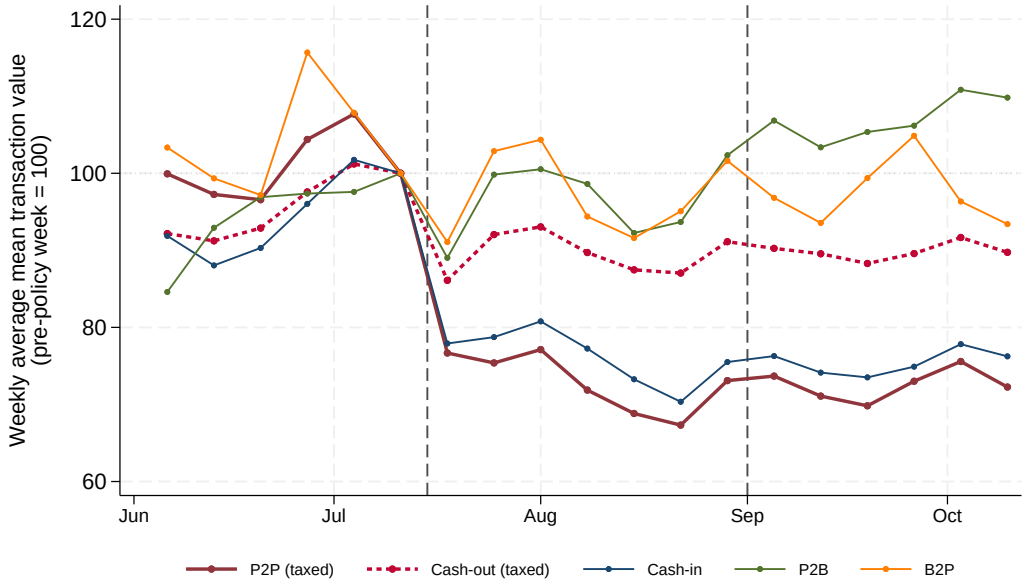
Notes: Undirected graph of on-network person-to-person transfers in each window (Appendix A.17). Reach and the support ratio are for the fixed cohort of senders with at least three transfers between June 1 and July 14 (placebo: April 1 to May 14); reach counts a sender with no link as zero, and the support ratio is over senders with at least two links. The largest connected part is computed over all accounts with at least one link in the window. Placebo change: the same measure from April 1–May 14 to May 15–July 1. Net of placebo: the pre-levy to high-rate change relative to the placebo change. Complete transaction records of each provider.

Table A20: New users at Tigo across the levy

	Pre-levy	High rate	Post-cut
New subscriber accounts per day	14,941	14,747	18,040
Same weeks of 2020	10,795	10,367	15,545
Share sending a transfer within 14 days	0.042	0.021	0.023
Share receiving a transfer within 14 days	0.247	0.286	0.280
Subscribers registered in the window (14-day shares)	409,717	512,742	512,352

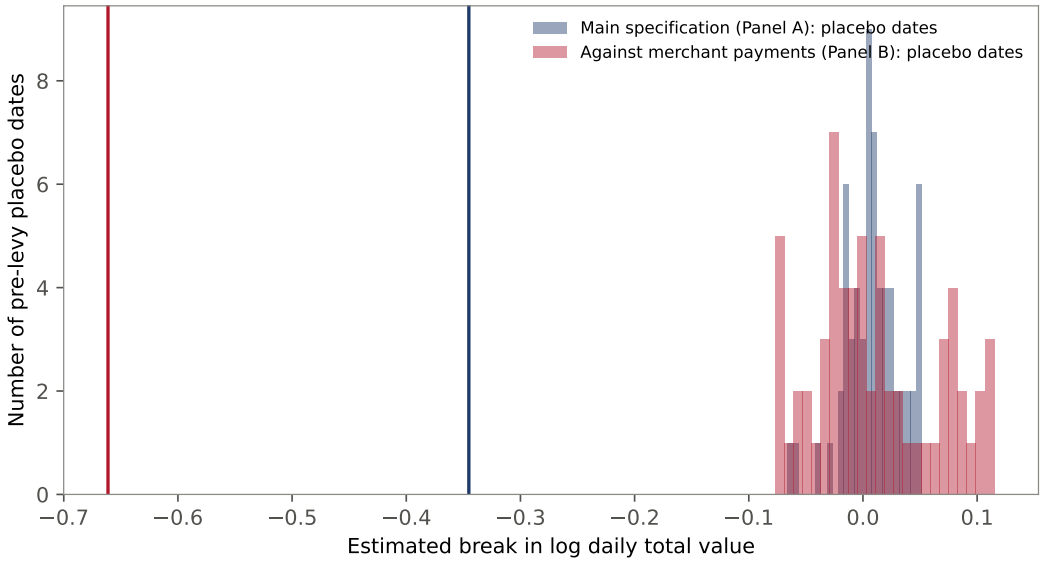
Notes: Subscriber accounts in Tigo's registry, by activation date. Sign-ups per day over June 1–July 14, July 15–August 31 and September 1–October 13, 2021; the 2020 row uses the same weeks 364 days earlier. Sending and receiving shares: on-network person-to-person transfers on the day of registration and the 13 days after, for subscribers registered June 1–30, July 15–August 17 and September 1–29. M-Pesa's registry fails a check on its dates and Airtel's data include no registry (Appendix A.17).

