

Digital Safety Nets*

Alexander Karaivanov (SFU) Benoît Mojon (BIS)
Luiz Pereira da Silva (Tokyo U.) Robert M Townsend (MIT)[†]

April 2026

Abstract

We analyze how digital technologies can be used to improve safety nets for insurance against idiosyncratic and aggregate income risks, tailored to deal specifically with common market frictions: limited commitment, private information, and transaction costs. We illustrate the gains from incentive-compatible risk-sharing schemes for groups of economic agents such as Thai households or Spanish firms. We assess the best-fitting financial regime for each group and quantify large welfare gains from improved insurance despite the existing obstacles to risk sharing. Our approach can be applied in many contexts to foster financial inclusion, alleviate poverty traps, and complement existing safety net policies and mechanisms in a cost-effective way. We provide blueprints for design and implementation.

Keywords: digital financial platforms and contracts, safety nets, targeted transfers, risk sharing, financial inclusion, limited commitment, private information

JEL codes: G10, G52, G23, D82, D53, O16

*This paper reflects the views of the authors and not necessarily the views of the Bank for International Settlements. We are grateful to Albert Pierres Tejada for excellent research assistance. Townsend acknowledges funding for the Thai data from the Bank of Thailand, the Thailand Research Fund, and the Eunice Kennedy Shriver National Institute of Child Health and Human Development (NICHD) grant R01-HD027638. Karaivanov acknowledges financial support from the Social Sciences and Humanities Research Council of Canada (SSHRC), grant 435-2024-0317. The authors are grateful for helpful discussions and comments from seminar participants at the BIS, the IMF, and from K. Bergant, B. Carton, G. Chen, J. Frost, L. Gambacorta, V. Gaspar, P-O. Gourinchas, B. Nolens, and H. Shin.

[†]Corresponding author: rtownsen@mit.edu; MIT Department of Economics, 50 Memorial Drive, Cambridge, MA 02142, USA

1 Introduction

Income risk, that is, volatility of income flows because of illness, job loss, macroeconomic shocks etc., remains prevalent for households, firms, regions, and countries. Various policies and institutions have emerged to curb these risks, ranging from financial markets that allow saving and borrowing to private and public social insurance and safety nets for unemployment, health risks and old age. These policies originated in today's developed countries and were later expanded to emerging and developing economies (EMDEs).¹ Arguably, they have proved reasonably effective in smoothing household income and expenditure and in stabilizing consumption and productive investment by the self-employed and small and medium-sized enterprises (SMEs). This was especially the case after World War II, from the 1950s to the 1980s in both advanced economies and EMDEs. Empirical evidence indicates that these policies contributed to the decline in poverty and inequality during this period (Morelli et al., 2015 and Alvaredo and Gasparini, 2015). Another important policy development is centered on increasing financial inclusion to the poor, i.e., creating incentives and regulations that improve their access to affordable financial services offered by a range of competing providers (see Princess Máxima, 2010 and CPMI and World Bank, 2016).

However, many countries have seen the trend in income inequality reverse direction since the 1980s (Atkinson and Bourguignon, 2015), with a steady increase in inequality of opportunities, income and wealth, as well as in geographical inequality, e.g., access to public goods.² Importantly, this increase partly reflects the fact that individuals and households with lower skills and/or incomes are subject to a combination of larger idiosyncratic and macroeconomic income risks. In OECD countries, not only is the average unemployment rate of low-skilled workers higher, but it also increases by more in recessions (Pereira da Silva et al., 2022). In EMDEs, where financial markets and social safety nets are less developed, idiosyncratic regional or cyclical income shocks may cause

¹A distinction should be made between broader social policies dating back to the 19th century and significantly expanded and formalized after the US New Deal in the 1930s and the Beveridge Report in Britain in the 1940s, and the social safety nets from the 1990s aimed at mitigating the impact on the poor from implementing macroeconomic structural adjustment policies. Among the latter are conditional cash transfers (CCTs) to households with income below a threshold and satisfying certain conditions, e.g., child school attendance.

²The post-1980 trend of rising income inequality has been addressed by a vast literature that explains it as the combination of skill-biased technological change and the effects of globalization on employment in advanced economies associated with the relocation of significant segments of productive activities. A complementary hypothesis with respect to household income risk is that the broad-based social policies dating back to the mid-20th century "ran out of financing fuel". They had been traditionally embedded into budget processes, acted as automatic stabilizers, and were a pillar of steady growth during the post-World War II period. However, both demographic and unemployment trends stretched them to their financing limits and they have become progressively less effective in covering risks at both the social and individual level. This characteristic is obviously more severe in EMDEs. As a consequence, the volatility and uncertainty of income has increased and even more so for low-skilled and low-income households, see Heathcote et al. (2020) for the U.S. and Pereira da Silva et al. (2022) for cross-country evidence.

the self-employed and workers in informal markets to slide into poverty traps which can persist, in some cases, across generations.³ Risks may be so large that households do not undertake entrepreneurial activities or limit the scale of family enterprises. This deepens economic downturns and lowers potential growth.

Income volatility can be triggered by various events or accidents, some of which idiosyncratic (unemployment, family break-up, poor health, etc.) and some related to the business cycle (Challe et al., 2017; Heathcote et al., 2020; Guvenen et al., 2022; Bilbie et al., 2023). Volatility increases economic insecurity and the risk of falling into poverty despite existing public safety nets. Economic insecurity is empirically more prevalent among the low-income groups, reinforcing the link with the rise in economic inequality.⁴ Therefore, smoothing income risks is a key component of policy considerations.

At the country level, crisis episodes often associated with large capital outflows compound income risks by reducing the fiscal resources devoted to institutional social safety nets. As various crises illustrate, fiscal adjustments have usually been implemented in an across-the-board fashion during an emergency. Likewise, as the euro area sovereign debt crisis has shown, advanced economies had also been increasingly subject to similar (albeit less severe) risks to the financing of their more comprehensive social safety nets. In a nutshell, existing safety nets are either affected by cyclical volatility, or their coverage is often limited, underdeveloped and ineffective, or potentially all these together (see Karaivanov et al., 2026).

Most of the aforementioned policies and institutions were designed before the availability of detailed information about individuals and households, and usually without data frequent enough to capture intra-year income volatility. Therefore, they had to be designed in a broad-based, macroeconomic manner. Even after the implementation of surveys with more granular data such as annual household surveys and the emergence of conditional cash transfers (CCTs), countries tended to rely on broad approaches and frameworks.⁵ In addition, the cost of reliable and timely information about individuals or households in relation to the state of the economy at any given time compounds these problems. It is difficult, if not impossible, to verify an individual's situation facing an adverse outcome, and whether a claim for assistance is justified and true, while addressing income volatility in a timely fashion. The empirical evidence points to the prevalence of

³The World Bank (2019) estimates that only 20% of the poorest people are included in social safety nets in low-income countries.

⁴See Latner (2019), Western et al. (2012) and Rohde et al. (2014).

⁵In many EMDEs, but also in advanced economies, such policy instruments tend to be subject to inertia and political economy pressures to support social policies at various levels of government. Despite efforts to increase social spending, it has been difficult to systematically evaluate their efficiency and improve coordination. Fiscal rigidities in legal systems make revisions complex, even when better procedures could bring efficiency gains. Social policy approaches in many developing countries have improved (for example, the use of conditional cash transfers or some measures implemented during the Covid-19 pandemic and the war in Ukraine) but remain predominantly dependent on financing constraints, centralized, and transfer-based.

gift-based assistance within extended families where trust plays a key role (Kinnan et al., 2024). This supports the conjecture that overcoming incomplete information and limits to commitment without the type of trust that prevails within family circles is a critical obstacle to developing effective contingent and targeted safety nets.

Recent technological developments allow sophisticated interventions that can process and use granular individual, household, and firm data at very low cost⁶ (Goldfarb and Tucker, 2019). Another development is the increased understanding of contracts that induce truthful self-reporting of unobserved individual circumstances. Automated contracts can also deal directly with commitment or moral hazard problems delivering constrained-optimal insurance. Furthermore, messages can be encrypted preserving privacy,⁷ and serve as an input into encoded risk-sharing agreements, aggregated to promptly detect geographic, business cycle or regional shocks. Multilateral smart contracts also allow irrevocable documentation of agreements, irrevocable escrow commitments to social insurance funds used for pooling of shocks, specified rules for the operation of a fund that deals with idiosyncratic and aggregate risk, and commitment to exclusion if the voluntary parts of rules are not followed. More recently blockchain technology and smart digital contracts could together implement incentive-compatible contracts rooted in mechanism design.

To reiterate, such data, understandings, and digital technologies can be used to reduce verification costs, reveal the true state of individuals' situations and provide better incentives – all of which could enable major progress in pooling idiosyncratic risks, regional and country-level shocks and reduce many impediments to the effectiveness of interventions and their rapidity, avoiding time lags and insecurity. This would make the set of available insurance arrangements less incomplete, preserve equal opportunities across economic agents and prevent unnecessary human tragedies, while at the same time providing a more effective stabilization of the business cycle, higher productivity, and improved financial flows.

