BIS Papers
No 139

Digital safety nets: a roadmap

by Alexander Karaivanov, Benoît Mojon, Luiz A Pereira da Silva and Robert M Townsend

Monetary and Economic Department

September 2023

JEL classification: G10, G52, G23, D82, D53, O16.

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

The views expressed are those of the authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

ISSN 1682-7651 (online)
ISBN 978-92-9259-683-5 (online)

Digital safety nets: a roadmap[1]

Alexander Karaivanov[†], Benoît Mojon[‡], Luiz A Pereira da Silva[‡] and Robert M Townsend[§]

Abstract

We show how new digital technologies can be used to improve safety nets for insurance against idiosyncratic and aggregate income risks, tailored to deal specifically with well-known obstacles to trade: limited commitment, moral hazard, unobserved states and payment transaction costs. We illustrate the gains from incentive-compatible voluntary risk-sharing schemes for groups of economic agents, eg Thai households and Spanish firms. We assess the currently best-fitting financial regime within each group and quantify large welfare gains from improved insurance despite the obstacles to trade. Our methods could be applied in various contexts to foster financial inclusion and complement existing broader safety net 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 classification: G10, G52, G23, D82, D53, O16.

1 This paper reflects the views of the authors and not necessarily the views of the BIS. 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 number R01 HD027638). We acknowledge fruitful discussions with and comments from participants to webinars at the BIS and the IMF and from Katharina Bergant, Benjamin Carton, Gong Chen, Jon Frost, Leonardo Gambacorta, Vitor Gaspar, Pierre-Olivier Gourinchas, Benedict Nolens and Hyun Song Shin.

† Simon Fraser University (SFU), akaraiva@sfu.ca

‡ Bank for International Settlements (BIS), benoit.mojon@bis.org and luiz.apereira@bis.org

§ Massachusetts Institute of Technology (MIT), rtownsen@mit.edu

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 then were extended to emerging and developing economies (EMDEs).[2] Arguably, they have proved reasonably effective in smoothing household income and expenditure and in stabilising 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 and up 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.[3]

However, many countries have seen the trend in income inequality reverse direction since the 1980s. The relation between the level of income and income inequality has taken a U-shaped curve[4] with a steady increase in inequality of opportunities, income and wealth, as well as in geographical inequality (eg access to public goods within the same jurisdiction). Moreover, this increase partly reflects that individuals and households with lower skills or incomes are subject to a combination of larger idiosyncratic and macroeconomic income risks. In countries that are members of the Organisation for Cooperation and Development (OECD), not only is the average unemployment rate of low-skilled workers higher, but it also increases by more in recessions, implying a hysteresis in income inequality (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 the self-employed and workers in informal markets to slide into poverty traps, which can persist, in some cases, across generations. Relatedly, risks may be so large that households do not undertake entrepreneurial activities or limit the scale of family-run enterprises. This deepens economic downturns and lowers potential growth.

2 A distinction should be made between broader social policies dating back to the 19th century and significantly expanded and formalised after the US New Deal in the 1930s and the legislation following the Beveridge Report in Britain in the 1940s; and social safety nets that became popular during the 1990s, aimed at mitigating the impact on the poor when developing countries and the Bretton Woods institutions began implementing structural adjustment policies. Among the latter programmes are conditional cash transfers (CCTs) to households eligible for a pecuniary transfer from the government when their income is below a certain level and satisfies certain conditions. Another important policy development centred around increasing "financial inclusion" to the poor, ie creating incentives and/or regulations that improve their access to useful and affordable financial services offered by a range of competing providers (see Princess Máxima (2010), CPMI and World Bank (2016)).

3 See Morelli et al (2015) and Alvaredo and Gasparini (2015).

4 See Atkinson and Bourguignon (2015), p xviii.

Income volatility can be triggered by events or accidents (unemployment, family break-up, poor health etc), some of which are related to the business cycle[5] but also by idiosyncratic shocks affecting workers and households. Volatility increases economic insecurity and the risk of falling into poverty despite existing public social safety nets. Higher economic insecurity is empirically more prevalent among the low-income groups. This reinforces the link between economic insecurity and the rise in economic inequality.[6] Therefore smoothing income risks is an important component of policy considerations.

Turning to country-level income shocks, crisis episodes often associated with large capital outflows compound macroeconomic income risks by reducing the fiscal financial resources devoted to institutional social safety nets. Under these circumstances, EMDEs have had to undertake fiscal adjustment when their macroeconomic stability came under pressure from volatile capital outflows and sudden stops. In such instances a reduction of their social safety nets represents an additional challenge. 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 older and 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 of these together (see Karaivanov et al (2023)).

Why is that so? The 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 (a growing premium for higher education) with the effects of globalisation on employment in advanced economies, associated with the relocation of significant segments of productive activities. A complementary hypothesis with respect to household income risks 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 stabilisers and were a pillar of steady growth during the post-World War II period. However, both demographic and unemployment trends in developed countries 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.[7] 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 United States and Pereira da Silva et al (2022) for cross-country evidence).

It is also true that most of the above-mentioned policies and institutions were designed before the availability of detailed information about individuals and households and usually without data that were 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 annual household surveys

5 See Challe et al (2017), Heathcote et al (2020), Guvenen et al (2022), Bilbiie et al (2023), among others and references therein.

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

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

with more granular data such as annual household surveys, and the emergence of conditional cash transfers (CCTs), the policies tended to rely on broad approaches and frameworks.[8] In addition, the cost of reliable and timely information about the relationship between the situation of individuals or households and the state of the economy at any given time compounded these problems. It is difficult, if not impossible, to verify an individual's situation facing an adverse outcome, and whether a claim for protection is justified and true, while addressing income volatility in a timely fashion. Empirical evidence points to the prevalence of gift-based assistance within extended families where trust plays a key role. 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 hindrance to developing effective contingent and targeted safety nets.

However, recent technological developments allow more sophisticated interventions that process more granular data on individual characteristics at very low cost[9] (Goldfarb and Tucker (2019)). Another development is the increased understanding of contracts that induce truthful self-reporting of unobserved individual circumstances. Contracts can also deal directly with commitment or moral hazard problems delivering constrained-optimal insurance arrangements. Further, honest messages can be encrypted, preserving privacy,[10] while allowing them to serve as an input into encoded risk-sharing agreements, and 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 new technology 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 a large share of idiosyncratic risks, regional shocks and country-level shocks. It could make the set of available insurance arrangements less incomplete, preserve equal opportunities

[8] In many EMDEs, but also in advanced economies, such policy instruments tend to be subject to inertia and political economy pressures to support social policies using agencies at various levels of government (local, regional, central). 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 to discuss, 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 with income thresholds or some measures implemented during the Covid-19 pandemic and the war in Ukraine) but remain predominantly dependent on financing constraints, centralised and transfer-based.

[9] Atkinson and Bourguignon (2015) recognise that technological progress will fundamentally change social policies. However, they focus much more on the effects of faster payments of lump sum amounts to beneficiaries than the design of policies that include new forms of insurance and contracts.

[10] Interested readers should refer to the discussion in D'Silva et al (2019) and Nilekani (2018) on the development of privacy-preserving financial infrastructure in India.

across economic agents and prevent unnecessary human tragedies, while at the same time providing a more effective stabilisation of the business cycle, higher productivity and improved financial flows.

Put another way, we propose, inter alia, a different approach to social safety nets,[11] 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 centralised system of controls and payments, we are suggesting an emphasis on mitigating income risk (especially idiosyncratic). This may include a *decentralised* implementation process and incentive-compatible contracts that collect contributions and pay indemnities. We argue that current technology can circumvent obstacles to the development of effective digital safety nets, which could help reduce both income inequality and volatility. These elements would provide a proof of concept and facilitate a shift in mindset among policymakers.

Complementing existing safety nets would enhance welfare and stabilise the income of low-income households, which in turn implies higher and more stable economic growth. Relatedly, insurance allows more flexible terms for repayment of credit and self-sustained conditional transfers funded with unconditional contributions. The economy would benefit from the enormous potential for creativity, productivity and innovation of groups that face differentiated, higher risks and inequality of opportunity.

Indeed, technologies are available today that could pool people's efforts and resources to reduce risk, with much lower implementation and verification costs and thus higher social rates of return as compared with traditional social safety nets. The new technologies and tools that we describe in this paper make it possible to design various components of social policies and financial inclusion programmes with characteristics that could easily complement those introduced in the 20th century. In summary, the key conceptualisation is how to address income risk, first by illustrating the large welfare gains that improved risk-sharing can deliver and then explaining how digital safety net platforms and apps can be coded and implemented to overcome well known information asymmetries or commitment hurdles.

This paper proceeds as follows. In Section 2, we explain the link between idiosyncratic risks and macroeconomic fluctuations. In Section 3, we illustrate improved risk-sharing schemes and document the state of affairs on the ground through survey evidence and salient case studies of digital innovation. In Section 4, we argue from mechanism design first principles that, although limited commitment and asymmetric information can limit risk-sharing, new technological tools could enable constrained-optimal arrangements that unlock considerable gains from voluntary participation. Section 5 presents a method for the assessment of the current state of affairs, typically characterised by limited risk-sharing, and provides quantitative estimates of the welfare gains to innovation in the context of Thai rural households and Spanish firms.

11 We show below that the approach can also apply to different jurisdictions and different economic actors (firms).

2. Idiosyncratic risk and macroeconomic fluctuations

In this section, we review both traditional idiosyncratic risks and heterogenous exposures to cyclical risks. We stress that uninsured idiosyncratic risk can translate into a loss that is material enough to transmit to the family and then to neighbouring 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 (eg 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 create or compound poverty traps and produce additional risks that might have potential local, regional and eventually macroeconomic effects. Clearly, the prevention of such chains of events is social welfare-enhancing.

2.1 Empirical evidence for limited insurance

A direct look at data on outcomes provides confirmation of the adverse consequences of income shocks and evidence of limited insurance. With complete risk-sharing, an individual's consumption should track the consumption of the community as a whole and be immune to that individual's income fluctuations. A direct look at income and consumption panel data over 10 years from villages in India (Townsend (1991)) 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 labourers, 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 Thai biannual cross-sectional surveys into a pseudo survey of regional and occupation cohorts helps track the income and consumption data across social groups. It shows that entrepreneurs and those in and around Bangkok, who have arguably left village-based social protection, bear the ups and downs of fluctuations and so would benefit from increased insurance (Townsend (2016)).

Additional data include the choice of occupations, sectors of chosen businesses, and the rates of return on those business. Specifically, if there were perfect mutual insurance within a village, then idiosyncratic fluctuations in returns would be entirely smoothed. 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.[12]

12 International investors in theory should allocate capital to countries with the highest return until returns across different countries are equalised.