Figure A3: Daily mean transaction value by transaction type around the two policy dates



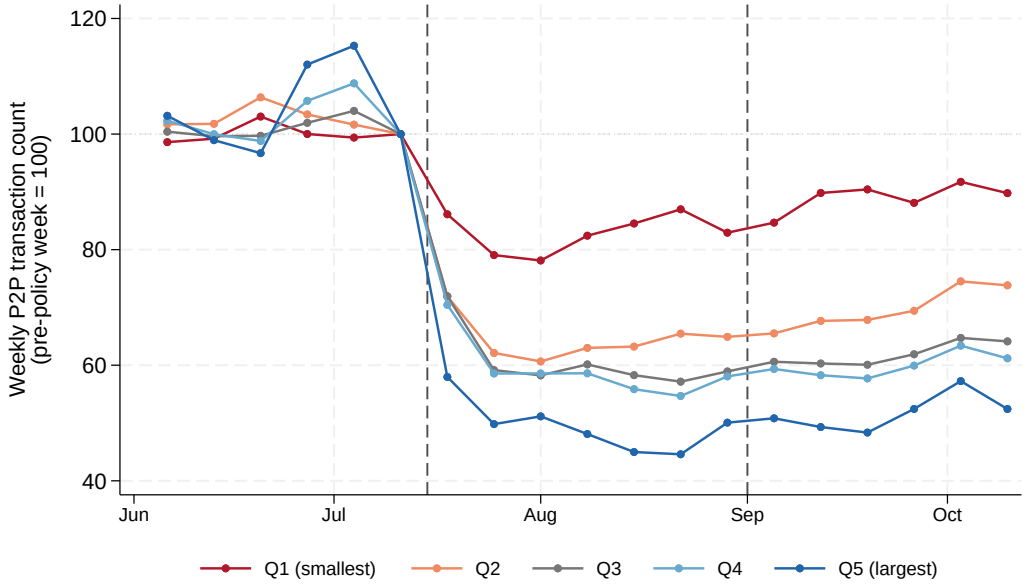
Notes: Weekly average of daily mean transaction value (daily value divided by daily count), normalized to 100 on July 8–14, 2021; series, providers, and window as in Figure 3. Business-to-person payments omit the three M-Pesa artifact days, on which duplication compresses mean value (Appendix A.1). Mean value falls for every series, taxed and untaxed, except merchant payments; Section 4 notes that this outcome is the one that does not clear placebo inference. Dashed lines: July 15 introduction and September 1 cut.

Figure A4: Randomization inference: the levy break against pre-levy placebo dates



