

Liquid Supply Chains

How Falling Underwriting Costs Could Turn World Trade into an Asset Class, and What That Would Mean for Growth

Christoph Gugelmann (Rhofin Inc.)¹

First draft: June 2026. This version: 11 June 2026. Comments welcome.

Disclosure of interest: the author is the founder and chief executive officer of Rhofin Inc., a commercial enterprise whose business is the subject of the proposal in Section 7. See the author note.

This paper is for research discussion only and is not an offer to sell securities, an investment recommendation, or legal advice.

Abstract

World trade rests on one of the safest large asset classes in finance, yet US\$2.5 trillion of annual demand for trade finance goes unmet, and the burden of the gap falls disproportionately on small and medium-sized enterprises and emerging-market firms. We argue that this is, in substantial part, a unit cost problem rather than a risk problem: the fixed cost of verifying, monitoring and documenting a shipment has historically exceeded the revenue available on small transactions, so rational lenders exclude the very borrowers whose marginal returns to capital are highest. We acknowledge that the headline loss data, being conditional on approval, cannot by themselves rule out the rival adverse-selection explanation, and we specify the test that discriminates between the two. Three concurrent shifts now change the cost structure of safe origination: agentic AI that lowers the marginal cost of underwriting and compliance, legal reforms that make digital trade documents enforceable collateral, and the digitization of logistics data. Together they make it possible to finance, collateralize and distribute trade exposures at the level of the individual shipment, a state we describe as a liquid supply chain. We set out the mechanism by which automation converts unpriceable operational risk into investable credit risk, trace the consequences for firms, trade and capital markets, and examine the principal failure modes, including model monoculture, the screening incentive failures familiar from past securitization booms, and the lessons of the 2021 Greensill collapse. We close with a measurable research agenda and an invitation to externally funded research partnerships to design the evaluation of a new origination channel before its first transaction.

Keywords: trade finance; working capital; SME credit; securitization; artificial intelligence; agentic AI; electronic trade documents; supply chains; credit rationing; financial technology

JEL classification: D82, F14, F36, G21, G23, O33

1. Introduction: the illiquidity of physical trade

A container of goods crossing the Pacific is, economically, a bundle of near-certain future cash flows. The goods exist, their movement is tracked in real time, the buyer is often an investment-grade retailer, and the historical loss rate on financing such flows is a fraction of one percent. Yet for most of the firms that ship them, those cash flows are frozen: the working capital embedded in goods in transit and invoices awaiting payment cannot be borrowed against, sold, or pledged at reasonable cost. A typical small exporter pays its supplier at or before shipment, waits roughly 45 days for the goods to move from factory gate to destination port on the main east-west lanes, delivers, and is then paid 60

¹ Christoph Gugelmann is the founder and chief executive officer of Rhofin Inc., a commercial enterprise developing AI-enabled trade-finance origination infrastructure. The author accordingly has a direct commercial and financial interest in the mechanisms this paper describes and in the conclusions it advances, and the paper should be read in light of that interest; the analysis and any errors are the author's own. Section 7 contains a standing invitation to externally funded research partnerships to design and conduct the evaluation of the paper's claims from the first transaction of a new origination channel, on terms guaranteeing full independence of analysis and publication. Correspondence: research@rhofin.com.

to 90 days after delivery on standard open-account terms. It is out of pocket for roughly 100 to 135 days or more on every order, and its growth is capped by cash rather than by demand.²

The central claim of this paper can be stated in one sentence: if verification costs fall while error rates do not rise, the minimum viable ticket size of trade finance falls, and the trade finance gap becomes experimentally testable as a unit-cost problem rather than a risk problem. Everything that follows unpacks that sentence: the evidence that motivates it, the mechanism that would carry it, the failure modes that could defeat it, and the measurements that would decide it.

We use the term *liquid supply chain* to describe the opposite state: one in which every eligible, verifiable shipment and invoice can be financed at the moment it is created, collateralized through enforceable digital documents, and distributed to capital markets in standardized form, much as a mortgage or a credit card receivable can be today. This paper asks what would happen to firms, to trade, and to the financial system if supply chains became liquid, and why the question has only recently become more than hypothetical.