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 international capital flows remain procyclical is overwhelming (eg 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 (Feldstein and Horioka (1980)). China does poorly in inter-regional flows relative to more advanced economies such as the United States. Internationally, countries in the EU, a monetary but not fiscal union, do poorly in sharing certain risks, as exemplified by inadequate cross-country unemployment insurance relative to what is achieved in the United States (Dolls (2019) and references therein).

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 analysed 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 programme, 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 to consumption and investment data points to substantial heterogeneity in outcomes linked to incomplete markets and obstacles to trade (Karaivanov and Townsend (2014)), a topic to which we return below.

2.2 Poverty traps

Another 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, this persists without outside intervention, in some cases even across generations, because these individuals have limited or no resources and their capital (physical or human) has a low return. Therefore, these self-reinforcing disadvantages make it extremely hard to escape poverty (Bowles et al (2006), Kraay and McKenzie (2014), Ghatak (2015), Balboni et al (2021), Banerjee and Duflo (2012)).

The literature has documented several mechanisms that lead to a poverty trap. For instance, Barrientos (2007) suggests that non-linear income dynamics, low asset endowments, and lack of access to credit are at the origin of poverty traps. Xuan 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 an "injury poverty trap". Liao et al (2022) show that, in the absence of health insurance, poor and vulnerable people tend to overstress their bodies to earn more money to invest in education. Rosenzweig and Binswanger (1993) and Carter (1997) showed how risk preferences can induce poor agricultural households that lack access to credit and insurance markets to choose low-return livelihoods as a way of self-insuring against weather risk. Again, those choices can also trap them in chronic poverty. Barrett et al (2016) also survey the range of the mechanisms generating poverty traps, such as poor nutrition and (mental and physical) health, endogenous behavioural patterns (eg risk and time preferences), poorly functioning capital markets, large uninsured risk exposure, and weak natural resource governance institutions.

2.3 Transmission to family and larger groups

Once a household has fallen into a poverty trap due to an idiosyncratic shock, this might have ripple effects and spread to larger groups at the local, regional or even national level. The idea of "contagion" is usually used in finance and 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 specific segment of the population, especially the most vulnerable and the poor. Various, partially overlapping mechanisms can be identified.

A more 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 labour. This creates amplified shocks that reverberate through supply chains and local labour markets, as revenues diminish. Consumption drops for those affected. Furthermore, broken transaction links persist over time, ie they are not easily replaced by alternative points of sale or employment, inflicting long-term structural damage and resulting in persistently lower incomes.

Another mechanism concerns 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 neighbourhoods 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)).

A third mechanism that links households with the macro-economy is occupational poverty traps, whereby the combination of borrowing constraints and lumpy production technologies means that poor individuals who start businesses that are too small can be trapped into earning returns of no more than subsistence level (Banerjee and Newman (1991)). Therefore, the concept of a poverty trap at the level of national economies is related to, and sometimes based on, microeconomic foundations at the household level. This helps provide justification for microfinance loans designed to allow households to lift themselves out of poverty by buying the fixed-cost assets necessary to operate a business.

A fourth mechanism that explains the transmission of local poverty traps to the macro-economy 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 geographic 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 enjoy a rising standard of living. This can be compounded by poor institutions at both local and central levels of government (Acemoğlu and Robinson (2012)). Poor households in remote rural regions living in isolation have a reduced number of available production technologies, which means the choice between lower-income and higher-income outcomes may be a more difficult discrete step.

2.4 Macroeconomic impacts

If an idiosyncratic shock can trigger the fall into poverty traps and their transmission to larger groups, and if the economic environment of such groups compounds the effects of such shocks, then this can help to explain why within-country income inequality has increased in recent decades. This rise has been widely discussed with different metrics and results (Gini, top decile, top 1% etc). From inequality, additional dynamics follow. Pereira da Silva et al (2022) suggest a two-way interaction between inequality and recessions. Higher levels of income inequality imply deeper recessions, and the latter tend to have a very persistent effect on income inequality. The income share of the wealthiest 10% generally increases after recessions, usually remaining high for years afterwards. In addition, the paper shows that greater inequality makes monetary policy less effective when used either to stimulate or slacken aggregate demand. Finally, fiscal policy has become less redistributive and less countercyclical, putting more onus on monetary policy as a tool for macroeconomic stabilisation.

As mentioned earlier, much of this greater within-country income concentration relates to structural factors such as technological progress and globalisation, which have greatly raised the returns to skills. However, reduced fiscal redistribution, through weaker tax progressivity and flatter tax systems, or cutbacks in unemployment insurance benefits, have not only increased after-tax inequality but have also made fiscal policy less countercyclical, hence eroding its stabilising effect. Lower tax progressivity means that, in expansions, rising incomes contribute less to government revenues. Similarly, lower unemployment benefits imply less government expenditure in recessions and smaller government revenues in expansions, as reduced benefits usually go hand in hand with lower contribution rates for unemployment insurance (see Pereira da Silva et al (2022) Chapter 3). These policy changes have also contributed to increasing income insecurity and volatility.

In summary, there is a transmission from idiosyncratic shocks, individuals and groups falling into poverty traps, regional geographical and even national macroeconomic effects, and the materialisation of these risks. On the other hand, the traditional macroeconomic policies, both fiscal and monetary, that are used for stabilisation purposes have been progressively downsized and, while maintaining some minimal social transfers, have lost their focus on income inequality, despite its propensity to be aggravated by downturns, and less responsive to income insecurity, volatility and the emergence of geographical poverty traps even in developed countries.

The implication is that the policy agenda should develop in three directions. On the microeconomic front, there is a need to break the transmission of individual or group risks into more widespread risks. This requires the design, methodology and implementation of policy interventions to be changed, especially as they involve fiscal instruments. As we argue above, the availability of data and new digital technologies reduces many impediments to improving the effectiveness of interventions and their rapidity, to avoid time lags and insecurity. The data include digital granular data at the individual, household or firm level. The tools include encryption, distributed ledgers, e-transfers and programmability. Data and new tools can transform social policies by better targeting and addressing risks and making insurance more complete. This can assist communities at the micro level while also providing more effective stabilisation of the business cycle.

Three more levers can be activated at the macroeconomic level. First, safety nets could, at least in principle, also be more effective across countries. This would imply revisiting our current international institutional setup for global safety nets (eg the network of multilateral, regional and bilateral institutions). Second, within a country, the familiar menu of structural policies to reduce inequality and the accompanying stabilisation policies could be revisited to better reflect what we know from the experience of the last decades (eg more skill-intensive technical change, capital endowments, income-tested transfers, minimum social revenue, inheritance taxation, more progressive income tax structure, role of a wealth tax etc). If we add income volatility, the timing of interventions is also of the essence to avoid staying too long at income levels that risk triggering a poverty trap. All this could be re-designed and implemented to help reduce both inequality and macroeconomic instability in a coordinated way. This is due to a positive feedback loop between reduction of inequality, which tends to amplify recessions, and macroeconomic stabilisation that curbs recessions, during which inequality rises and many fall into poverty traps. Third, self-funding insurance with premiums can be integrated into borrowing contracts to provide more flexibility in repayment and lower debt levels, which in turn can be a drag on growth, restrain policy, and act as a source of systemic risk.

3. Improving risk-sharing schemes

In this section, we show via stylised illustrative examples how economic agents can engage in self and mutual insurance schemes that help them smooth their consumption against income shocks. Importantly, such tools are often unavailable, and advances in new technologies have not yet developed solutions to expand risk-sharing opportunities to their full potential.

3.1 A stylised example

To set the stage, we introduce a simple stylised economy with two agents, A and B, who are, for example, kinship-related and enter various mutual arrangements for sharing risk. Each agent receives a stream of income over 10 periods. Each period's income is a random draw from a discrete uniform distribution taking values between 6,000 and 114,000 Thai baht. The incomes of A and B are assumed to be uncorrelated. Hence, in each period an agent can expect that their income will be on average 60,000 Thai baht (equivalent to USD 5,500 per annum in PPP terms), and that the economy's total average income will be 120,000. However, the actual income may be as low as 6,000 or as high as 114,000 for each agent and as low as 12,000 and as high as 228,000 for the economy.

In Table 1a, Table 1b and Graphs 1.A to 1.F, we report one possible realisation of income streams for agents A and B over 10 periods and analyse the effects of six different ad hoc smoothing/risk-sharing schemes.[13] Agent A is depicted with blue lines, agent B with red lines. For both agents, gross income is plotted in dashed lines and disposable income (after arrangement/transfers) in solid lines.

In the first scheme ("flat tax and transfers"), instances of low income for one of the agents (income below 30,000) trigger a transfer to that agent, which is funded by collecting 10% of both agents' gross income. For instance, in period 1, agent B receives 8,4000 (a net transfer equal to 10% of agent A's income) and in period 3, agent A receives 9,000 (10% of agent's B income).[14] If both agents draw an income below 30,000, which never happens in our 10-period example, there would be no transfer between them. As shown in Graph 1.A, the scheme's benefits in terms of disposable income smoothing and risk-sharing are limited to the ability to lift income from very low levels.

In the second scheme, ("saving only"), smoothing is feasible only via individual savings by each agent. We assume a zero interest rate for simplicity. As shown in Graph 1.B, this scheme allows both agents to smooth their disposable income only after periods in which they can accumulate savings. This is the case for agent A, who

[13] While the schemes in Table 1 and Graph 1 are ad hoc, in the sense that they use simple tax or sharing rules based on exogenously set cutoffs and are not computed as solutions to dynamic maximisation problems, they illustrate the main features of optimal constrained risk-sharing contracts such as history dependence, contingent transfers and aggregate savings. We analyse constrained-optimal schemes in Section 5.

[14] Both agents contribute 10% of their income but the agent with gross income less than 30,000 gets back the taxes they contributed, so that the difference of their gross and disposable income equals to the tax collected from the high-income agent.

can smooth their negative income shock in period 2 and period 7 but not in period 3 and period 4. Agent B begins with a very low disposable income in periods 1 and 2 in which they cannot accumulate savings but can maintain a disposable income of 60,000 in periods 6 and 7 in spite of very low gross income thanks to the savings accumulated from receiving a high income in previous periods.

Effects on disposable income of various self or mutual insurance schemes

In thousands of Thai baht Table 1a

Period	Gross income		Disposable income					
			Flat tax and transfers		Saving only		Transfers	
	A	B	A	B	A	B	A	B
1	84	18	75.6	26.4	60	18	67.5	34.5
2	54	42	54	42	60	42	51	45
3	12	90	21	81	30	60	31.5	70.5
4	48	66	48	66	48	60	52.5	61.5
5	102	108	102	108	60	60	103.5	106.5
6	90	30	90	30	60	60	75	45
7	48	18	43.2	22.8	60	60	40.5	25.5
8	54	66	54	66	60	60	57	63
9	114	12	102.6	23.4	60	30	88.5	37.5
10	114	72	114	72	60	60	103.5	82.5
			Graph 1.A		Graph 1.B		Graph 1.C	

Note: the table reports the gross income flow of agents A and B as a randomly drawn value between 6 and 114.