In terms of policy, we therefore propose a different approach to social safety nets, complementary to existing ones. Instead of an almost exclusive emphasis on social transfers (broad or targeted) with a principal (the state) defining eligibility thresholds and a centralized system of controls and payments, we suggest an emphasis on mitigating income risk, particularly idiosyncratic and breaking the transmission of individual or group risks into more widespread risks. This can include a decentralized implementation process and incentive-compatible digital contracts that collect contributions and pay indemnities, including self-sustained conditional transfers funded with contributions. We thus argue

⁶Atkinson and Bourguignon (2015) recognize that technological progress will fundamentally change social policies. However, they focus on the effects of faster payments rather than the design of policies that include new forms of insurance and financial contracts.

⁷For example, see the discussion in D'Silva et al. (2019) and Nilekani (2018) on the development of privacy-preserving financial infrastructure in India.

that current technology can circumvent obstacles to the development of effective digital safety nets which could help reduce both income inequality and volatility. Complementing existing safety nets would enhance welfare and stabilize the income of low-income households, which in turn implies higher and more stable economic growth. Related to this, insurance allows more flexible terms for loan repayment in credit contracts and lower debt levels, which in turn can be a drag on growth and a source of systemic risk. The economy would benefit from the enormous potential for creativity, productivity and innovation of groups that face differentiated, higher risks and inequality of opportunity.

The new technologies and tools that we describe here make it possible to design various components of social policies and financial inclusion programs with characteristics that could easily complement those introduced in the 20th century. In summary, the key conceptualization is how to address income risk, by illustrating the large welfare gains that improved risk sharing can deliver and offering a proof-of-concept, using insights from mechanism design, of how digital safety net platforms and apps that are incentive-compatible and cost-effective can be coded and implemented to overcome information asymmetries or commitment obstacles to trade.

2 Idiosyncratic Risk and the Need for Improved Safety Nets

In this section, we review the empirical evidence and stress that uninsured idiosyncratic risk can translate into costs material enough to transmit to the family and then to neighboring groups through the loss of income and/or capital. These disruptions increase income volatility and can be caused by various chain events: the loss of solidarity within groups (e.g., within villages due to migration to cities), the rupture of a critical supply chain in the production process, etc. If the cumulative losses are large, they can generate or compound poverty traps and produce additional risks with potential local, regional and macroeconomic effects. Clearly, the prevention of such chains of events is welfare enhancing.

2.1 Empirical evidence for limited insurance

A direct look at outcome data confirms the adverse consequences of income shocks and provides evidence of limited insurance. With complete risk sharing, an individual's consumption should track the consumption of the community as a whole and be uncorrelated with the individual's income fluctuations. A direct look at income and consumption panel data over 10 years from villages in India, Townsend (1994) shows that incomes are not covariate, so that insurance is possible, and although there is a substantial amount of local insurance there remain particularly vulnerable within-village groups, for example, wage laborers, for whom consumption drops in economic downturns.

Richer nationally-representative Thai data show how parts of society can be more exposed to economic fluctuations. Aggregating biannual cross-sectional surveys into a pseudo-survey of regional and occupation cohorts helps track income and consumption data across social groups. These data show that entrepreneurs and people in and around Bangkok who have arguably left village-based social protection bear the ups and downs of fluctuations and so would benefit from enhanced insurance (Townsend, 2016).

Additional data include the choice of occupations, sectors of chosen businesses, and the rates of return of those businesses. Specifically, if there existed perfect mutual insurance within a village, idiosyncratic fluctuations in returns would be entirely smoothed out. Higher expected returns would not be needed to compensate for risks not borne, consistent with the Thai village monthly panel data (Samphantharak and Townsend, 2018). But for common village-level aggregate risk, the higher the co-movement of the return on the households' chosen technologies with the village aggregate return, the higher is that household's average return in the data, to compensate for the aggregate risk, meaning that the latter is relatively uninsured. However, as village aggregate shocks are not covariate across villages, such shocks are potentially insurable when pooling across villages, but this is not currently happening. By extension, the impact of country-level income shocks could be reduced by risk sharing.

A related literature on insurance, or the lack thereof, examines capital flows and remittances. In theory, if a household were holding a portfolio of real and financial assets, we should see transfers and incoming capital flows that are inversely related to the household's underlying income. The same logic holds for regions, with a distinction between income from local production and non-local income from remittances, transfers, and dividends. Again, there should be outflows when a region or sector is doing well and inflows when it is doing poorly. But the evidence that capital flows remain procyclical is overwhelming (e.g., Forbes and Warnock, 2012; BIS, 2017; Koepke, 2019; and Aldasoro et al., 2023). Related to this, with diversified portfolios savings and investment should be based on productivity only. Investment should not drop with temporary income shortfalls. These benchmarks have been used extensively in studies of inter-regional and international risk sharing (e.g., Feldstein and Horioka, 1980) showing that regions or countries can significantly improve risk sharing compared with the status quo.⁸ For example, China does poorly in inter-regional flows compared to advanced economies such as the United States. Countries in the EU, a monetary but not fiscal union, also do less well in sharing unemployment-related risks relative to what is achieved in the United States (e.g., see Claveres and Stráský, 2018 or Dolls, 2024).

⁸One could interpret the founding mission of the Bretton Woods institutions (IMF, World Bank), many United Nations specialized agencies (WHO), and other regional institutions (e.g., ESM, AMRO) as helping foster international risk sharing. However, many consider the global economy as still falling short of having an effective global safety net (e.g., Carstens, 2021; Adrian et al., 2022).

These findings suggest that existing markets and risk-sharing institutions are only partially adequate. The degree to which intermediaries contribute to insurance in Thailand has been analyzed by Alem and Townsend (2014) using consumption and financial transaction data. Although an agricultural bank (BAAC) does help in consumption-smoothing, other financial institutions contribute little and virtually no institution covers shocks that adversely impact investment. The BAAC seems to have been less than successful in its introduction of a rainfall-indexed crop insurance program, arguably because the meteorological data used were not sufficiently granular and there was too much basis risk. Related research estimating the best-fitting financial and information settings using consumption and investment data points to substantial heterogeneity in outcomes linked to incomplete markets and obstacles to risk sharing (Karaivanov and Townsend, 2014).

2.2 Transmission to family and larger groups

Idiosyncratic shocks can have ripple effects and spread to larger groups at the local, regional, or even national level. The idea of “contagion” is often used in finance or epidemiology but it can be also translated into socio-economic terms and applied to the cascading, non-linear effects of uninsured risks. In addition, a systemic shock such as a pandemic can disproportionately affect a segment of the population, especially the most vulnerable and the poor. Various partially overlapping mechanisms can be identified.

A granular look at Thai villages is revealing. Households without other relatives in the village are generally not engaged in gift-giving insurance networks (Kinnan et al., 2023). When these households experience large unexpected and uninsured expenses from illness, they cut back on material input purchases and hired labor. This creates amplified shocks that reverberate through supply chains and local labor markets, as revenues diminish. Consumption drops for those affected. Furthermore, broken transaction links persist over time, i.e., they are not easily replaced by alternative points of sale or employment, inflicting long-term structural damage and resulting in persistently lower incomes.

Another transmission mechanism works through the risk profile of groups that are similar enough to create a “domino effect” once a member is hit by a shock. Durlauf (1994) shows that local spillover effects could be interacting with the income-based stratification of neighborhoods to transmit parental economic status from generation to generation. Bird (2007) finds that contagion from individual situations to larger groups depends to a large extent on the group’s sharing of common characteristics such as the absence of nurturing, lack of investment in human capital and a dearth of opportunities, all of which make people much more susceptible to falling into poverty traps when they are exposed to negative shocks (see also Chetty et al., 2014).

2.3 Poverty traps

Additional justification for the need for improved risk sharing to reduce income volatility and its consequences is the extensive literature on poverty traps: individuals or groups can fall into self-reinforcing situations that cause and maintain poverty. Once a household is in a poverty trap, it persists absent outside intervention, in some cases even across generations, because the individuals have limited or no resources and their capital (physical or human) has low return. These self-reinforcing disadvantages make it very hard to escape poverty (Bowles et al., 2006; Kraay and McKenzie, 2014; Ghatak, 2015; Balboni et al., 2021; Banerjee and Duflo, 2012).

The literature documents several mechanisms that can lead to a poverty trap.⁹ Barrientos (2007) emphasizes the role of non-linear income dynamics, low asset endowments, and lack of access to credit. Thanh (2005) examines the vicious circle of poverty and ill health. Idiosyncratic events, such as an accident or injury, could have severe and lasting consequences for family income producing a poverty trap. Liao et al. (2022) show that, in the absence of health insurance, poor and vulnerable people tend to over-stress their bodies to earn more money to invest in education. Rosenzweig and Binswanger (1993) and Carter (1997) show how risk preferences can induce poor rural households lacking access to credit and insurance to choose low-return livelihoods as self-insurance against weather risk, but these choices can trap them in chronic poverty.

A second mechanism is occupational poverty traps, whereby the combination of borrowing constraints and lumpy production technologies implies that poor individuals who start businesses that are too small can be trapped into earning returns of no more than the subsistence level (Banerjee and Newman, 1991). This provides justification for microfinance loans designed to allow households to lift themselves out of poverty by buying lumpy capital assets necessary to operate a business.