The argument proceeds as follows. Section 2 documents a puzzle: an asset class with exceptionally low realized losses that capital nonetheless refuses to fund at scale. We state the puzzle together with its principal rival explanation, adverse selection, and explain why the two cannot be separated on existing data. Section 3 identifies the three converging shifts, technological, legal and informational, that change the underlying cost structure. Section 4 sets out the economic mechanism: the conversion of operational risk into priceable credit risk, and the analogy, with its important limits, to the standardization episodes that liquefied mortgages and containerized freight. Section 5 traces the consequences through firm, trade and capital-market channels, and Section 6 confronts the risks, including the one recent, directly relevant failure: Greensill. Throughout, we treat the claim that AI changes trade finance as a hypothesis to be tested, and we specify the measurements that would falsify it. Section 7 closes with the research agenda and an explicit invitation to researchers to design those tests before the data exist.

2. The puzzle: a safe asset that capital will not fund

The Asian Development Bank estimates the global trade finance gap, defined as demand for trade financing that financial institutions decline to meet, at US\$2.5 trillion per year, unchanged from 2023 and up from about US\$1.5 trillion in 2015.³ The gap is approximately 10 percent of global merchandise trade,⁴ and it is concentrated where it does the most damage: 41 percent of trade-finance applications from small and medium-sized enterprises are rejected,⁵ even though SMEs constitute roughly 90 percent of the world's businesses and more than half of global employment.⁶ The most recent survey reports, for the first time, SME rejection rates close to those of large corporates, a convergence the ADB cautions is provisional; even taken at face value, it does not dissolve the

² Illustrative of standard market terms. Door-to-port times on the major east-west lanes average roughly 45 days: Flexport's Ocean Timeliness Indicator reports China-US West Coast at 37 days, China-US East Coast at 53.3 days and China-North Europe at 58.4 days for the week to 23 March 2026 (Flexport, *Ocean Timeliness Indicator*, 26 March 2026, flexport.com/research). Suppliers are frequently paid at or before shipment, and open-account terms of 60 to 90 days after delivery are common; on the prevalence of post-shipment open-account terms see the FCIB survey evidence in Antràs and Foley, n. 23 below, whose own transaction data show shorter terms (7 to 45 days after arrival), so the 60-to-90-day illustration should not be attributed to that source.

³ Beck, S., A. DiCaprio, A. Pandey and A. Malaket, "ADB Global Trade Finance Gap Survey," *ADB Briefs* No. 378 (Manila: Asian Development Bank, December 2025), DOI: 10.22617/BRF250557-2. The survey of over 110 trade-finance providers estimates global unmet demand at US\$2.5 trillion in 2025, unchanged from 2023 (pp. 1-2); ADB first estimated the gap at about US\$1.5 trillion in 2015 (p. 1).

⁴ ADB Briefs No. 378, n. 3 above, p. 2: the gap represents around 10 percent of global merchandise trade flows in this 2025 edition, down from about 10.6 percent in 2023; the survey uses merchandise trade as the denominator for consistency across editions (p. 2, n. 7 therein).

⁵ ADB Briefs No. 378, n. 3 above, p. 5 and Figure 2: SME rejection rate of 41 percent (40 percent as charted), down from 45 percent in 2023 and for the first time close to the 40 percent recorded for large and mid-cap corporates; the ADB cautions the convergence may partly reflect discouraged demand and warrants continued monitoring and research.

⁶ World Bank, "Small and Medium Enterprises (SMEs) Finance," worldbank.org, n.d.: SMEs represent around 90 percent of all businesses and account for more than half of global employment; in emerging economies formal SMEs contribute up to 40 percent of GDP.

puzzle, because the gap’s absolute incidence remains concentrated among SMEs and emerging-market firms, and because the argument of this paper concerns the minimum viable ticket size rather than differential rejection rates as such.

What makes the gap remarkable is the risk profile of the asset being rationed. The ICC Trade Register, the most comprehensive dataset on trade-finance performance, covering more than US\$25 trillion of exposures, reports default rates persistently below 0.3 percent across major products, with exposure-weighted loss rates on letters of credit of the order of ten basis points or less.⁷ Trade finance is short-tenor, self-liquidating and secured by goods with observable market value. By the standards of corporate or consumer credit, it is among the safest lending banks do.