[1] **Flat tax and transfers:** within each period, a 10% tax is collected from both agents and transferred to the agent with income below 30.

[2] **Saving only**: each agent saves income in excess of 60 and can use the savings to smooth disposable income in periods when their gross income is less than 60.

[3] **Transfers**: the agent with higher income in a period transfers to the other agent 25% of the difference in their gross incomes.

Source: Authors' calculations.

In the third scheme, in Graph 1.C we analyse an ad hoc transfer scheme where the two agents agree to transfer to one another a quarter of their income difference (but do not save). This scheme illustrates the benefits of partial risk-sharing. It is also ad hoc, in the sense that the quarter difference rule is not endogenously derived from underlying obstacles to trade. Nevertheless, this scheme is our first example of a smart contract, that is a program or algorithm that determines an outcome, namely transfers, conditional on inputs, namely incomes, doing the math simply as "if…, then…" lines of code. The solid lines (disposable income) in Graph 1.C are, by construction, much smoother than the gross income lines, as the agents insure one another and lie closer to the average income level of 60. However, the graph also shows that fluctuations remain, ie the agents are not able to fully smooth their disposable income over time and income shocks. For example, both have high income in period 5 and lower incomes in periods 6 and 7. They would clearly have preferred to save in period 5 for "rainy days".

Effects on disposable income of various self or mutual insurance schemes

In thousands of Thai baht | Table 1b

Period	Gross income		Disposable income								
			Type 1 savings and transfers		Savings pot	Type 2 savings and transfers		Savings pot	Type 3 savings and transfers		Savings pot
	A	B	A	B		A	B		A	B	
1	84	18	60	30	12	84	18	0	60	42	0
2	54	42	57	51	0	54	42	0	54	42	0
3	12	90	24	60	18	12	90	0	42	60	0
4	48	66	60	60	12	48	66	0	54	60	0
5	102	108	73.5	78	70.5	57	63	90	60	60	90
6	90	30	60	60	70.5	90	30	90	60	60	90
7	48	18	60	60	16.5	75	45	36	60	60	36
8	54	66	60	60	16.5	54	66	36	60	60	36
9	114	12	84	40.5	18	111	9	42	60	60	42
10	114	72	84	60	60	81	39	108	60	60	108
			Graph 1.D			Graph 1.E			Graph 1.F		

Note: the table reports the gross income flow of agents A and B as a randomly drawn value between 6 and 114.

[1] **Type 1 savings and transfers (Limited commitment)**: transfers as in Table 1a combined with saving in periods when individual gross income is larger than 60. Pooled savings can be used in subsequent periods when individual income after transfer is less than 60. In any period, transfers cannot exceed 12 and transfers+savings cannot exceed 30. Transfers from past savings are split equally between the agents up to their disposable income reaching 60.

[2] **Type 2 savings and transfers**: transfers are based on observed aggregate income. Half of aggregate income in excess of 120 is requested from each agent and pooled; the pooled funds are then used in any following period in which aggregate income is less than 120 (each agent receives half of the difference between 120 and aggregate income, if feasible).

[3] **Type 3 savings and transfers**: each agent contributes to a savings pool any income above 60 and receives a transfer equal to the difference between 60 and their income when the income is less than 60, subject to the available pooled funds.

Source: Authors' calculations.

In Table 1b and Graph 1.D, we report an ad hoc scheme that combines aggregate savings and transfers to illustrate that such combination can deliver smoother disposable income than either the saving only or within-period transfers schemes (Graphs 1.C and 1.B) alone. Part of aggregate income can be saved and past savings (the column "savings pot") may be drawn upon to smooth disposable income. We impose caps on transfers and savings motivated by limited commitment as an obstacle to trade. The key idea is that, in each period, both agents will voluntarily comply with the contract only as long as their outside option is less favourable. In this economy, the natural alternative for each agent could be "saving-only". An agent who draws a high income could thus be reluctant to transfer 25% of their income to the other agent or make a large contribution to a common savings pool available for both agents in the future.[15]

As a result of combining savings and transfers, periods of low aggregate gross income are better smoothed than in the transfer-only scheme shown in Graph 1.C. In addition, the smoothing of disposable income does not depend only on one's own savings. For instance, in period 9, agent B benefits from a transfer that is drawn from the aggregate savings and the high income of agent A. Importantly, such a scheme, although still ad hoc, is designed to account for incentive compatibility and voluntary participation by each agent. In addition, it may be easily implemented using a smart contract that takes income draws as inputs and maps them into agent-to-agent transfers and/or savings pot contributions.

We next illustrate the importance of the observability of individual income, which we have so far taken for granted. It may be that only aggregates are observed externally, eg average income in the region from profits and wage disbursements, and not individual earnings. Specifically, in Graph 1.E ("Type 2 savings and transfers") we assume that only aggregate income is observed. Half of aggregate income in excess of 120,000 is requested from each agent and pooled; the pooled funds are then used in any following period in which aggregate income is less than 120,000 (each agent receives half of the difference between 120,000 and aggregate income, if feasible). This can be again implemented via a smart contract which uses inputs data to determine outcomes. However, in this case the data are aggregated. Likewise, the outcome, transfers, cannot target individuals. Henceforth each agent contributes in the same proportion and receives the same amount in each period. In the first four periods, when aggregate income is lower than 120,000, no income is saved nor redistributed. In periods such as period 10 when aggregate income is 186,000, each agent is requested to contribute 33,000 to the scheme in spite of the large difference of income between the two agents. Agent B's after-tax income is only 39,000.

As a comparison, in the Type 3 savings and transfers scheme (Table 1b and Graph 1.F), contributions to the scheme and indemnities are agent-specific. Only agents whose gross income is larger than 60,000 contribute savings to the fund and only those whose income is lower than 60,000 receive an indemnity. As shown in Graph 1.F and in the last column of Table 1b, this scheme is highly effective in sharing risk within periods and over time. From period 5 onward, income is perfectly insured. Income could have been perfectly smoothed earlier too if the insurance scheme could

15 We cap each period's transfers and "transfers + savings" that a single agent is required to pay at 12,000 and 30,000, respectively. We then check the utility streams for each agent in each period, comparing staying in the Graph 1.D scheme against deviating to saving only (Graph 1.B) and verify that neither agent has incentive to deviate (renege on the requested transfer or savings pot contribution) for standard preference specifications (log and constant relative risk aversion utility).

Effects on disposable income of various self or mutual insurance schemes

In thousands of Thai baht | Graph 1

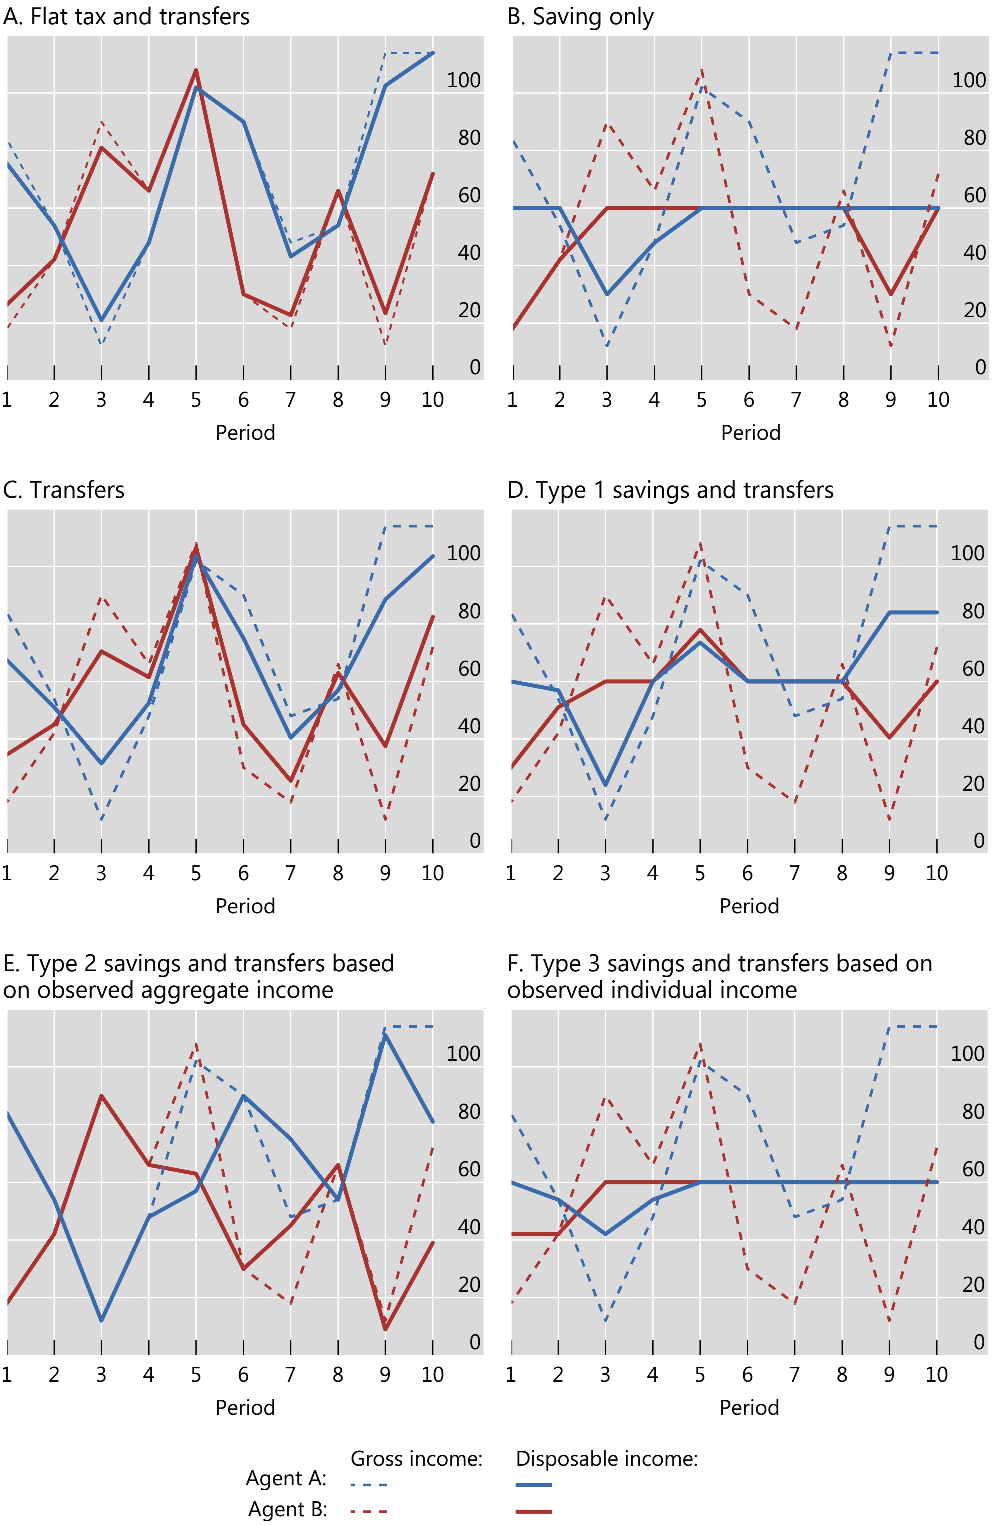

Note: for more details, see notes to Tables 1a and 1b.