Another mechanism is the concentration of risks in specific geographic areas. Jalan and Ravallion (2002) define a “geographic poverty trap” as occurring when the characteristics of a region are such that a household’s consumption cannot rise over time, while an otherwise identical household living in a different better endowed area would have a rising standard of living. This can be compounded by poor institutions at the local or central levels of government (Acemoglu and Robinson, 2012). Poor households in remote rural regions have a reduced set of available production technologies, which means that the choice between lower-income and higher-income outcomes can be a more difficult discrete step.

⁹Barrett et al. (2016) survey the possible mechanisms, such as poor nutrition and mental and physical health, endogenous behavioral patterns (e.g., risk or time preferences), poorly functioning capital markets, large uninsured risk exposure, and weak natural resource governance institutions.

2.4 Financial inclusion

For a large part of the world population, neither self-insurance via savings nor mutual insurance are available. This is particularly salient for farmers in EMDEs, who face the triple challenge of being too isolated, too small-scale for traditional financial intermediaries, and too risky to access financial services (Benami and Carter, 2020). As mentioned before, many EMDEs have put in place some version of a conditional cash transfers program (CCT). In other emergency circumstances (e.g., Covid-19), exceptionally large state-contingent transfers have been made. However, the latter programs were often implemented under the urgency of a crisis. Consequently, they tend to be poorly targeted and frequently transfer funds with lags that might be sufficient to put an SME out of business. In this respect, digital platforms and smart contracts offer considerable potential for improvement. According to the World Bank’s Global Findex Database (2021), 1.4 billion adults remain unbanked. The same data also reveal that only 40% of adults in EMDEs are saving and, among those who borrow money, nearly half are borrowing from their family, 32% consider it would be very difficult to access emergency money, while 9% consider it would be impossible. These facts suggest the great potential that developing financial inclusion could imply for protecting the poorest in society from income risks and poverty traps.

Recent years have seen major progress in financial inclusion thanks to the spreading of smartphones, financial apps and the development of large-scale public policies that foster digital infrastructure (e.g., Aadhaar and UPI in India, see D’Silva et al., 2019). The main novelty is that the marginal cost of giving access to these apps is negligible, opening the way to an untapped means of risk sharing and income smoothing for hundreds of millions. For instance, in sub-Saharan Africa, saving through a mobile phone app, which did not exist back in 2017, is now used by 20% of adults in countries such as Ghana, Kenya, Senegal, Uganda and Zambia (World Bank, 2021).¹⁰ This progress is highly promising. Indeed, among adults that have a digital or traditional account, more than 50% used their digital account to save money and 20% borrowed money. These proportion of savers and borrowers are much smaller among the unbanked.

Indeed, in an increasing number of countries (e.g., Argentina, Brazil, China, India and Indonesia), digital apps for payments have been a starting point for the dramatic spread of new technology among the population. Regarding low-cost fast payments for small amounts (the vast majority of financial services demanded by the poor), Brazil’s Pix reached millions of users, accounting for 75% of the adult population within a year. In several instances, for payment providers, digital apps and data are coupled not only with savings and investment instruments, but also with risk sharing through flexible contingent

¹⁰Another striking example is how M-PESA facilitated risk sharing in Kenya. Jack and Suri (2014) show that M-PESA users were able to shield their consumption from income shocks while the consumption of non-users dropped by 7%. Users tapped remittances to an extent not available to non-users.

credit arrangements, and in the case of China, health insurance reaching 100 million. Brazil's ongoing work on a Digital Real features programmability with delegation to the private sector. The progress of financial inclusion in India is also remarkable (D'Silva et al. (2019)). Those skeptical of voluntary participation in risk-sharing digital platforms should be reassured. Several countries are moving forward. For example, the Bank of Thailand's policies and strategies for a sustainable digital economy include open infrastructure, open data and free competition. Nevertheless, innovation and technological updates are not universal.

The spread of digital financial apps is not a panacea. Although the main benefits are highly intuitive, with transaction costs reduced considerably, digital finance apps do not per se overcome issues of trust, asymmetric information and limited commitment. In the case of M-PESA users, the main source of risk sharing was family members. Likewise, Karaivanov and Townsend (2014) show that expenditure-smoothing in Thai villages relies on family networks. This type of risk sharing is thus not accessible to large parts of society, notably those without family ties or those who left their village-based social networks when moving to cities.

3 Obstacles to Risk Sharing

The previous section argues the benefits of smoothing disposable income with improved risk sharing and safety nets. However, there are well-known obstacles that lie in the way of ensuring participation and cooperation in such schemes. We next discuss and distinguish the main obstacles to risk sharing. Namely, limited commitment with lack of trust is different from private information compounded by limited communication. Likewise, new tools that enable commitment and self-reporting with encryption are available, with smart contracts which can be implemented on a digital platform with escrow accounts and multilateral programmed code.

3.1 Limited commitment and lack of trust

In practice, agents who participate in an income insurance scheme are happy to receive an indemnity. However, the payment of premia out of high-income states is not guaranteed. It is well known that in a static single-period insurance contract there is no incentive to pay the premium. In informal systems such as kinship groups there can be social sanctions or penalties that help ensure participation, but this mechanism is absent when dealing with strangers (or with other countries in the case of international risk sharing), despite the known advantages of wider insurance pools.

Table 1. Ad hoc risk sharing with an escrow account (in thousands of Thai baht)

Period	Gross income		Escrow account		Disposable income	
	Agent A	Agent B	Agent A	Agent B	Agent A	Agent B
t	54	42	13.5	10.5	40.5	31.5
$t+1$	12	90			36	90

Note: in this ad hoc scheme each of the agents puts in escrow one quarter of their income in the first period and understands that if their income is low in the second period, they will receive the escrowed amount back and, if the income of the other agent is large, also an indemnity equal to the amount put in escrow by the other agent.

3.1.1 Use of escrow

One solution to the problem of limited commitment is to require the premium to be paid in advance. This can be illustrated with a two-period contract. As an approximation, in Table 1 we consider a two-agent model in which both agents have similar income incomes in the first period (t) but their relative positions in the next period ($t+1$) are random with one agent having much larger gross income than the other. Each agent would like to plan to receive an insurance indemnity in the second period upon a low realized income, at the cost of paying a premium upon a high realized income. Full insurance, i.e., splitting the difference in incomes, would be ideal. However, if there is a limited commitment problem, the high-income agent would not pay.

One possible solution consists of paying the insurance premium in advance regardless of income at that time. Imagine the following ad hoc risk-sharing scheme in the two-period model of Table 1 where the agents put in escrow one quarter of their realized period- t income to hedge against low income in period $t + 1$. Here we consider the risk-sharing scheme as a platform, not only with ledgers as in social insurance account balances, but also using new digital technologies which allow agents, including strangers, to participate. An individual agent’s contribution to the insurance fund is put in escrow. Specifically, a trusted third party agrees to hold money contributions in escrow while minting a token that is programmable via smart contract code. The savings function is thus automated. Likewise, transfers as indemnities received by the platform participants are executed according to risk-sharing rules embodied in the code. Conversion of tokens back into money from the trusted third party is also subject to programmed rules. Some tokens may be required to remain in escrow, other amounts can be disbursed, according to pre-programmed and agreed rules. The key point is that escrow accounts are under the control of the code and not individual discretion.

To continue with the two-period example, suppose it is agreed that in period $t+1$, the agent with the lower income receives back their own escrow deposit from period t plus, as indemnity, the money put in escrow by the other agent in period t . In Table 1, in the columns labeled “Disposable income”, we illustrate the agents’ payoffs if they participated

in such insurance mechanism. As the table shows, the mechanism is highly effective to insure against very low-income draws. In period $t + 1$, agent A would have disposable income of 36,000 instead of 12,000. With preferences that are sufficiently risk-averse, it is highly likely that individuals would be willing to participate in schemes that allow them to hedge effectively against low realizations of future income. Notably, concavity implies that the marginal utility gain from receiving a transfer upon realization of a low income is higher than the utility cost of paying an insurance premium when receiving a high income. Applying a log utility function to disposable income makes this point obvious.

Risk aversion is therefore critical. It helps explain why agents would pay an insurance premium in high-income states in return for receipt of an indemnity in low-income states. The gain from such a mutualized insurance pool is higher overall welfare and efficiency.

Such arrangements do face the problem that the high-income agent who is supposed to contribute might prefer to withdraw from the arrangement instead. But, as such high incomes do not occur frequently, after controlling for aggregates, one should expect to see full risk sharing over period of time, and then a reset (partial risk sharing) when the income for an agent is very high. Such an agent is induced to stay in the arrangement by getting higher future consumption on average.

The use of escrow can mitigate this problem. However, full insurance, while preferable in theory is unlikely to be attainable even with escrow. Indeed, full insurance could require a large premium that in the escrow scheme would be subtracted ex-ante from gross income in period t . When income is low prepaying a given premium amount is less likely to be acceptable to the agent. Such circumstances will limit insurance. We gain the intuition that a voluntary prepayment of collateral would depend on the income process.