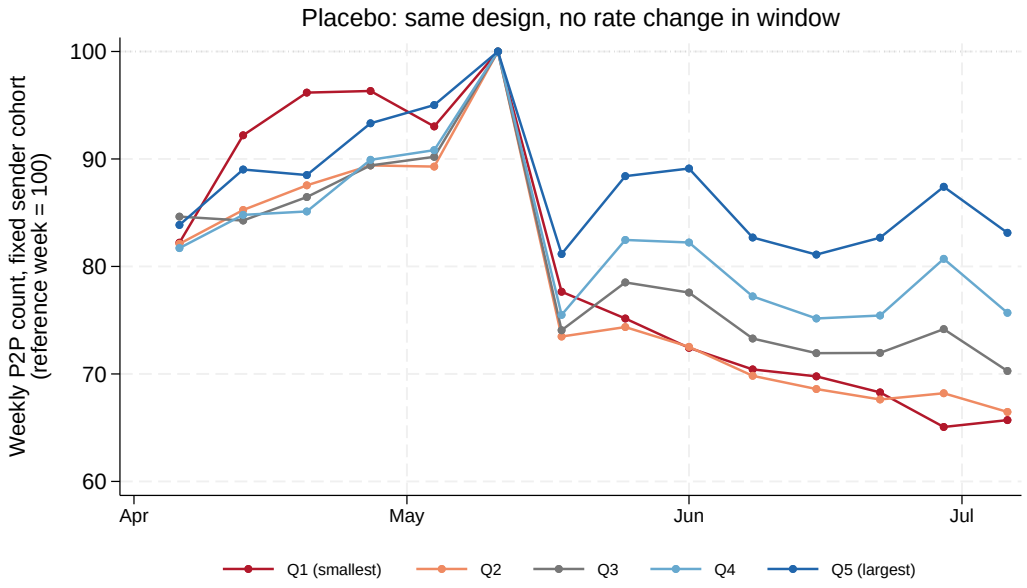
Notes: Distribution of the estimated break in log total transacted value when the intervention date is set to each pre-levy date admitting a full ±21-day window (64 placebo dates, all windows ending on or before July 14, 2021, so none overlaps the levy), estimated on the provider-by-type series with series and provider-by-day fixed effects. Navy: the main specification (Panel A of Table 2); red: taxed transfers against merchant payments (Panel B). Vertical lines mark the estimates at the true July 15 break, window-matched and therefore not identical to the coefficients in Table 2, which are estimated on June 1 to August 31. No placebo attains the true break in either panel.

Figure A5: Weekly person-to-person transaction count by pre-policy size quintile



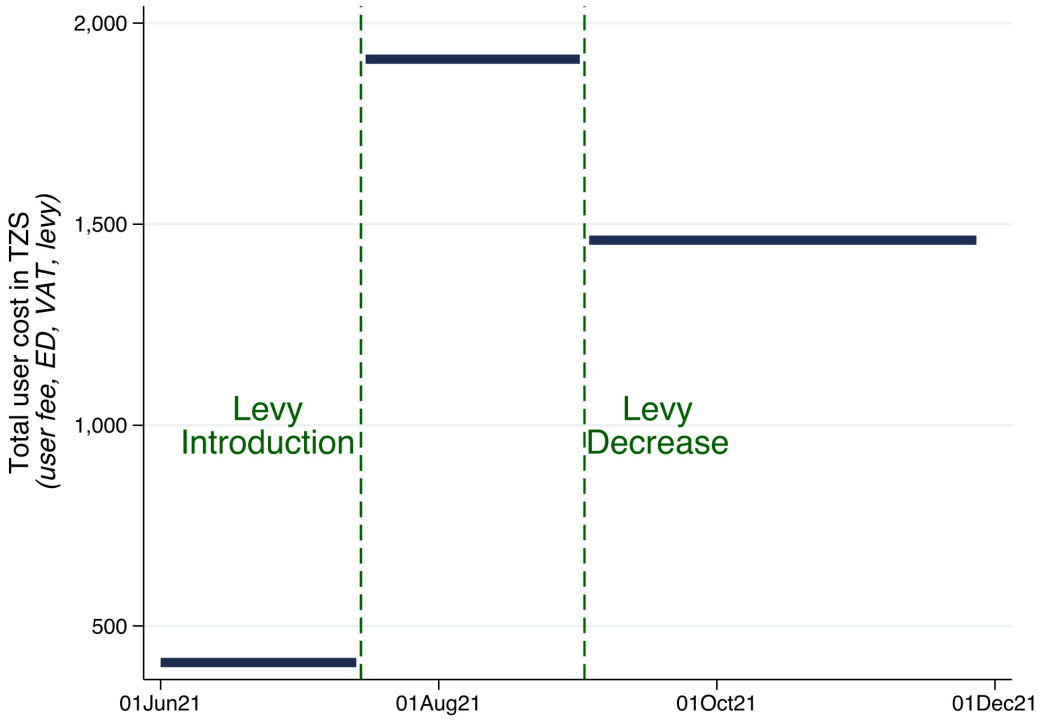
Notes: Weekly count of taxed transfers in each pre-policy size quintile, clean assignment, normalized to 100 on July 8–14, 2021; quintiles are defined on the universe distribution of taxed-transfer values in June 1 to July 14 and applied to all transactions; three providers, June 1 to October 13, non-overlapping seven-day bins; colors run from red (smallest) to blue (largest). Each transaction is assigned by its own value, so transfers that shrink move between bins, depressing counts at the top and inflating them at the bottom; Figure 8 follows a fixed cohort of senders instead, and the gradient there largely disappears (Appendix A.6). Dashed lines: July 15 introduction and September 1 cut.