Source: Authors' calculations.

have borrowed during the first four periods when aggregate income remained below 120,000. The key implication is that each agent's disposable income is smoothed around a constant. Here we take this constant to be the same (60,000) for both agents for simplicity. However, income equalisation per se is not the objective, and risk-sharing; heterogeneity across the agents (eg in productivity or ability) can easily be incorporated.

Although ad hoc, these examples clearly illustrate the benefits of risk-sharing between agents and of smoothing income risks over time and states of the world via (individual or aggregate) savings and transfers. The examples also provide intuition for the advantage that comes with observability of individual income shocks.

3.2 How these examples match the situation on the ground

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 programme (CCT). In other emergency circumstances (eg Covid-19), exceptionally large state-contingent transfers have been made. However, the latter programmes 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. With the types of scheme described in Graphs 1.B to 1.F, household consumption and investment expenditure could be greatly insulated from idiosyncratic income shocks.[16]

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 (eg 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)).[17] This progress is highly promising. Indeed, among adults that have a

[16] Beyond households, any two agents, including two villages, two regions or two countries, may also benefit from improved risk-sharing.

[17] Another 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

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 (eg 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 (eg 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 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 sceptical 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.

To be sure, the lack of risk-sharing across economic agents encompasses more than financial inclusion in EMDEs. The persistence of the Feldstein-Horioka puzzle indicates that countries can improve risk-sharing compared with the status quo. One could interpret the founding mission of the Bretton Woods institutions (IMF, World Bank), many United Nations specialised agencies (WHO etc), and other regional institutions (eg ESM, AMRO) as helping to foster international risk-sharing. However, many consider the global economy as still falling short of having an effective global safety net (eg Carstens (2021), Adrian et al (2022)). Another interesting example is unemployment re-insurance across the member states of the euro area. Dolls (2019) and Claveres and Stráský (2018), among others, show how ad hoc re-insurance schemes for euro area members may provide a highly effective means of risk-sharing and income smoothing over time and in the cross section, while remaining self-funded and without giving rise to one-way transfers (see also Karaivanov et al (2023)).

We therefore now turn to how insights from mechanism design can be used to help design insurance schemes that are incentive compatible and cost-effective.

while the consumption of non-users dropped by 7%. Users tapped remittances to an extent not available to non-users.

4. Obstacles to risk-sharing

The previous section has shown the benefits of our illustrative schemes for smoothing disposable income with some form of risk-sharing. But there are well known difficulties that lie in the way of participating and cooperating in such schemes. This section makes a distinction among obstacles to trade. 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.

4.1 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 premiums out of high-income states is not guaranteed. It is well known that in static, single-period insurance contracts there is no incentive to pay the premium. In informal systems, as in 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 well known advantage of wider insurance pools.

Use of escrow accounts

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 motivation, we return to our two-agent example. As an approximation, in Table 2 we consider a two-period model in which both agents have similar income incomes in the first period (t), (eg as in period 2 in Table 1a) and their relative positions in the next period (t+1) are random, with one agent having much larger gross income than the other. Thus, each agent would like to plan to receive an insurance indemnity in the second period on a low realised income at the expense of paying a premium on a high realised income. Full insurance, ie, splitting the difference in incomes, would be ideal. However, if there is a limited commitment problem, the high-income agent would not pay.

Ad hoc risk-sharing with an escrow account

In thousands of Thai baht — Table 2

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.

Source: Authors' calculations.

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 2 where the agents put in escrow one quarter of their realised income in period t 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 with the new technologies which allow agents, including strangers, to participate. An individual agent's contribution to the insurance fund is put in escrow. Specifically, a government or trusted third party agrees to hold money contributions in escrow while minting a token that is programmable via smart contract code. Specifically, the savings function is thus automated. Likewise, transfers as indemnities received by the participants in the platform 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, not individual discretion.

To continue with the above two-period example, suppose it is agreed that in period $t+1$, the agent with the lower income gets 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 2, 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.[18] 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 realisations of future income. Notably, the gain from receiving a transfer on 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 understand 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 mutualised 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 the period t income is low, paying a given premium is less likely to be acceptable to the agent. Such

[18] We could similarly apply the same logic starting from the gross incomes in period 8 of Table 1a (as period t), in which case Agent B would have income of 42,000 instead of 12,000 in period t+1.

circumstances will limit insurance. We gain the intuition that a voluntary pre-payment of collateral would depend on the income process of each agent.

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 realised income is relatively low or, alternatively, using accumulated past saved premiums. But unlike in the saving-only setting, here the agents are committed to future payouts from others in low-income states. For example, Dolls (2019) illustrates what such schemes may look like for EU unemployment re-insurance, and see also Section 5.

Exclusion and the benefits of public infrastructure

Rather than requiring collateral or escrow to deal with reneging, another possibility is to exclude agents from future risk-sharing upon failure to follow the rules, ie 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, 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 agents to abide by the rules, ie, they will choose to pay the required premium when income is high.[19]

A real-world example of this type of 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 can potentially act as monopolists and extract rents, in which case much of the benefit from risk-sharing accrues not to the public but rather to the provider.[20] 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 organisations could provide their own not-for-profit insurance platform. Such a platform could implement computer codes that fit the insurance schemes to the income characteristics of a 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 vouched for by a trusted third party, that the code executes as

[19] While such an exclusion can be very effective as 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.

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

intended, ie implementing the agreed sharing rule and state-contingent programmed premiums and indemnities. But the platform itself does not require a trusted third party.

Each participant would have an account on the common ledger, which would be used to pay out premiums in advance or to receive inflows as indemnities. Pre-paid deposits are locked, then tokenised, into e-objects, and thus sequestered for premiums 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 realised. An additional rule can be the exclusion of an agent when the ex post premiums are not paid. This requires a public commitment to such governance. Rules for governance are an important aspect of digital platforms. A hybrid scheme can use both pre-paid collateral and sanctions, hence requiring less collateral.

4.2 Private information

Although digital data allow granular targeting, as noted earlier, the exact situation, eg the income or savings of an agent, may nevertheless remain unknown to others, as an unobserved type or choice 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, ie 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. Further, the code can run offline, not on a blockchain, avoiding costly validation, once the objectives and performance targets are agreed. Private sector entities can utilise 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 3). Suppose each agent contracts not with others individually but through a risk-neutral intermediary/platform absorbing risk as a benchmark case, or as a large mutual fund that pools funds. As a starting point, imagine at period t the possibility for not just individual savings for those with high incomes, but also for individual borrowing for those with low incomes. With underlying income at t, 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 here, there is no additional insurance possible at t+1. Limited commitment in repayment of a loan at t+1 is not considered to be obstacle here, as in this section the focus is instead on private information, but one can use, as earlier, collateral and exclusion, coupled with private information as detailed below.

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

In thousands of Thai baht

Table 3

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 B/L and risk-sharing, we show the effects of reducing the reimbursement by one quarter of the income difference in period t+1.

Source: Authors' calculations.

Such borrowing and lending is, however, not the best information-constrained arrangement. It is better to move partway toward full insurance, to an ad hoc risk-sharing rule. Townsend (1982) shows that a constrained optimal hybrid, combining risk-sharing with borrowing and lending may bring higher welfare to the agents than strict borrowing and lending alone. The essential idea is that, when the agent claims low income, the incoming transfer as indemnity is larger than what a loan would be in the pure borrowing/ lending scheme. Likewise, the amount invested as premium at a claimed high income would be higher than the savings in the pure borrowing/lending scheme. The loan is repaid in the second period but less than 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 that transfers are attenuated. To be incentive-compatible, what happens in one period is tied to future periods. The agent with private information internalises the 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 input into the if/then statements of the code. Which branch is applicable depends on the past history of messages. These are stored on the platform as an encrypted database. There is a common underlying state of what is in each of the participant accounts, consistent with history on the platform, but accounts can be partitioned and kept private with encryption, and with 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, ie with public information on incomes, would predict that if agents insure one another by agreeing that the repayment of loans should depend on the income flow at the period when the loan is repaid. To fix ideas, assume independent and identically distributed income shocks between the two agents, where the agent that draws a high income in t is lending to the agent that draws a low income in t. It is optimal in ex ante terms that, if the lender in period t again has a higher income in period t+1 than the borrower, the repayment of the loan contracted with the borrower at t would thus be reduced by a factor that depends on the preferences and the amount of income risks of the two agents. Likewise, a borrower with lower income at t+1 repays less.

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 with the risk-neutral principal above. An ad hoc example of this for the two agents is reported in Table 3.

The main intuition in Townsend (1982) is that adding risk-sharing to borrowing and lending can be Pareto-improving in ex ante terms. Table 3 provides the intuition. While in period t, it makes sense that agent B borrows from agent A, but paying back this loan in period t+1 could be very costly in utility terms for agent B because their gross income of 18000 is very low. Hence, the possibility of moderating the repayment as a function of realised income in period t+1 is highly preferable.

It might appear that trade for a given agent is 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 with encryption from others and from the platform. Nevertheless, with encryption, appropriate randomisation 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 realisations 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 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. The 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 can be added up, and if low in the aggregate, used as a claim for an incoming village-level indemnity from other villages, or if high, the village would pay a premium.

4.3 A two-period illustration

We 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 example. Suppose income in each period can take two possible values, either y_H=5 (high income) or y_L=1 (low income), each with a probability of one half. For simplicity, suppose agents have log utility of consumption, ln(c) and there is no discounting across periods. Agents maximise their expected utility of consumption for the two periods subject to specific constraints imposed by the financial regime they are in.

We consider a wide 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 intermediary/platform subject to an information constraint, namely the agent's income realisation is unobserved by the platform), *limited commitment* (the agents interact with an intermediary/platform subject to a commitment constraint, namely they cannot commit not to renege on a payment and go to the autarky outcome if 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 non-contingent asset with period return/ gross interest rate equal to 1 (zero net interest).