More generally, interacting over more than two periods, it remains the case that if current disposable income after paying a premium is below some threshold, then an agent would not participate in insurance for the next period, as the utility loss is too painful despite the potential future gain. One can imagine various modifications of the scheme, with the insurance premium scaled back if current realized income is relatively low or, alternatively, using accumulated past saved premia. But unlike in a saving-only setting, here the agents are committed to future payouts from others in low-income states (see Section 3.3). Karaivanov et al. (2026) and Dolls (2024) illustrate such risk-sharing schemes as applied to EU unemployment re-insurance.

3.1.2 Exclusion and the benefits of public infrastructure

Rather than requiring collateral or escrow to prevent renegeing, another possibility is to exclude agents from future risk sharing upon failure to follow the rules, i.e., not following the stipulation to pay in. Intuitively, for an arbitrarily large number of periods such exclusion implies a large welfare loss to the agent from reduced consumption smoothing,

to be traded off against keeping current high income. When the former is bigger than the latter the threat of exclusion from the scheme provides an effective incentive for the agents to abide by the rules, i.e., they would choose to pay the required premium when income is high.¹¹

A real-world example of such incentive mechanism is the potential loss from being excluded from trading platforms. Alibaba's mutual health insurance was successful in part because failure to pay the annual health premium exposed the household to losing access to all Alipay services including the popular Alipay e-payments service.

Although private sector entities can enforce such sanctions, there is a downside, as large providers may potentially act as monopolists and extract rents, in which case much of the benefit from risk sharing accrues not to the public but to the provider.¹² Not all private entities do this in practice but the risk exists. One option is to incorporate private entities into a public programmable platform infrastructure, with governance and regulatory stipulations. Another possibility is that the state, communities, or non-profit organizations could provide their own not-for-profit insurance platform. Such a platform could implement computer code that fits the insurance schemes to the income and preference characteristics of a specific group of economic agents. Importantly, trust in the insurance scheme may be reinforced if the non-profit objectives are clearly spelled out.

Participation would be agreed to voluntarily by all in advance, when convinced via simulation or pilot, or if vouched for by a trusted third party that the code executes as intended, i.e., implements the agreed sharing rule and state-contingent programmed premia and indemnities. The platform itself, however, does not require a trusted third party.

Each participant would have an account on a common ledger which would be used to pay out premia in advance or to receive inflows as indemnities. Prepaid deposits are locked, then tokenized into e-objects, and thus sequestered for premia only, as per the agreed rules. Contributions could also be made, as in an ex-post mutual fund, and these could henceforth be instantaneous at the time income is realized. An additional rule could be the exclusion of an agent when due ex-post premia are not paid. This requires a public commitment to such governance. Thus, rules for governance are an important aspect of digital platforms. Finally, a hybrid scheme can use both prepaid collateral and sanctions hence requiring less collateral.

¹¹While exclusion can be very effective as an incentive device, it can be time-inconsistent, in the sense that there exist gains from trade and the parties may wish to re-negotiate ex post. Competition among insurance providers may also undermine the ability to exclude agents.

¹²Theoretically, exclusivity could be a beneficial feature in information- or commitment-constrained settings but this needs to be balanced against the usual welfare gains from competition.

3.2 Private information

Although digital data allow granular targeting, as noted earlier the exact situation, e.g., the income or savings of an agent, may nevertheless remain unknown to others as an unobserved type or unobserved choice/action of the agent. That is, an agent may not want to share some private information for fear of expropriation by the state or by those in the local community. The mechanism-design remedy is to use hybrid risk-sharing contracts with incentives to induce truthful self-reporting, i.e., honest messages about unobserved underlying states or actions. Furthermore, these messages and underlying balances on ledgers would be encrypted so as to conceal them from others, although the code can operate on the encrypted messages just as it would on the original unencrypted ones using homomorphic encryption. That is, the original underlying message space is one-to-one with the encrypted space but the encrypted messages are impossible to decipher. It is not necessary to have a (potentially not trusted) intermediary as the interface between messages and outcomes. All is done in the code, even though the inputs into the code are encrypted. Furthermore, the code can run offline, not on a blockchain, thus avoiding costly validation once the objectives and performance targets are agreed. Private sector entities can utilize code with encryption as an alternative institutional implementation so that they too have no access to the underlying private information.

Again, a simple two-period model with high and low incomes helps to illustrate how this would work (see Table 2). Suppose, as a benchmark case, that each agent contracts not with others individually but through a risk-neutral intermediary/platform absorbing risk or a large mutual fund that pools funds. As a starting point, imagine that at period t those with high incomes can save and those with low incomes can borrow individually. With underlying period- t income (high or low), unobserved, it is up to the agent to decide whether and how much to borrow from or lend to the platform. A market-determined interest rate can be used. The agents take the borrowing or lending position in the first period t voluntarily. It is also assumed that transfers at period $t + 1$ are determined entirely by the agents' positions in period t . In the two-period model shown in Table 2 there is no additional insurance possible at $t + 1$. Limited commitment in repayment of a loan at $t + 1$ is not considered to be an obstacle here, as in this section we focus instead on private information but, as earlier, one can use collateral and exclusion, as detailed below.

Such borrowing or lending is, however, not the best information-constrained arrangement. It is more efficient to move partway toward full insurance, to an ad hoc risk-sharing rule. Townsend (1982) shows that a constrained-optimal hybrid arrangement, essentially combining risk sharing with borrowing and lending, can yield higher welfare to the agents than strict borrowing and lending. The essential idea is that when an agent claims low income the incoming indemnity transfer is larger than what a loan would be in the pure borrowing/lending scheme. Likewise, the amount invested as premium when high income

Table 2. Constrained-optimal combination of borrowing and lending and risk sharing as in Townsend (1982)

Period	Gross income		Borrowing/lending (B/L)		Disposable income	
	Agent A	Agent B	Agent A	Agent B	Agent A	Agent B
t	90	30	-15	+15	75	45
t+1	48	18	+15	-15	63	3
Combining B/L and risk sharing						
t	90	30	-15	+15	75	45
t+1	48	18	+7.5	-7.5	55.5	10.5

Note: in this ad hoc scheme the agents lend one another one quarter of their income difference in period t and repay it in period $t+1$, assuming a zero interest rate. The repayment in period $t+1$ can be reduced, to take into account the large gross income difference in that period. In the case combining borrowing/lending and risk sharing, we show the effects of reducing the reimbursement by one quarter of the income difference in period $t+1$.

is claimed would be higher than the savings in the pure borrowing/lending scheme. The loan is repaid in the second period but less compared to in the case of pure borrowing, and the return on savings is lower than in the case of pure investing. The scheme is designed to be incentive-compatible so that messages about the gross incomes received are truthful. The intuition is that period-by-period full risk sharing would be best but it is not attainable, so transfers are attenuated. To be incentive-compatible, what happens in one period is tied to future periods. The agent with private information internalizes this trade-off, which would itself be a distortion relative to the first-best world, and so the tool is used sparingly. Note, in particular, that the history of incomes is an inherent part of the platform in this example, and it is an input into the *if/then* statements of the digital code. Which branch is applicable depends on the past history of messages which are stored on the platform in an encrypted database. There is a common underlying state of what is in each of the participants' accounts, consistent with their history on the platform but accounts can be partitioned and kept private using both homomorphic encryption and multi-party computation.

This logic can be extended to the case of two agents and no risk-neutral principal. The optimal achievable insurance contract without information problems, i.e., with public information on incomes, would predict that the agents insure one another by agreeing that the repayment of loans should depend on the income flow in the period when the loan is repaid. To fix ideas, assume independent and identically distributed income shocks between the two agents where the agent who draws a high income at time t is lending to the agent who draws a low income at t . It is optimal in ex-ante terms that if the lender in period t once again has a higher income than the borrower in period $t + 1$, then the repayment of the loan contracted at t would be reduced by a factor that depends on the preferences and the amount of income risk for the two agents.

Income information here is private but as before each agent can be given an incentive nevertheless to announce truthfully in period t if their current message constrains future outcomes at $t + 1$, as is in the constrained-optimal hybrid contract with a risk-neutral principal described above. An ad hoc example for two agents is displayed in Table 2.

The main intuition in Townsend (1982) is that adding risk sharing to borrowing and lending can be Pareto-improving in ex-ante terms. Table 2 provides the intuition. While in period t it makes sense that agent B borrows from agent A, paying back this loan in period $t + 1$ could be very costly in utility terms for agent B because her gross income of 18,000 is very low. Hence, the possibility of moderating the repayment as a function of realized income in period $t + 1$ is optimal.

It might appear that gains from trade for a given agent are still limited in the second period $t + 1$, but now, with two agents, the outcomes can depend on what the *other* agent announced at t , which can be kept private from others and from the platform via encryption. Nevertheless, appropriate randomization of the actual transfer that each person receives is still possible, so that any individual history can be concealed. It might be tempting for an agent to claim low income if that agent knew the overall state including the history of all agents, in order to claim a (larger) indemnity or pay a lower premium. However, the same message may incur a heavy penalty for some realizations of the history of what the other agent claimed at the same period t . The intuition is that insurance is limited when the history is known, akin to letting an agent buy insurance after knowing that there will be a loss; likewise, insurance is possible when keeping some of the history concealed.