Figure A6: Placebo: the cohort design applied to a window with no rate change



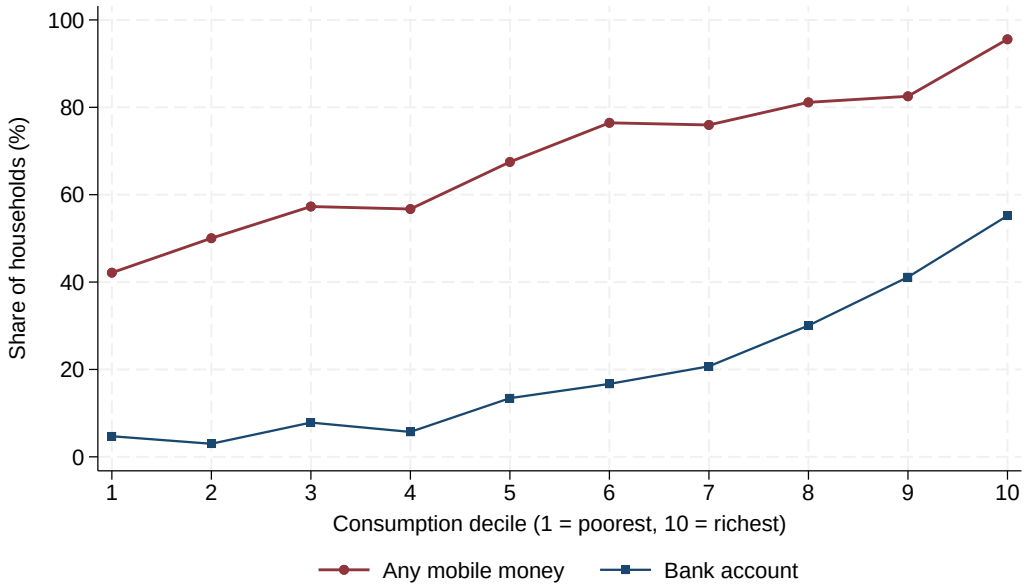
Notes: The design of Figure 8 applied to a window with no rate change: senders with at least three person-to-person transactions between April 1 and May 14, 2021 are sorted into quintiles on their mean transfer over that window and followed through July 14, normalized to 100 on May 8–14 (161,484 senders, after the Tigo duplication correction of Appendix A.1). Mean reversion is present and ordered by quintile, but in the opposite direction to the levy result: the largest quintile lies above the smallest. The level falls to 65–89 with no policy change, which calibrates how much of the post-levy level reflects selection on the sort window.

Figure A7: Total user cost of a representative 40,000 TZS transfer



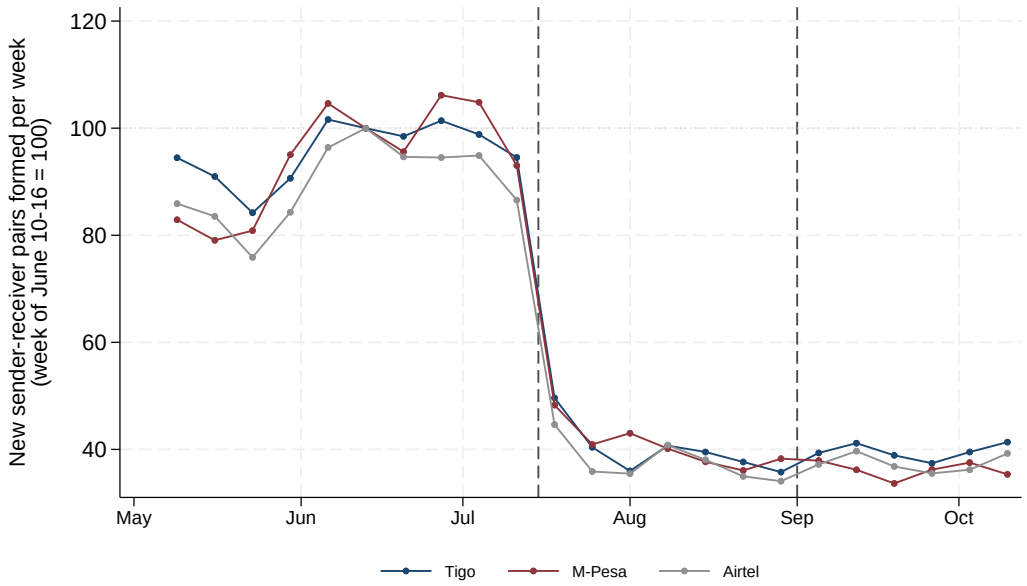
Notes: Total user cost combining the provider's service fee, the 10 percent excise duty, the 18 percent VAT, and the levy, computed for a 40,000 TZS person-to-person transfer. The dashed lines mark the introduction of the levy on July 15, 2021 and the cut on September 1, 2021.