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

In example units Table 4

Transfers	Period 1		Period 2				Expected utility
	τ_H	τ_L	$\tau_{H,H}$	$\tau_{H,L}$	$\tau_{L,H}$	$\tau_{L,L}$	
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.

Source: Authors' calculations.

Table 4 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. Here a transfer is signed so that positive 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 (eg 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) and the transfers in all periods and states are zero. In contrast, full insurance calls for a premium of –2 whenever income is high and indemnity of +2 whenever income is low, so that consumption equals the expected income (3) in all times and states.

In the saving-only regime, the agent saves when first-period income is high (transfer –1.4 in state H) and then receives back the same saved amount (+1.4) regardless of the income shock 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 low (state L), the agent is unable to smooth consumption and the transfers in states L, LH and LL equal 0.

In the borrowing and savings regime, the agent saves (transfer –1.4) when first-period income is high (received back in period 2, as in saving only) but is also able to borrow (+0.3) when the first-period income is low, which is repaid (transfer –0.3) in the second-period states LH and LL (in the latter state the agent is vulnerable). Commitment to repay 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 second-period income (otherwise an agent would claim the income level yielding a larger transfer); however, they can depend on the first-period income (ie there is history dependence, see Townsend (1982)). There is more smoothing than in the saving-only and the borrowing and savings regimes. Note the opposite signs of first and second-period transfers, although the magnitudes are attenuated in the second period. Specifically, the agent receives a larger indemnity, transfer of +0.9, if first-period income is low (partial insurance), followed by a premium, transfer of –0.2, required in states LH and LL. If the first-period income is high, the agent pays a premium, transfer of –1.8, and receives back +1.1 in second-period states HH and HL. Note that in the hidden income regime the platform can enforce payment in states LH and LL, since we assume that there is no commitment problem, but having to pay a premium in the low-income state LL is hurtful for the agent.

In the limited commitment regime, there is a larger first-period premium (–2.1) than 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) but the agent does receive an indemnity (financed by the large first-period premium) if second-period income is low, including in state LL. The limited commitment regime thus deals with vulnerability for low-income draws in both periods, unlike any of the other constrained regimes.

In the "hidden income + limited commitment" regime, the optimal contract deals with both the asymmetric information and commitment problems at the same time. This naturally results in lower expected utility for the agent than having each obstacle alone, but it still achieves more smoothing and higher expected utility than in the saving-only and the borrowing and saving regimes.

Finally, moral hazard can be included too, by extending the model with an unobserved action (eg effort) affecting the income probabilities. As a result, partial insurance obtains, which reduces effort (the essence of the moral hazard problem) but this trade-off is optimised and there is still a gain in expected utility and consumption-smoothing relative to autarky or saving only.

5. An algorithm for assessing improved risk-sharing schemes

The previous section illustrated with simple examples the benefits and obstacles to risk-sharing. This section describes the Karaivanov and Townsend (2014) approach (henceforth KT) to assessing the practical possibilities of implementing risk-sharing arrangements through data and theory. Specifically, the approach can determine whether existing financial arrangements on the ground are exogenously incomplete, as in the buffer stock savings or borrowing/lending schemes set out in the examples of Section 3, or endogenously incomplete, constrained by various explicit obstacles to trade. We also describe how the approach can be used to quantify the welfare gains from improved risk-sharing platforms in instances where on-the-ground arrangements are limited.

5.1 The Karaivanov-Townsend approach

A highlight of the KT approach is that it 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. 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 (eg 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, whichever is available. Variables at the household level can include consumption or expenditure, income, capital and/or investment. Karaivanov and Townsend (2014) analysed communities of farm and non-farm businesses in rural and urban areas of Thailand. An extension analysed 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 (autarchy, saving-only, and non-contingent debt subject to natural borrowing limit), as well as financial/information environments with obstacles to trade[21] that limit insurance, including constraints stemming from limited commitment, unobserved income or moral hazard with an unobserved action, with the former two discussed in Section 3.[22]

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

[22] More 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

In the underlying economic environment in KT, risk-averse households run SMEs with capital and labour inputs. Income is endogenous and subject to idiosyncratic random shocks. The inclusion of capital recognises that business investment is jointly determined with household consumption in utilising 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 (eg 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,[23] 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 setting in the villages is typically saving-only or non-contingent borrowing and lending, while the best-fitting financial setting in the urban areas is an endogenously constrained setting with information constraints, eg moral hazard. This conclusion mirrors other work that analysed 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 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.

methods allow combinations of obstacles as part of the same environment, for example moral hazard with unobserved output or limited commitment.

23 For the exogenously incomplete markets settings (eg saving-only), one maximises the agents' utility subject to resource and borrowing constraints. For the obstacle-constrained financial/information settings, one maximises 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 actuarily fair insurance.

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.

5.2 Quantifying gains from improved risk-sharing

5.2.1 Rural Thai households

We use the KT approach described in Section 5.1 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 realisation, 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 (see Section 3.2). 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 overindebted in the first place, but all subject to limited commitment (and more

Data summary – Thai rural households without kin

Graph 2

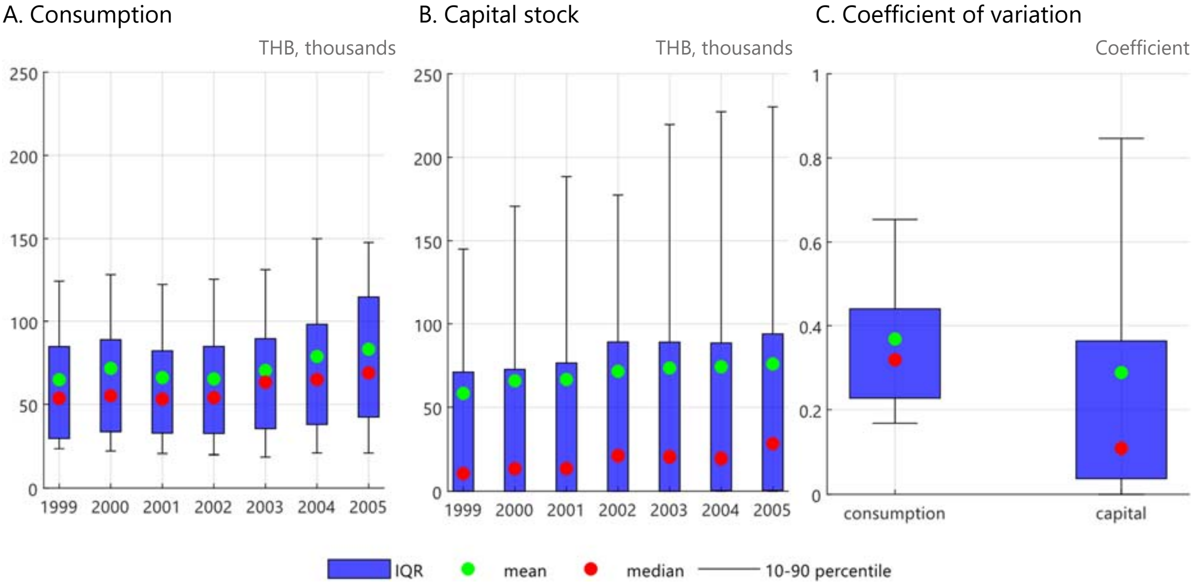

Note: the left-hand and centre panels plot the cross-sectional distribution (inter-quartile range (IQR), mean, median, 10th and 90th percentile) of consumption and capital stock across the 140 households. The right-hand panel plots the cross-household distribution of the coefficient of variation of consumption 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.

Source: Authors' calculations.

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, eg the financial regime that best characterises 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. Graph 2 shows summary statistics for the sample. There is substantial variation in the consumption and capital stock in the cross section of households, for various year, (Panels 2.A and 2.B). Even more telling, there is substantial variation for a household over time (Panel 2.C). The median coefficient of variation of consumption of approximately 0.35 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

Gains from insurance – example time paths

In model units, 1 model unit=179,172 Thai baht Graph 3

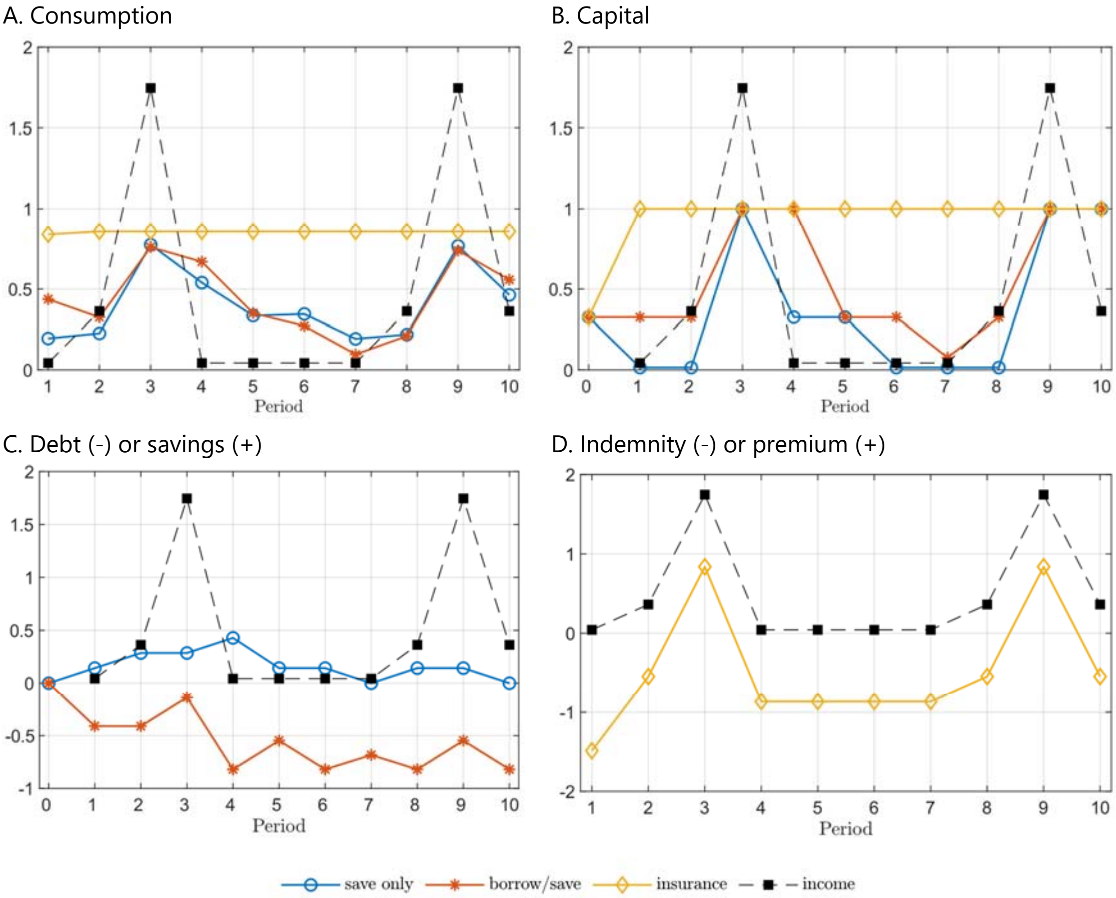

Source: Authors' calculations.