A powerful encryption tool is multi-party computation, in which individual states can be concealed while added up and revealed as an aggregate to everyone. This aggregate can be used to execute an insurance option with others. For example, several agents in a village can self-report an unobserved component of income. These components are added up, and if low in the aggregate, be used as a claim for an incoming village-level indemnity from other villages, or if high, the village would pay a premium.

3.3 A two-period illustration

We compare and illustrate the impact of financial, information, and commitment constraints on the ability of agents to smooth consumption and share income risk via a two-period computed numerical example displayed in Table 3. Suppose income in each period can take two possible values, either $y_H=5$ (high income) or $y_L=1$ (low income), each with probability one half. For simplicity, suppose the agents have log utility of consumption, $\ln(c)$ and there is no discounting across the periods. The agents maximize their expected utility of consumption for the two periods subject to different constraints corresponding to the financial regime they are in.

We consider a broad range of financial regimes: *autarky* (no smoothing, consumption equals income in each period and state), *saving only* (an agent can only smooth via savings in a non-contingent asset), *borrowing and saving* (an agent can save or borrow in a non-contingent asset), *hidden income* (agents interact with an insurance platform subject to an information constraint, namely the agent’s income realization is unobserved by the platform), *limited commitment* (the agents interact with a platform subject to a commitment constraint, namely they cannot commit not to renege on a payment and switch to autarky if this is in their interest), *hidden income + limited commitment* (both the information and commitment constraints are present), and *full insurance* (the first-best arrangement allowing unconstrained state-contingent risk-sharing transfers). The agents, in the saving only or borrowing and saving regimes, or the platform, in the other non-autarky regimes, are assumed to have access to a non-contingent asset with period return (gross interest rate) equal to 1.

Table 3. Optimal risk sharing with different financial, information or commitment constraints (two-period example)

Transfers	Period 1		Period 2				Expected utility
	τ_H	τ_L	τ_{HH}	τ_{HL}	τ_{LH}	τ_{LL}	
Autarky	0	0	0	0	0	0	1.61
Full insurance (first best)	-2	2	-2	2	-2	2	2.20
Saving only	-1.4	0	1.4	1.4	0	0	1.73
Saving and borrowing	-1.4	0.3	1.4	1.4	-0.3	-0.3	1.75
Hidden income	-1.8	0.9	1.1	1.1	-0.2	-0.2	1.87
Limited commitment	-2.1	0.7	0	1.9	0	0.7	2.03
Hidden income+limited commitment	-1.7	0.5	1.2	1.2	0	0	1.85

Note: the table shows the optimal transfers τ_{ij} , defined as the difference between consumption and income in state ij where i denotes first-period income (H, high or L, low) and j denotes second-period income. The computation assumes no discounting, zero interest on saving or borrowing, and log utility.

Table 3 shows the optimal *transfers* (the difference between consumption and income in any given period and income state) computed for each of the different financial settings. Transfers are signed, so that a positive value means incoming to the household and negative means out-going from the household. Saving is out-going, in the sense of not being available for current consumption, so the “transfer” in that case (e.g., state H) is negative. In many of the regimes the transfer value shown includes both savings and insurance premium.

In the autarky regime consumption always equals income (no smoothing is possible) thus the reported transfers in all periods and states are zero. In contrast, full insurance calls for a premium of -2 (half of the difference between the high and the low incomes) whenever income is high and an indemnity of $+2$ whenever income is low so that consumption equals the expected income (3) in all time periods and states.

In the saving-only regime, the agent saves when first-period income is high (transfer

-1.4 in state H) and receives back the same saved amount (transfer $+1.4$) regardless of the income state in the second period (states HH and HL). That is, the first and second period transfers are equal in absolute value but have opposite signs. If first-period income is instead low (state L), the agent is unable to smooth consumption and therefore the transfers in states L, LH and LL equal 0.

In the saving and borrowing regime the agent saves (transfer -1.4) when first-period income is high (state H), which is then received back in period 2 states HH and HL, as in the saving-only regime. However, unlike before, the agent is now able to borrow (transfer/loan $+0.3$) when first-period income is low (state L), which is then repaid (transfer -0.3) in second-period states LH and LL. In the latter case the agent is vulnerable and consumes less than y_L . Commitment to repay the loan is assumed; if commitment were absent, the outcome would be the same as in the saving-only regime.

In the hidden income regime, the truth-telling constraint implies that the second-period transfers cannot depend on the second-period income – otherwise an agent would claim the income level yielding a larger transfer. However, these transfer can and optimally do depend on the first-period income, i.e., there is history dependence, see Townsend (1982). More smoothing is attained than in the saving-only and the saving and borrowing regimes. Note the opposite signs of first- and second-period transfers depending on the income state, although the magnitudes are attenuated in the second period. Specifically, the agent receives a larger indemnity, a transfer of $+0.9$, if first-period income is low (partial insurance), followed by paying a premium (transfer of -0.2) in states LH and LL. If instead the first-period income is high (state H), the agent first pays a premium (transfer of -1.8) and then receives back $+1.1$ in second-period states HH and HL. Note that in the hidden income regime the platform can enforce payment of premia in states LH and LL since we assume that there is no commitment problem. Having to pay a premium in the low-income state LL is hurtful for the agent.

In the limited commitment regime, if income is high there is a larger optimal first-period premium (-2.1) compared to in full insurance and a partial indemnity ($+0.7$) if the first-period income is low. In the second period, the risk-sharing platform cannot enforce paying a premium because of the commitment problem (see states HH and LH with transfer 0), however, the agent still optimally receives an indemnity (financed by the large first-period state H premium) if her second-period income is low, including in the state LL. The limited commitment regime thus deals with vulnerability at low-income draws in both periods, unlike any of the other constrained regimes.

In the combined “hidden income + limited commitment” regime, the optimal contract has to deal with both asymmetric information and commitment problems at the same time. Naturally, this results in lower expected utility for the agent than when facing each obstacle on its own, however, note that the constrained optimal risk-sharing contract still achieves more consumption smoothing and higher expected utility than the saving only

and the saving and borrowing regimes.

Finally, moral hazard as an obstacle to risk sharing can be included too, by extending the model to allow for an unobserved action (e.g., effort) affecting the income probabilities. As a result, partial insurance obtains which reduces effort (the crux of the moral hazard problem) but this trade-off is optimized and there is still a gain in expected utility and consumption-smoothing relative to autarky or saving only.

4 Assessing and Implementing Improved Risk-sharing Schemes

4.1 An algorithm for estimating and assessing obstacles to risk sharing

The approach of Karaivanov and Townsend (2014), hereafter KT, allows a quantitative assessment of the gains from participating in an insurance platform and describes the mechanisms through which such a platform can overcome, or at least greatly mitigate, obstacles due to limited commitment or asymmetric information. Specifically, the approach can determine whether existing financial arrangements on the ground are exogenously incomplete, as in buffer-stock savings or borrowing/lending schemes, or endogenously incomplete, constrained by various explicit obstacles to trade. The approach uses structural estimation to compare data against numerically computed solutions of various financial market settings, with the end goal of determining the precise type of friction (e.g., borrowing constraints, information or commitment obstacles) that best matches the data. Identifying the underlying financial setting (obstacle to trade) and model parameters allows for informed and more reliable evaluation of counterfactuals and their associated welfare gains, for example, a switch to a less constrained financial environment.

The data used in the KT approach can be cross-sectional, time-series, or panel, as available. Variables at the household level can include consumption or expenditure, income, capital and/or investment. Karaivanov and Townsend (2014) analyzed communities of farm and non-farm businesses in rural and urban areas of Thailand. An extension analyzed banked and unbanked businesses in Spain (Karaivanov et al., 2019). Preceding work included retrospective surveys on wealth and distinguishing across alternative models of occupational choice subject to financing constraints (Paulson et al., 2006).

The financial settings studied in KT include the exogenously incomplete market settings discussed above (autarky, saving-only, and non-contingent debt subject to natural borrowing limit), as well as financial/information environments with obstacles to trade¹³ that limit insurance, including constraints stemming from limited commitment, unobserved income or moral hazard with an unobserved action, with the former two discussed

¹³By obstacle to trade, we mean any real-world information, commitment or other constraints that render the full-information optimal risk sharing infeasible.

in Section 3.¹⁴

In the underlying economic environment in KT, risk-averse households run SMEs with capital and labor inputs. Income is endogenous and subject to idiosyncratic random shocks. The inclusion of capital recognizes that business investment is jointly determined with household consumption in utilizing the available (disposable) income, and potentially subject to financing constraints. Including capital and its dynamics over time is also particularly relevant for the possible occurrence and persistence of poverty traps, as discussed in Section 2. Indeed, households who cannot invest because of low-income shocks may deplete their working capital and reduce their wealth and consumption over several years. The KT environment is dynamic, with multi-period contracts and extends to infinite horizon planning. The computational approach allows arbitrary functional forms for preferences and business technology (e.g., non-parametrically calibrated from the data) while standard parametric versions such as constant relative risk aversion (CRRA) preferences and Cobb-Douglas production functions are also featured, with parameters estimated from the data.