Figure A8: Mobile money use and bank account ownership by consumption decile, Tanzania (TNPS Wave 5, pre-policy)



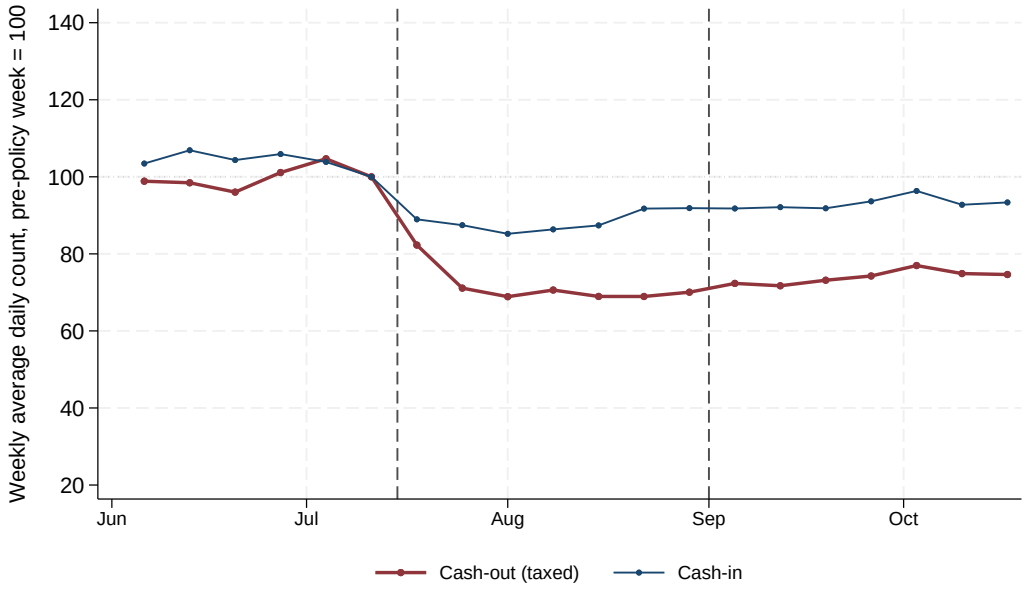
Notes: Share of households reporting any mobile money use (M-Pesa, Airtel Money, or Tigo Pesa) and any bank account, by decile of real consumption per adult equivalent. Source: Tanzania National Panel Survey Wave 5, restricted to households interviewed before July 2021 (N = 2,638). Consumption per adult equivalent is the official LSMS aggregate (expenditure module, deflated to a national reference). Deciles are constructed using cross-sectional sampling weights. Mobile money penetration is already high in the bottom consumption decile (42 percent), while bank account ownership remains scarce until the upper deciles (5 percent in the bottom decile, 55 percent in the top), so bottom-decile households are concentrated on mobile money with no formal banking substitute.

Figure A9: New sender–receiver pairs formed per week, fixed cohorts by provider



Notes: Count of on-network person-to-person sender–receiver pairs appearing for the first time in each week, computed on the complete transaction records of each provider for the same fixed cohorts as Figure 11, normalized to 100 on the week of June 10–16, 2021. A pair counts as new in the week of its first transfer since April 1, 2021; the series begins in mid-May so that pairs already active before the sample window do not appear as new. Seven-day non-overlapping bins; window ends October 13. Definitions in Appendix A.17.

Figure A10: Weekly average daily volume: cash-out vs. cash-in around the levy introduction and rate cut



Notes: Weekly average of the daily count of cash-out (taxed) and cash-in (untaxed) transactions, normalized to 100 on the reference week of July 8–14, 2021; non-overlapping seven-day bins. M-Pesa and Tigo only (the Airtel extract has no cash-in series), so the window runs to October 20, stopping before the M-Pesa business-to-person artifact days of October 23–28 (Appendix A.1). Dashed lines: July 15 introduction and September 1 cut.