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.

Graph 3 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. (Graph 2). 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.

Graph 3 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, insurance implies an additional gain from production efficiency (and thus in expected income, although in this simulation and the graph the effect of higher capital on income is suppressed).[24] 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 Graph 3 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 2.C) or by paying insurance premiums or receiving indemnity payments (Panel 2.D), depending on realised income in each period. Note that for insurance, households frequently receive indemnities while those who experience high income pay premia.

Graph 4 further explores the poverty trap idea. The insurance platform/setting can help financially constrained households who are unable to invest, because of an inability or restricted ability to borrow. Graph 4 shows four such example households from the model-generated data. They all have low initial capital stock, which in the save-only 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 Graph 4) 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

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

Gains from insurance – overcoming poverty traps

Capital, in model units, 1 model unit=179,172 Thai baht

Graph 4

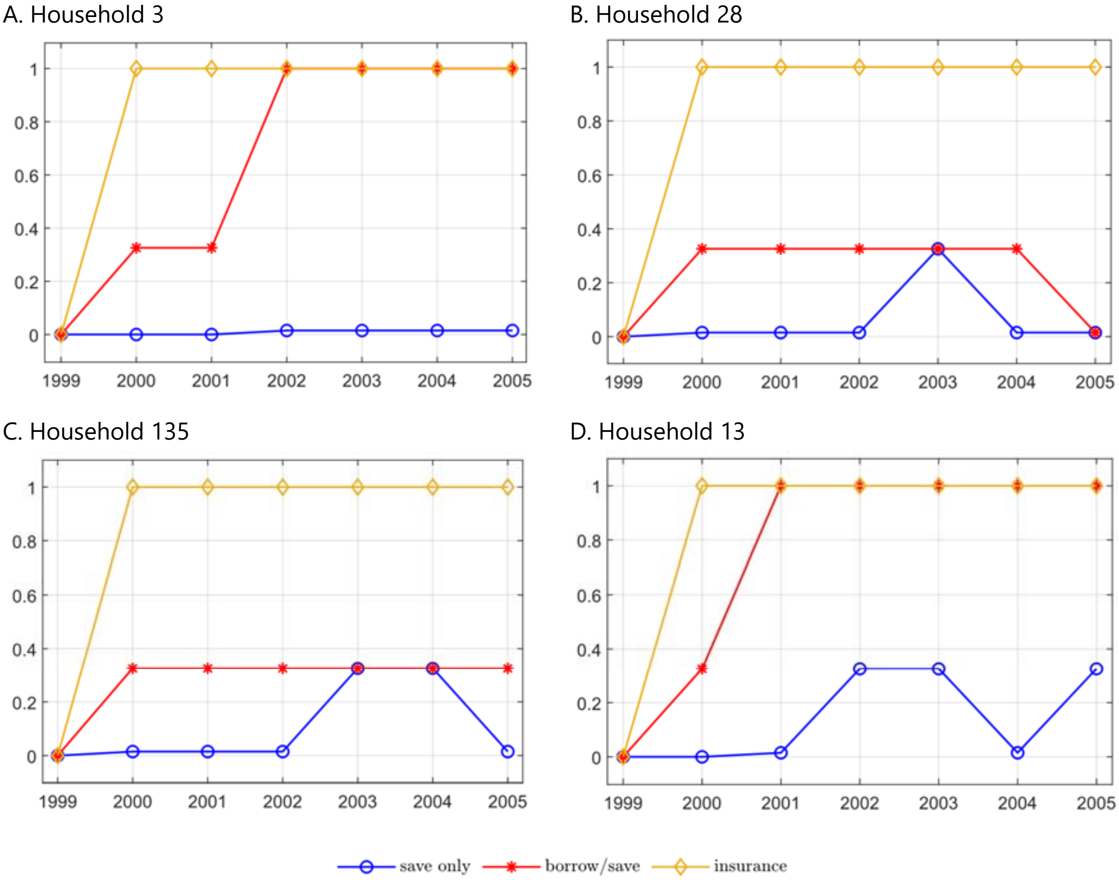

Source: Authors' calculations.

consumption streams). Graph 5 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 graph 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 Graph A3.

In Graph 6 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

Gains from insurance – willingness to pay

In thousands of Thai baht

Graph 5

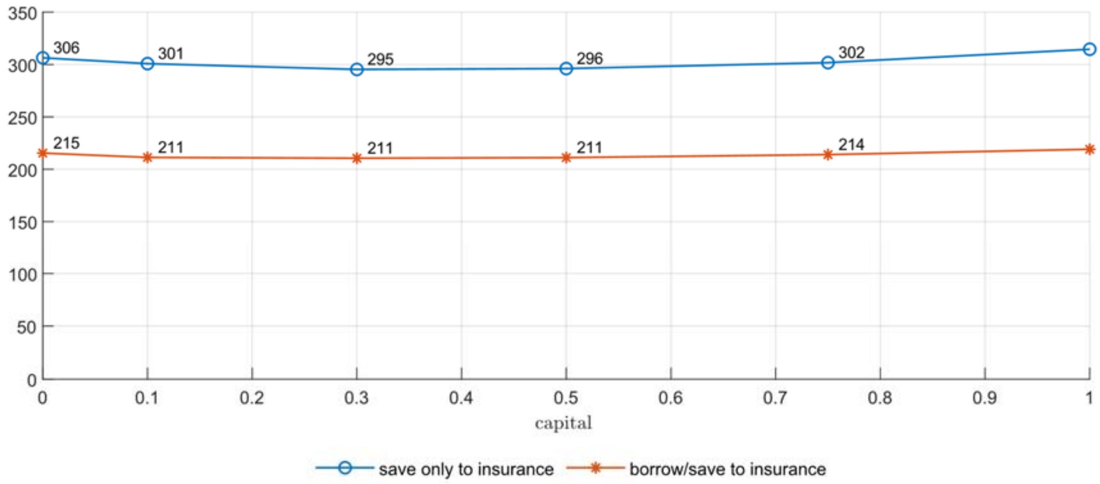

Source: Authors' calculations.

household.[25] Graph 6 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 Graph 3. 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 Graph A1 in the Appendix). Indeed, the median estimated increase in consumption from improved insurance is almost 200% in 1999 and 90% in 2005 for the save-only setting (Panel A1.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.

Graph 7 illustrates the mechanism through which the insurance platform achieves consumption- and investment-smoothing. For the KT parameter estimates

25 In 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.

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

Gains from insurance, consumption and capital-smoothing

In thousands of Thai baht Graph 6

Consumption

A. Saving only B. Saving and borrowing C. Insurance

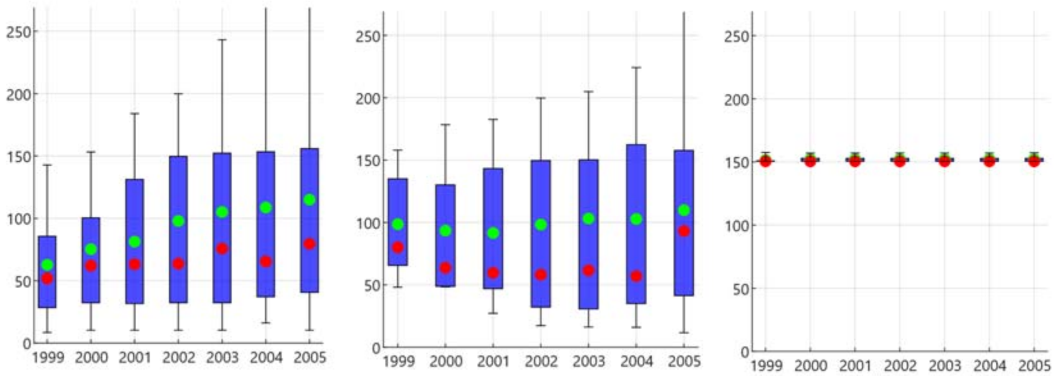

Capital

D. Saving only E. Saving and borrowing F. Insurance

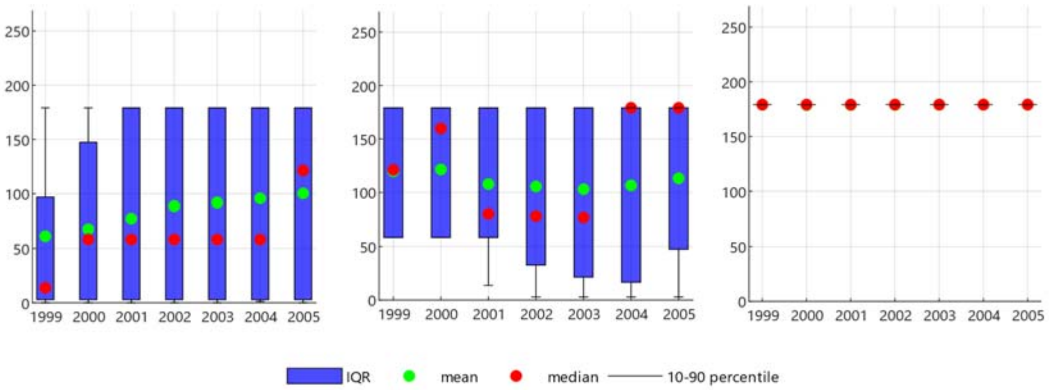

Source: Authors' calculations.

of new investment is undertaken to raise the households' capital stock to the optimal level (see Graph 6). Panel 7.B shows the distribution of the premiums (in dark colour) 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.[26]

[26] The 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 the endogenously modelled income process (see KT). For a *specific* sample of households and income realisation paths, eg such as those featured here, the platform may incur a deficit or surplus.