The main insight in KT is combining linear programming with maximum likelihood estimation. This allows for a direct mapping between the numerical solutions for the different financial settings, already in probabilistic form,¹⁵ and the likelihood (a measure of the models' fit with the data) which may be unavailable using other solution and estimation methods. The method allows for measurement error, estimation of the distribution of unobserved state variables, and the use of data from transitions before households reach a steady state.

Crucially, KT find that Thai households running SMEs who live in rural areas face different financing constraints from households in urban areas. The best-fitting financial regime in the villages is typically saving-only or non-contingent borrowing and lending, while the best-fitting financial regime in the urban areas is endogenously constrained by information constraints (moral hazard). This conclusion mirrors other work that analyzed the Townsend Thai data such as Paulson et al. (2006) and Ahlin and Townsend (2007).

This evidence shows, that for specific groups of households, access to safety nets is much more limited. These groups would benefit most from improving access to better income insurance. The welfare of such households, either in terms of insulating consumption and preserving their investment into their working capital (indeed, many are

¹⁴More recent work by Ru and Townsend (2022), includes costly state verification, as in Townsend (1979), in which output/income can be verified at a cost. Additionally, these methods allow combinations of obstacles as part of the same environment, for example moral hazard with unobserved output or limited commitment.

¹⁵For the exogenously incomplete markets settings (e.g., saving-only), one maximizes the agents' utility subject to resource and borrowing constraints. For the obstacle-constrained financial/information settings, one maximizes the profits of an intermediary (insurance platform) subject to a specified utility level for the agents with a parametric distribution estimated from the data or set to generate zero ex ante profits for the intermediary as in actuarially fair insurance.

self-employed) could increase considerably. In addition, to the extent that the cyclical fluctuations of income can also be smoothed via inter-temporal risk-sharing schemes, recessions would be less deep and possibly less frequent.

The welfare of households altogether combines the level of consumption and its smoothness from period to period. As is well known, households typically value such smoothness, notably to avoid having to cut consumption drastically in some periods. This preference is typically associated with the notion that households are averse to income risk. It also confirms the preference for avoiding insecurity and volatility of income.


4.2 Quantifying gains from improved risk sharing

4.2.1 Rural Thai households

We use the KT approach described in Section 4.1 to showcase new computed examples that demonstrate the gains from improved safety nets and risk sharing in a developing country context. We compare a very constrained financial setting, in which households can only save, with a setting in which households can both save or borrow at a fixed interest rate, and with an insurance platform setting with contingent transfers (premiums or indemnities) subject to a limited commitment friction. In the latter case, a household may decide to leave the insurance scheme, for example, after high income realization, and move to the save-only setting. This commitment problem may constrain the degree of feasible risk sharing by requiring that incentive-compatible participation be respected. Likewise, the commitment problem can reduce borrowing to saving-only, if households refuse to pay back loans voluntarily. The insurance regime allows borrowing, including premiums paid into an insurance pool, then rebates when a household would be unable to pay back a loan, also with accounts sequestered in escrow, preventing households from being over-indebted in the first place, but all subject to limited commitment (and more generally obstacles to trade) and to initial voluntary participation and voluntary contributions.

We quantify the gains from improved insurance for a sample of 140 rural Thai households with limited family connections, who own and run small businesses. Previous work (Karaivanov and Townsend, 2014; Ru and Townsend, 2022) has shown that this group tends to be the most vulnerable, e.g., the financial regime that best characterizes their ability to smooth income shocks corresponds either to *savings only* I or to *savings with limited borrowing*. Conversely, these are households who are the most likely to benefit from new tools which allow improved information flows and trust among strangers. Figure 1 displays summary information for the sample. There is substantial variation in the consumption and capital stock in the cross section of households, for various year, (Panels A and B). Even more telling, there is substantial variation for a household over time (Panel C). The median coefficient of variation of consumption of approximately 0.35

Figure 1. Data summary – Thai rural households without kin




Note: the left-hand and centre panels plot the cross-sectional distribution (inter-quartile range (IQR), mean, median, 10th and 90th percentile) of consumption and capital stock across the 140 households. The right-hand panel plots the cross-household distribution of the coefficient of variation of consumption and capital computed over the time period 1999–2005. The average income of approximately 60K Thai baht for this sample was equivalent to USD 5.5K in PPP terms.

means that for the median household consumption is below two standard deviations of the mean once every 12 years. As a median, half the households bear more risk than that. The implications for capital should not be ignored either. Variability in investment reflects sensitivity of investment to cash flow, that is, households being constrained from taking advantage of productive opportunities (Samphantharak and Townsend, 2018) for lack of funding. This implies potentially large welfare gains from smoothing consumption and capital investment, which we quantify and illustrate next.

Figure 2 plots the model-simulated time paths of income, consumption and capital (top panels) for an example household with zero initial savings or debt and positive initial capital stock. We use the estimated parameters from KT for the saving-only setting and the sample of 140 households without kin network in 1999. (Figure 1). We plot the model-simulated time paths for the two exogenously constrained settings (save-only, in blue and borrow/save, in red) and for the insurance setting with limited commitment, in yellow. In this example an exogenous income process is imposed (the dashed black line) common to all the graphs.

Figure 2 shows that the insurance scheme achieves a significantly smoother consumption profile over time than the save-only and borrow/save settings, implying a welfare gain. In addition, the insurance platform allows the households' capital stock to reach its optimal level very quickly, essentially “borrowing” from future or higher income states. In less perfect regimes capital stays low, reflecting a kind of poverty trap. Thus, insur-

Figure 2. Gains from insurance — example time paths



Note: in model units. 1 model unit = 179,172 Thai baht.

ance implies an additional gain from production efficiency (and thus in expected income, although in this simulation and the figure the effect of higher capital on income is suppressed).¹⁶ The insurance scheme is also able to maintain a high level of consumption and capital over consecutive periods of low income, unlike the save-only or borrowing-constrained settings. The bottom panels in Figure 2 show the mechanisms through which consumption and capital stock smoothing is achieved in each setting – by running up or down a buffer stock of savings or debt (Panel C) or by paying insurance premia or receiving indemnity payments (Panel D), depending on realized income in each period. Note that for insurance, households frequently receive indemnities while those who experience high income pay premia.

Figure 3 further explores the poverty trap possibility. The insurance platform/setting can help financially constrained households who are unable to invest, because of an inability or restricted ability to borrow. Figure 3 shows four such example households from the model-generated data. They all have low initial capital stock, which in the save-only

¹⁶See Appendix Graph A2, in which we report simulated results allowing income to be endogenously determined by each financial setting, as in KT (2014).

setting implies low expected income, low consumption, and inability to invest. Saving only is particularly damaging. Borrowing is helpful overall, especially for two of the four households. For all of them, these obstacles are removed by participating in the insurance platform, which allows these household to quickly build up their capital stock (plotted on Figure 3) and attain higher expected income and consumption (not plotted).


The gains from improved risk sharing in the insurance setting can be quantified by asking what amount a household would be willing to pay (from their wealth or savings) to permanently move from a restricted financial setting (save-only or borrow/save) into the insurance setting and be as well-off as in their current arrangement, in terms of the levels of consumption that they can expect over time (the technical term used by economists is present-value-expected utility of future consumption streams). Figure 4 computes and plots these willingness-to-pay amounts in Thai baht as a function of initial capital in hand, k . The gains from insurance expressed in terms of willingness to pay to enter the scheme are very large, about 300,000 baht for the save-only setting and about 215,000 baht for the borrow/save setting. In comparison, median yearly consumption in the sample is about 59,000 baht. The estimated gains do depend on the risk-aversion value used in the simulations (for readers familiar with functional forms for utility functions, the figure uses a CRRA function with risk-aversion coefficient estimate 2.9) but remain large compared with median yearly consumption if the households were less risk-averse than estimated: see Appendix Figure A3.

In Figure 5 we further illustrate and quantify the distribution of gains from improved insurance for all 140 Thai rural households over the period 1999 to 2005. We use the actual income data for each household and year (mapped onto a discrete grid for the computation, see KT) and the actual initial capital stock for each household.¹⁷ Figure 5 shows that, for the same estimated parameters used to simulate the three settings which feature relatively high risk-aversion (for readers familiar with utility functions, here we use a CRRA coefficient of 2.9) and hence a strong dislike for consumption variation across time and income states, the insurance platform smooths the households' consumption and capital almost perfectly, confirming the intuition from the example in Figure 2. The degree of smoothing is significantly larger than that in the save-only and borrow/save settings, where the ability to smooth consumption and business capital is constrained by the restricted financial instruments available to the households. Note also that the levels of consumption and capital in the insurance setting are higher for the majority of households, as compared with the other two financial settings.

The consumption gains of moving to the insurance scheme are large (see Figure A1 in the Appendix). Indeed, the median estimated increase in consumption from improved


¹⁷In this simulation the initial debt/savings for each household are set to zero in the save-only and borrow/save settings and the initial household present value (promised) utility is set to the level achieving zero ex ante profits for the insurance scheme in the limited commitment insurance setting.

Figure 3. Gains from insurance – overcoming poverty traps – capital



Notes: Household capital stock in model units, 1 model unit=179,172 Thai baht.