Insurance premiums and indemnities

Graph 7

A. Insurance indemnity (-) or premium (+)

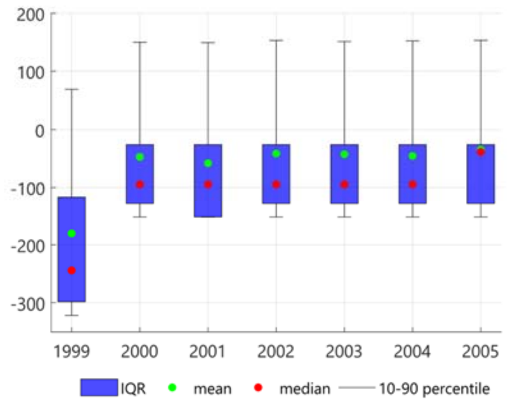

B. Premiums and indemnities by household and year

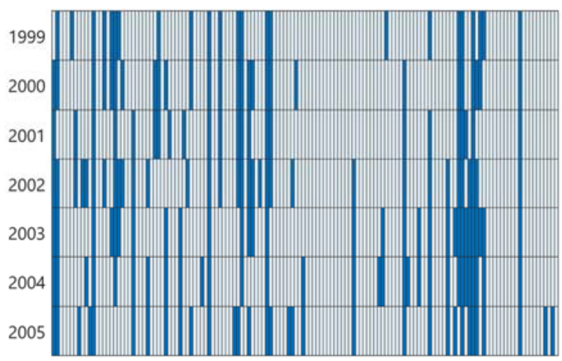

Note: Insurance premium (if positive) or indemnity (if negative) is defined as household income – (consumption + investment). In the right-hand panel, each vertical bar denotes a household; a dark colour means that a household is a net contributor (paying a premium) in a given year; a light colour means that the household is a net receiver (receives an indemnity).

Source: Authors' calculations.

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 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 5.2.1, 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.

Graph 8 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 Graph 6. 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. Graph 8 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 5.2.1 and

Graph 6. 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.

Gains from improved risk-sharing, Spanish firms

Capital, in thousands of euros

Graph 8

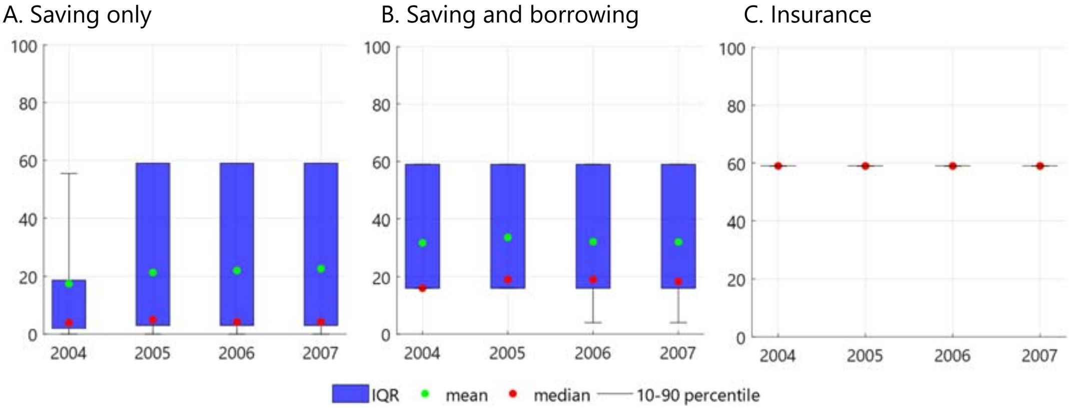

Source: Authors' calculations.

6. Conclusions

We reach four main conclusions. First, the gains from new designs for social safety nets for risk-sharing under the proposed framework can be very large vis-à-vis currently existing social policies and/or reliance on individual savings and absence of pooling risks. These gains would be largest for households and SMEs who have limited family networks and no access to financial services. The gains from sharing macroeconomic risks could also be significant across regions in a country and across countries; see for instance the application of our framework to evaluate an optimal unemployment re-insurance scheme for the euro area in Karaivanov et al (2023). Second, better insurance against income shocks could limit the occurrence of poverty traps and imply a more stable business cycle and less inequality of opportunity and income. Another important result is that such a scheme allows for income-smoothing, which would significantly reduce income insecurity and volatility, in turn contributing to micro and macroeconomic stability. Third, we illustrated how the Karaivanov and Townsend (2014) approach can be used to characterise the financial setting of groups of economic agents and the welfare gains associated with new and improved types of safety net. Finally, we describe how such safety nets can be implemented using financial apps or smart contracts on digital platforms coded to overcome the asymmetric information or commitment obstacles that have, so far, limited the spread of welfare-improving safety nets.

The welfare gains we quantify come from multiple metrics. The first is better smoothing of idiosyncratic shocks, hence less variable individual consumption. Related to this, investment can be based more closely on productivity and made less sensitive to fluctuations in cash flow. In addition, comparing levels with and without the innovation, 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, ie the amount a household is willing to pay to join the insurance platform.

From a broader policy perspective, our results can be used as an important innovation and complement to existing social policy design and as basis for safety nets in a broader, cross-country perspective. Currently, at the national level, in addition to broad social policies, the best type of safety net is generally some form of CCT. It is usually a centralised top-down allocation of public resources to households or firms, which is always subject to cyclical fluctuations and political disputes over fiscal space. Indeed, existing policies often run into moral hazard issues and potential confrontation over entitlements since their financing comes from the overall tax pool.

Even with improved targeting and more information about adequate income thresholds for eligibility, most existing social safety net approaches mitigate risks only when they become systemic in a macro-stabilisation perspective. Despite the downsizing in recent years, with some exceptions, of components of broader social policies, the deployment of additional safety nets has played an important role following large shocks (eg the GFC, Covid-19, the collateral effects of the war in Ukraine). However, these safety nets have limitations. They fail to cover some unexpected idiosyncratic events that could be insured in a decentralised, bottom-up, voluntary and cooperative approach. In addition, despite built-in automatic stabilisers, current safety nets sometimes have implementation lags due to the need to identify macroeconomic triggers to justify their deployment (eg crises, large capital outflows etc) and these lags (and/or the aftermaths of recessions that have lasting effects) might 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 digital systems could complement the current tools with rapid risk-sharing assistance mechanisms that could provide near real-time payments before official transfers are released.

Obviously, 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 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 policy debate, both within the country and at the international level (BIS and 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 minimise the importance of improving the existing social safety nets and do not suggest that the macroeconomic 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 safety nets and their funding in favour of a "risk-sharing only" approach. In fact, given the large potential efficiency gains from risk-sharing, a promising programme would be how to make both types of policy and approach complementary, so that risk-sharing becomes integral part of policy design. The benefits of new technologies would increase economic opportunities for local SMEs and low-income households, which are most frequently subject to the income shocks that feed inequality and volatility.

References

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.Ahlin, C, and R Townsend (2007): "Using repayment data to test across models of joint liability lending", *The Economic Journal*, vol 117, F11–F51.

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*, vol 142, no 103733.Acemoğlu, D and J Robinson (2012): "Why nations fail: the origins of power, prosperity, and poverty", *Finance and Development*, vol 49, no 1.

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*, vol 183, no 1, pp 91–103.

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.

Atkinson, A and B Bourguignon (2015): *Handbook of income distribution*, vol 2, North Holland.

Balboni, C, O Bandiera, R Burgess, M Ghatak and A Heil (2021): "Why do people stay poor?", working paper, London School of Economics.

Banerjee, A, and A Newman (1991): "Risk-bearing and the theory of income distribution", *The Review of Economic Studies*, vol 58, no 2, pp 211–35.

Bank for International Settlements (2017): "Understanding globalisation", *Annual Economic Report*, June, Chapter 6.

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.

Barrientos, A (2007): "Does vulnerability create poverty traps?", University of Sussex, Institute of Development Studies Chronic Poverty Research Centre, Working Papers, no 76.

Benami, E, M Carter (2021): "Can digital technologies reshape rural microfinance? Implications for savings, credit, & insurance", *Applied Economic Perspectives and Policy*, vol 43, no 4, December, pp 1196–220.

Bilbiie, F, G Primiceri and A Tambalotti (2022): "Inequality and business cycles", *Janeway Institute Working Papers*, University of Cambridge.

Bird, K (2007): "The intergenerational transmission of poverty: An overview", Overseas Development Institute Chronic Poverty Research Centre, Working Papers, no 99, October.

Bowles, S, S Durlauf and K Hoff (2006): "Poverty traps", *Economic Development Quarterly*, vol 23, no 2, pp 172–73.

Carstens, A (2021): "Rethinking the global financial safety net", speech at the Sixth High-Level Regional Financing Arrangements (RFA) Dialogue, 12 October 2021.

Carter, M (1997): "Environment, technology, and the social articulation of risk in West African agriculture", *Economic Development and Cultural Change*, vol 45, no 3, pp 557–91.

Challe, E, J Matheron, X Ragot and J Rubio-Ramirez (2017): "Precautionary saving and aggregate demand", *Quantitative Economics*, vol 8, no 2, pp 435–78.

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*, vol 129, no 4, pp 1553–623.

Claveres, G and J Stráský, (2018): "Stabilising the euro area through unemployment benefits re-insurance scheme", *OECD Economics Department Working Papers*, no1497, OECD Publishing.

Dolls, M, (2019): "An unemployment re-insurance scheme for the eurozone? Stabilizing and redistributive effects", *EconPol Policy Reports*, no 10, Ifo Institute – Leibniz Institute for Economic Research at the University of Munich.

D'Silva, D, Z Filková, F Packer and S Tiwari (2019): "The design of digital financial infrastructure: lessons from India", *BIS Papers*, no 106.

Durlauf, S (1994): "Spillovers, stratification, and inequality", *European Economic Review*, no 38, pp 836–45.

Forbes, K and F Warnock (2012): "Capital flow waves: Surges, stops, flight, and retrenchment", *Journal of international economics*, vol 88, no 2, pp 235–51.

Goldfarb, A and C Tucker (2019): "Digital Economics", *Journal of Economic Literature,* vol 57, no 1, March, pp 3–43.

Ghatak, M (2015): "Theories of poverty traps and anti-poverty policies", *The World Bank Economic Review*, vol 29(S1), pp S77–S105.

Guvenen, F, L Pistaferri and G Violante (2022): "Global trends in income inequality and income dynamics: New insights from GRID", *Quantitative Economics*, vol 13, pp 1321–60.

Heathcote, J, F Perri and G Violante (2020): "The rise of US earnings inequality: does the cycle drive the trend?", *Review of Economic Dynamics*, vol 37, August, pp 181–204.

Jack, W, and T Suri (2014): "Risk sharing and transactions costs: evidence from Kenya's mobile money revolution", *American Economic Review*, vol 104, no 1, pp 183–223.

Jalan, J and M Ravallion (1998): "Transient poverty in postreform rural China", *Journal of Comparative Economics*, vol 26, no 2, pp 338–57.

Ji, Y, S Teng and R Townsend (2022): "Dynamic bank expansion: spatial growth, financial access, and inequality", working paper, MIT.

Karaivanov, A and R Townsend (2014): "Dynamic financial constraints: distinguishing mechanism design from exogenously incomplete regimes", *Econometrica*, vol 82, pp 887–959.

Karaivanov, A, J Saurina and R Townsend (2019): "Family firms, bank relationships and financial constraints: a comprehensive scorecard", *International Economic Review*, vol 60, pp 547–93.

Karaivanov, A, B Mojon, L Pereira da Silva, A Pierres Tejada and R Townsend (2023): "Incentive-compatible unemployment reinsurance for the euro area", *BIS Working Papers*, (forthcoming).

Koepke, R (2019): "What drives capital flows to emerging markets? A survey of the empirical literature", *Journal of Economic Surveys*, vol 33, no 2, pp 516–40.

Kinnan, C, K Samphantharak, R Townsend and D Vera Cossio (2020): "Propagation and insurance in village networks", *NBER Working Papers*, no 28089.

Kraay, A and D McKenzie (2014): "Do poverty traps exist? Assessing the evidence", *Journal of Economic Perspectives*, vol 28, no 3, pp 127–48.

Latner, J (2019): "Economic insecurity and the distribution of income volatility in the United States", *Social Science Research*, vol 77, pp 193–213.

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*.

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, pp 593–696.

Nilekani, N (2018): "Data to the people: India's inclusive internet", *Foreign Affairs*, vol 97, no 19.

Paulson, A, R Townsend and A Karaivanov (2006): "Distinguishing limited liability from moral hazard in a model of entrepreneurship", *Journal of Political Economy*, vol 114, no 1, February, pp 100–44.

Pereira da Silva, L, E Kharroubi, E Kohlscheen, M Lombardi and B Mojon (2022): "Inequality hysteresis and the effectiveness of macroeconomic stabilisation policies", BIS, May.

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*, vol 60, no 1, May, pp S159–76.

Rosenzweig, M, and H Binswanger (1993): "Wealth, weather risk and the composition and profitability of agricultural investments", *The Economic Journal*, vol 103, no 416, pp 56–78.

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.

Samphantharak, K and R Townsend (2018): "Risk and return in village economies", *American Economic Journal: Microeconomics*, vol 10, no 1, February, pp 1–40.

Townsend, R (1979): "Optimal contracts and competitive markets with costly state verification", *Journal of Economic Theory*, vol 21, no 2, October, pp 265–93.

——— (1982): "Optimal multiperiod contracts and the gain from enduring relationships under private information", *Journal of Political Economy*, vol 90, no 6, pp 1166–86.

——— (2016): "Village and larger economies: The theory and measurement of the Townsend Thai project", *Journal of Economic Perspectives*, vol 30, no 4, pp 199–220.

Xuan Thanh, N (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.

Western, B, D Bloome, B Sosnaud and L Tach (2012): "Economic insecurity and social stratification", *Annual Review of Sociology*, vol 38, pp 341–59.

World Bank (2021): Global Findex Database: Financial Inclusion, Digital Payments, and resilience in the Age of COVID-19.

Appendix

Gains from insurance – consumption growth

Consumption gains, in percent | Graph A1

A. Saving only to insurance | B. Saving and borrowing to insurance

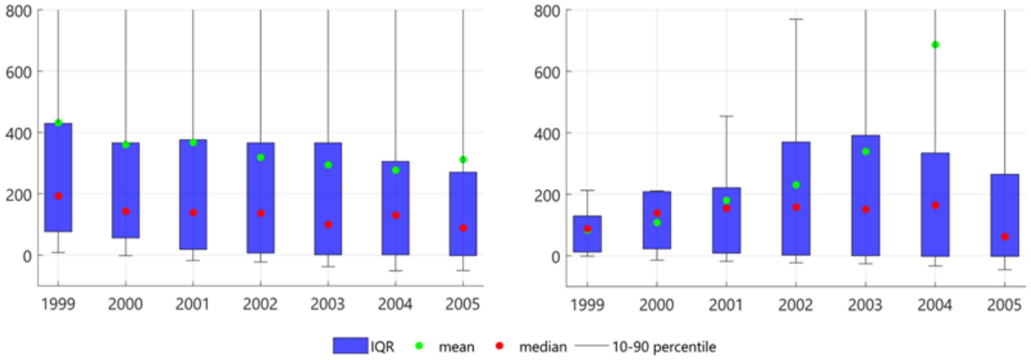

Note: the graph plots the distribution of consumption growth, defined as the ratio of the consumption level in the insurance setting to the consumption level in the restricted setting (save-only or borrow/save) computed for each household and year.

Source: Authors' calculations.

Gains from insurance – capital accumulation

In thousands of Thai baht | Graph A2

A. Saving only | B. Saving and borrowing | C. Insurance

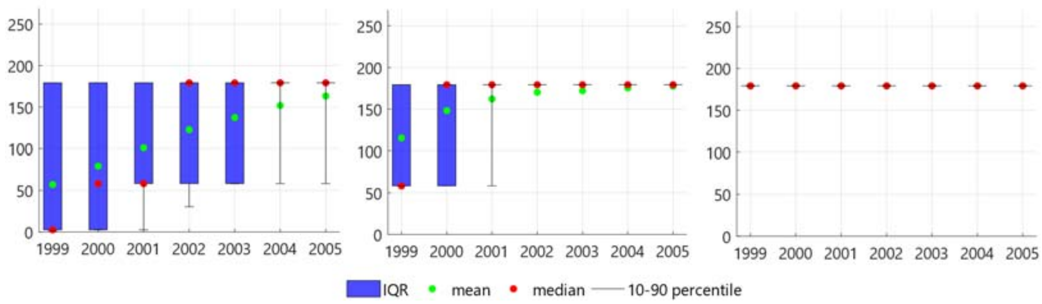

Note: the graph plots the distribution of end-of-period capital stock for each of the three settings (save only, borrow/save and insurance). Each household is initialised at its 1999 capital from the data, zero debt/savings and the ex ante utility yielding zero insurer profits. The income process is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

Source: Authors' calculations.

Willingness to pay for insurance (robustness)

In thousands of Thai baht | Graph A3

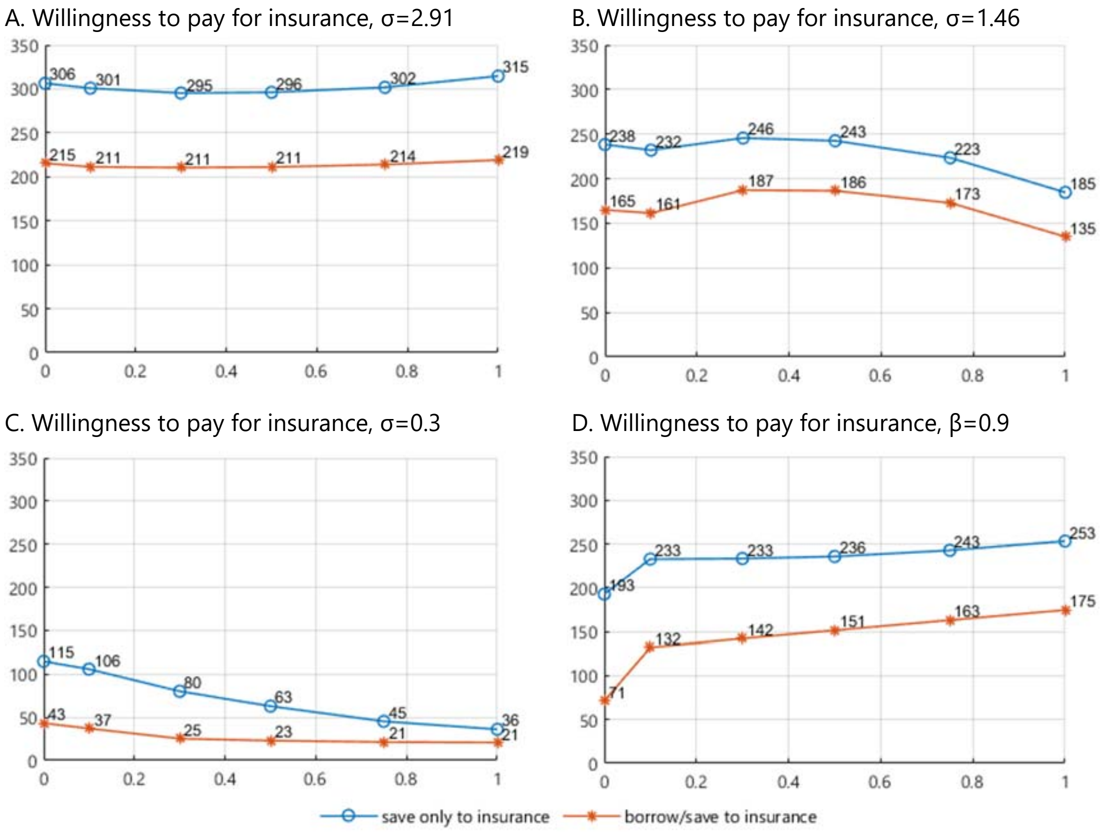

Note: we plot the agents' willingness to pay to move to the insurance setting, as in Graph 5, but computed for different values for risk aversion in the agents' preferences, the baseline estimate (panel A) or lower values (panels B and C) or for smaller discount factor (panel D).

Source: Authors' calculations.

Previous volumes in this series

No	Title	Issue date
BIS Papers No 138	Financial stability risks from cryptoassets in emerging market economies	August 2023
BIS Papers No 137	Building an integrated surveillance framework for highly leveraged NBFIs – lessons from the HKMA	July 2023
BIS Papers No 136	Making headway – Results of the 2022 BIS survey on central bank digital currencies and crypto	July 2023
BIS Papers No 135	The energy transition and its macroeconomic effects	May 2023
BIS Papers No 134	Global tightening, banking stress and market resilience in EMEs	April 2023
BIS Papers No 133	The two-regime view of inflation	March 2023
BIS Papers No 132	Information governance in sustainable finance	December 2022
BIS Papers No 131	Central banking after the pandemic: challenges ahead	December 2022
BIS Papers No 130	Pricing of climate risks in financial markets: a summary of the literature	December 2022
BIS Papers No 129	The role of non-bank financial institutions in cross-border spillovers	December 2022
BIS Papers No 128	Central bank digital currencies in Africa	November 2022
BIS Papers No 127	Historical monetary and financial statistics for policymakers: towards a unified framework	September 2022
BIS Papers No 126	Corporate digital identity: no silver bullet, but a silver lining	June 2022
BIS Papers No 125	Gaining momentum – Results of the 2021 BIS survey on central bank digital currencies	May 2022
BIS Papers No 124	The design of a data governance system	May 2022
BIS Papers No 123	CBDCs in emerging market economies	April 2022
BIS Papers No 122	The monetary-fiscal policy nexus in the wake of the pandemic	March 2022
BIS Papers No 121	Covid-19 and the monetary-fiscal policy nexus in Africa	February 2022
BIS Papers No 120	Virtual banking and beyond	January 2022

All volumes are available on the BIS website (www.bis.org).