Figure 4. Gains from insurance – willingness to pay



Notes: The units are thousands of Thai baht.

insurance is almost 200% in 1999 and 90% in 2005 for the save-only setting (Panel A) with the median increase for the other years in between these numbers, and between 158% in 2001 and 62% in 2005 for the borrow/save setting. Some households with particularly low consumption in the constrained settings register gains of 800% or more. Since the insurance setting smooths out the cross-sectional and time variation of consumption, a relatively small fraction of households end up with a lower level of consumption in the insurance setting. However, they still attain higher expected utility than in the other settings because of the hugely reduced variability of their consumption.

Figure 5. Gains from insurance – consumption and capital smoothing




Notes: The units are thousands of Thai baht.

Figure 6 illustrates the mechanism through which the insurance platform achieves consumption- and investment-smoothing. For the KT parameter estimates used in the simulations, we see that most households in the insurance setting receive a net transfer (indemnity). The indemnity is larger in the first period (1999) as a bulk of new investment is undertaken to raise the households' capital stock to the optimal level (see Figure 5).

Panel 7.B shows the distribution of the premia (in dark color) or indemnities (in white) across all 140 households and the seven years of data. We see that very few (only four) households are net contributors in all years; the rest of the households are either receivers of indemnity in some of the years or (if their income is consistently low) in all years.¹⁸

Figure 6. Insurance premia and indemnities



Note: Insurance premium (if positive) or indemnity (if negative) is defined as household income minus (consumption plus investment). In the right-hand panel, each vertical bar denotes a household; a dark color indicates a net contributor (paying a premium) in a given year, while a light color indicates a net receiver (receiving an indemnity).

4.2.2 Spanish firms

A second application of the KT approach demonstrating the gains from improved risk sharing is based on Karaivanov, Saurina and Townsend (2019), hereafter KST, who examine the effect of financial constraints on business investment and cash flow using Spanish firm data. The authors classify firms according to whether they are family-owned or belong to a family-linked network, versus not family-owned, and according to the firms' number of banking relations (with none, one, or several banks). Estimating alternative financial settings via maximum likelihood, the authors find that family firms are less financially constrained and better able to allocate funds and smooth investment across states of the world and time than non-family firms, especially compared with unbanked non-family firms.


We thus use the sample of unbanked non-family firms ($N=14,152$) in the period 2004–07 and the structurally estimated parameters for the save-only setting from KST, to

¹⁸The simulation assumes zero ex ante profit (no ex ante deficit or surplus) for the insurance platform, based on contracting with a large number (continuum) of households facing an endogenously modelled income process (see KT). For a specific sample of households and income realization paths, e.g., such as those featured here, the platform may incur a deficit or surplus.

evaluate the gains from the smoothing of cash flow shocks on the firms' business capital investment. Unlike in our application to Thai households in Section 4.2.1, the firms are assumed to be risk-neutral in KST and hence the focus is on investment and capital-smoothing and not on consumption-smoothing.

Figure 7 illustrates and quantifies the gains from improved insurance by comparing the distribution of firms' capital stocks in three alternative financial settings: save-only, borrow/save, and limited-commitment constrained insurance, analogous to Figure 5. We use the actual cash flow series for each firm and year mapped onto a discrete grid for the computation (see KST) and the actual initial capital stock for each firm. Figure 7 shows that, for the same estimated parameters used to simulate all settings, the insurance setting results in almost perfect smoothing of firms' capital and investment, confirming our findings from Section 4.2.1 and Figure 5. In contrast, in the save-only and borrow/save settings, the firms' ability to smooth investment and the resulting capital level is constrained, which results in some firms being unable to attain or maintain the optimal level of capital investment when facing negative cash flow shocks.

Figure 7. Gains from improved risk sharing, Spanish firms - Capital



Notes: The units are thousands of Euro.

5 Conclusions

We reach four main conclusions. First, the gains from new designs for risk-sharing safety nets under the proposed framework can be large vis-à-vis the currently existing social policies or reliance on individual savings and absence of pooling. The gains are the largest for households and SMEs with limited family networks and no access to financial services.¹⁹ Second, improved insurance against income shocks can limit the occurrence

¹⁹The gains from sharing macroeconomic risks could also be significant across regions or countries; e.g., see the application of our framework to an incentive-compatible unemployment re-insurance scheme

of poverty traps and reduce inequality of opportunity and income. It also allows for income-smoothing, which could significantly reduce income insecurity and volatility, in turn contributing to macroeconomic stability. Third, we show how the Karaivanov and Townsend (2014) approach can be used to characterize the financial setting of different groups of economic agents and quantify the welfare gains associated with new and improved mechanism-design based safety nets. Finally, we outline how digital safety nets can be implemented using financial apps or smart contracts on digital platforms coded to address asymmetric information or commitment obstacles.

The welfare gains we quantify come from multiple metrics. The first is a better smoothing of idiosyncratic shocks, hence less variable individual consumption. Related to this, investment can be based more closely on productivity and less sensitive to fluctuations in cash flow. Comparing levels with and without improved insurance, the capital stock can rise substantially, helping people to escape or avoid a poverty trap. With the resulting output increase, consumption would also rise. In the model these gains are captured by ex-ante expected utility increases but we translate them into real terms, i.e., the amount a household is willing to pay to join the insurance platform.

From a policy perspective, our results can be used to inform innovation and complement existing policy design and as basis for safety nets in a broader perspective, improving on centralized top-down allocations of public resources to households or firms that are often subject to cyclical fluctuations, fiscal constraints or political disputes. Even with improved targeting and information about income thresholds for eligibility, existing social safety net approaches often only mitigate risks when they become systemic in a macro-stabilization perspective (e.g., the GFC, Covid-19, the war in Ukraine). These safety nets have limitations and fail to cover some unexpected idiosyncratic events that could be insured in a decentralized, bottom-up, voluntary and cooperative manner. In addition, despite built-in automatic stabilizers, the current policies have implementation lags due to the need to identify macroeconomic triggers to justify their deployment (e.g., crises, large capital outflows etc.) and these lags, and/or the lasting aftermaths of recessions, may themselves cause local poverty traps. Delays in implementation sometimes suffice to create income insecurity and excessive volatility for vulnerable households, including the self-employed and SMEs. In contrast, an approach that uses programmable digital platforms can complement the current policy tools with rapid risk-sharing assistance mechanisms that provide near real-time payments before official transfers are released.

Several caveats are in order. First, the gains from improved risk sharing, despite the existing obstacles, depend on household preferences and on SME technologies. If we have overestimated risk aversion, then we have overestimated the welfare gains. However, we have done robustness checks, which show substantial gains nevertheless. Relatedly, the

for the euro area in Karaivanov et al. (2026)

investment gains can be substantial, as illustrated in the case of Spanish firms. Second, there may also be a learning-by-doing transition when new platforms are introduced, with the need for potential initial subsidies, but we leave that for another paper and perhaps controlled trials. A third caveat is the institutional context in which the innovation would take place, which we have not yet explored. In Thailand, one idea is to allow voluntary participation at the village level through the pre-existing and near universal Million Baht village funds which provide credit. Of course, Thailand is intended as illustrative of the possible gains, and other countries could consider public, private or public-private platforms that are increasingly a part of the policy debate, both within the country and internationally (BIS and the IMF). Finally, any intervention that is scaled up can have general equilibrium effects on prices as shown in Ji et al. (2022) where interventions take the form of permitted bank branch expansion. In the data, and in the calibrated model, the expansion is spread out over time, leading to cross-market inequality. Related are welfare gains or losses due to general equilibrium effects. While everyone gains if there are no pre-existing commercial bank branches in a market and these gains are substantial for all, some agents in markets with pre-existing branches can take a loss.

Finally, we do not minimize the importance of improving the existing social safety nets and do not suggest that the debate about broad social policies and their overall resource envelope is unnecessary. This paper does not by any means suggest a substitution and/or a reduction of existing nets and their funding in favor of a “risk sharing only” approach. In fact, given the large potential efficiency gains from improved risk sharing, a promising program would be to make both approaches complementary, so that risk sharing becomes integral part of policy design. The benefits of digital platform technologies can expand economic opportunities for local SMEs and low-income households which are most frequently subject to the income shocks that feed inequality and volatility.

References

- [1] Adrian, T, F Grinberg, T Griffoli, R Townsend and N Zhang (2022), “A multi-currency exchange and contracting platform”, *IMF Working Papers*, no 2022/217, International Monetary Fund.
- [2] Ahlin, C and R Townsend (2007), “Using repayment data to test across models of joint liability lending”, *The Economic Journal* 117:F11–F51.
- [3] Aldasoro, I, P Beltrán, F Grinberg and T Mancini-Griffoli (2023), “The macro-financial effects of international bank lending on emerging markets”, *Journal of International Economics*, 142, no 103733.
- [4] Acemoglu, D and J Robinson (2012), “Why nations fail: the origins of power, prosperity, and poverty”, *Finance and Development*, 49(1).

- [5] Alem, M and R Townsend (2014), “An evaluation of financial institutions: Impact on consumption and investment using panel data and the theory of risk-bearing”, *Journal of Econometrics*, 183(1):91–103.
- [6] Alvaredo, F and L Gasparini (2015), “Recent trends in inequality and poverty in developing countries”, in A Atkinson and B Bourguignon (eds), *Handbook of Income Distribution*, vol 2, pp 697–805.
- [7] Atkinson, A and B Bourguignon (2015), *Handbook of Income Distribution*, vol 2, North Holland.
- [8] Balboni, C, O Bandiera, R Burgess, M Ghatak and A Heil (2021), “Why do people stay poor?”, working paper, London School of Economics.
- [9] Banerjee, A and A Newman (1991), “Risk-bearing and the theory of income distribution”, *The Review of Economic Studies*, 58(2):211–235.
- [10] Bank for International Settlements (2017), “Understanding globalisation”, *Annual Economic Report*, June, Chapter 6.
- [11] Barrett, C, T Garg and L McBride (2016), “Well-being dynamics and poverty traps”, *GRI Working Papers*, no 222, Grantham Research Institute on Climate Change and the Environment.
- [12] Barrientos, A (2007), “Does vulnerability create poverty traps?”, University of Sussex, Institute of Development Studies Chronic Poverty Research Centre, Working Papers, no 76.
- [13] Benami, E and M Carter (2021), “Can digital technologies reshape rural microfinance? Implications for savings, credit, & insurance”, *Applied Economic Perspectives and Policy*, 43(4):1196–1220.
- [14] Bilbiie, F, G Primiceri and A Tambalotti (2022), “Inequality and business cycles”, *Janeway Institute Working Papers*, University of Cambridge.
- [15] Bird, K (2007), “The intergenerational transmission of poverty: An overview”, Overseas Development Institute Chronic Poverty Research Centre, Working Papers, no 99, October.
- [16] Bowles, S, S Durlauf and K Hoff (2006), “Poverty traps”, *Economic Development Quarterly*, 23(2):172–173.
- [17] Carstens, A (2021), “Rethinking the global financial safety net”, speech at the Sixth High-Level Regional Financing Arrangements (RFA) Dialogue, October 2021.


- [18] Carter, M (1997), “Environment, technology, and the social articulation of risk in West African agriculture”, *Economic Development and Cultural Change*, 45(3):557–591.
- [19] Challe, E, J Matheron, X Ragot and J Rubio-Ramirez (2017), “Precautionary saving and aggregate demand”, *Quantitative Economics*, 8(2):435–478.
- [20] Chetty, R, N Hendren, P Kline and E Saez (2014), “Where is the land of opportunity? The geography of intergenerational mobility in the United States”, *Quarterly Journal of Economics*, 129(4):1553–1623.
- [21] Claveres, G and J Stráský (2018), “Stabilising the euro area through unemployment benefits re-insurance scheme”, *OECD Economics Department Working Papers*, no 1497, OECD Publishing.
- [22] Dolls, M (2024), “An unemployment re-insurance scheme for the eurozone? Stabilizing and redistributive effects”, *European Economic Review* 170(104872).
- [23] D’Silva, D, Z Filková, F Packer and S Tiwari (2019), “The design of digital financial infrastructure: lessons from India”, *BIS Papers*, no. 106.
- [24] Durlauf, S (1994), “Spillovers, stratification, and inequality”, *European Economic Review*, 38:836–845.
- [25] Forbes, K and F Warnock (2012), “Capital flow waves: Surges, stops, flight, and retrenchment”, *Journal of International Economics*, 88(2):235–251.
- [26] Goldfarb, A and C Tucker (2019), “Digital Economics”, *Journal of Economic Literature*, 57(1): 3–43.
- [27] Ghatak, M (2015), “Theories of poverty traps and anti-poverty policies”, *The World Bank Economic Review*, 29(S1):S77–S105.
- [28] Guvenen, F, L Pistaferri and G Violante (2022), “Global trends in income inequality and income dynamics: New insights from GRID”, *Quantitative Economics*, 13:1321–1360.
- [29] Heathcote, J, F Perri and G Violante (2020), “The rise of US earnings inequality: does the cycle drive the trend?”, *Review of Economic Dynamics*, 37:181–204.
- [30] 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.
- [31] Jalan, J and M Ravallion (1998), “Transient poverty in postreform rural China”, *Journal of Comparative Economics*, 26(2):338–357.

- [32] Ji, Y, S Teng and R Townsend (2022), “Dynamic bank expansion: spatial growth, financial access, and inequality”, working paper, MIT.
- [33] Karaivanov, A and R Townsend (2014), “Dynamic financial constraints: distinguishing mechanism design from exogenously incomplete regimes”, *Econometrica*, 82:887–959.
- [34] Karaivanov, A, J Saurina and R Townsend (2019), “Family firms, bank relationships and financial constraints: a comprehensive scorecard”, *International Economic Review*, 60:547–593.
- [35] Karaivanov, A, B Mojon, L Pereira da Silva, A Pierres Tejada and R Townsend (2026), “Incentive-compatible unemployment reinsurance for the euro area”, *European Economic Review* (forthcoming).
- [36] Koepke, R (2019), “What drives capital flows to emerging markets? A survey of the empirical literature”, *Journal of Economic Surveys*, 33(2):516–540.
- [37] Kinnan, C, K Samphantharak, R Townsend and D Vera Cossio (2020), “Propagation and insurance in village networks”, *NBER Working Papers*, no 28089.
- [38] Kraay, A and D McKenzie (2014), “Do poverty traps exist? Assessing the evidence”, *Journal of Economic Perspectives*, 28(3):127–148.
- [39] Latner, J (2019), “Economic insecurity and the distribution of income volatility in the United States”, *Social Science Research*, 77:193–213.
- [40] Liao, P, X Zhang and Z Wanlu (2022), “Endogenous health risks, poverty traps, and the roles of health insurance in poverty alleviation”, *Health Economics Review*.
- [41] Morelli, S, T Smeeding and J Thompson (2015), “Post-1970 trends in within-country inequality and poverty: rich and middle-income countries”, in A Atkinson and B Bourguignon (eds), *Handbook of Income Distribution*, vol. 2, p.593–696.
- [42] Nilekani, N (2018), “Data to the people: India’s inclusive internet”, *Foreign Affairs*, 97(19).
- [43] Paulson, A, R Townsend and A Karaivanov (2006), “Distinguishing limited liability from moral hazard in a model of entrepreneurship”, *Journal of Political Economy*, 114(1):100–144.
- [44] Pereira da Silva, L, E Kharroubi, E Kohlscheen, M Lombardi and B Mojon (2022), *Inequality hysteresis and the effectiveness of macroeconomic stabilization policies*, Bank for International Settlements.

- [45] Rohde, N, K Tang and D Prasada Rao (2014), “Distributional characteristics of income insecurity in the U.S., Germany, and Britain”, *Review of Income and Wealth*, 60(1):S159–S176.
- [46] Rosenzweig, M and H Binswanger (1993), “Wealth, weather risk and the composition and profitability of agricultural investments”, *The Economic Journal*, 103:56–78.
- [47] Ru, H and R Townsend (2022), “Kinship networks enhanced by a formal financial intervention: costly state verification and the Thai Village Fund”, working paper, MIT.
- [48] Samphantharak, K and R Townsend (2018), “Risk and return in village economies”, *American Economic Journal: Microeconomics*, 10(1):1–40.
- [49] Thanh, NX (2005), “The injury poverty trap in rural Vietnam: causes, consequences and possible solutions”, *Epidemiology and Public Health Sciences*, Department of Public Health and Clinical Medicine, Umeå University.
- [50] Townsend, R (1979), “Optimal contracts and competitive markets with costly state verification”, *Journal of Economic Theory*, 21(2):265–293.
- [51] Townsend, R (1982), “Optimal multiperiod contracts and the gain from enduring relationships under private information”, *Journal of Political Economy*, 90(6):1166–1186.
- [52] Townsend, R (1994), “Risk and Insurance in Village India”, *Econometrica*, 62(3):539–591.
- [53] Townsend, R (2016), “Village and larger economies: The theory and measurement of the Townsend Thai project”, *Journal of Economic Perspectives*, 30(4):199–220.
- [54] Western, B, D Bloome, B Sosnaud and L Tach (2012), “Economic insecurity and social stratification”, *Annual Review of Sociology*, 38:341–359.
- [55] World Bank (2021), *Global Findex Database: Financial Inclusion, Digital Payments, and resilience in the Age of COVID-19*.


6 Appendix

Figure A1. Gains from insurance — consumption growth




Note: The figure plots the distribution of consumption growth, defined as the ratio of the consumption level in the insurance setting to the consumption level in the restricted setting (save-only or borrow/save), computed for each household and year.

Figure A2. Gains from insurance – capital accumulation



Note: The figure plots the distribution of end-of-period capital stock for each of the three settings (save only, borrow/save, and insurance). Each household is initialized at its 1999 capital from the data, with zero debt or savings and the ex ante utility yielding zero insurer profits. The income process is endogenous, as implied by the model solution. Units are thousands of Thai baht.

Figure A3. Willingness to pay for insurance (robustness)



Note: We plot agents' willingness to pay to move to the insurance setting, as in Figure 5, computed for different values of risk aversion in agents' preferences: the baseline estimate (panel A), lower values (panels B and C), and a smaller discount factor (panel D). Units are thousands of Thai baht.