Three caveats must discipline this evidence before anything is built on it. First, and most importantly, the Trade Register’s loss rates are *conditional on approval*. Low realized losses on the approved book are exactly what an effective screening process would produce if rejected applicants were unobservably riskier than approved ones; the data are therefore consistent both with the unit-cost explanation advanced here and with a classical adverse-selection account in which the rejected 41 percent is rejected because it deserves to be. The two hypotheses cannot be separated on the existing public data. They can, however, be separated by experiment, because they make opposite predictions about the same observable: if the gap is a unit-cost phenomenon, segments that become bankable when verification costs fall should perform at loss rates comparable to the incumbent book; if it is adverse selection, losses in the newly approved segments should rise sharply toward the level that justified their exclusion. Section 7 makes this the central discriminating measurement of the research agenda. Second, the lowest headline loss figures belong to letters of credit, which are bank-intermediated and doubly secured; the unmet demand is concentrated in open-account receivables and pre-shipment finance, products with thinner documentation and a different fraud profile, and the inference from one product’s loss history to another’s must be made with care. Third, the gap estimate itself is a survey measure of self-reported rejected and discouraged demand, with no price attached; some fraction of any such gap is demand that is unviable at any compliant price, and the US\$2.5 trillion figure should be read as an upper bound on the addressable problem rather than a market size.⁸

With those caveats stated, the unit-cost hypothesis runs as follows. Originating a compliant trade-finance exposure requires document verification, sanctions and trade-based money-laundering screening, counterparty due diligence, collateral perfection and ongoing monitoring. These costs are largely fixed per transaction: in ADB’s 2016 survey, 90 percent of banks cited anti-money-laundering and know-your-customer compliance as an impediment to expanding trade finance, and the constraint has persisted, with more than 70 percent of banks in the 2021 survey still ranking AML/KYC requirements as the single biggest hindrance; the 2025 survey reports funding and macroeconomic constraints overtaking KYC at the top of the ranking, but compliance concerns recur survey after survey and reducing compliance costs remains among the measures banks would most prioritize to narrow the gap.⁹ When the fixed cost of safe origination exceeds the fee income available on a US\$50,000 shipment, declining that shipment is the rational decision, whatever its credit quality. The

⁷ International Chamber of Commerce, *ICC Trade Register Report 2024* (Paris: ICC, with Global Credit Data and Boston Consulting Group, 2024), analyzing 2023 data from 22 banks covering over 52 million transactions and more than US\$25.7 trillion of exposures, approximately 16 percent of traditional global trade-finance flows. The 2025 Market Commentary reports overall default rates below 0.3 percent across major trade-finance products; exposure-weighted loss rates reported for letters of credit are of the order of 0.10 percent of exposure (import) and 0.02 percent (export).

⁸ The gap is constructed from provider-reported rejections and discouraged demand, not from observed demand at a market-clearing price (methodology: ADB Briefs No. 378, n. 3 above, pp. 1-2). It indicates the scale and incidence of rationing, not an addressable market; the share financeable at compliant cost and acceptable risk is the empirical question this paper proposes to measure.

⁹ Asian Development Bank, *2016 Trade Finance Gaps, Growth, and Jobs Survey* (Manila: ADB, 2016): 90 percent of surveyed banks cited anti-money-laundering and know-your-customer requirements as impediments to expanding trade finance, and 77 percent cited Basel III requirements. ADB, *Toward Inclusive Access to Trade Finance* (Manila: ADB, 2022): in the 2021 survey more than 70 percent of banks ranked AML/KYC requirements the biggest hindrance and around 60 percent the Basel requirements. The 2025 survey (n. 3 above, p. 6) reports funding constraints (ranked first by over 22 percent) and macroeconomic and counterparty risk displacing KYC concerns (ranked first by just over 15 percent) at the top, while compliance issues recur survey after survey and “reduce compliance costs” remains among the top measures respondents would prioritize to narrow the gap (pp. 6, 12-13 and Figure 8).

gap, on this reading, is not mispriced risk; it is a minimum viable ticket size, and the SMEs below it are excluded by arithmetic rather than by riskiness.

The same cost structure would explain a second anomaly: the near-total absence of capital markets from working-capital finance. Mortgages, which are long dated, complex and historically loss prone, are funded substantially through securitization, while trade receivables and in-transit inventory, despite shorter duration and lower losses, remain overwhelmingly on bank balance sheets.¹⁰ Global factoring turnover reached €4.0 trillion in 2025, but it is dominated by bank-affiliated factors in Europe; North America accounts for roughly 4 percent of world volume.¹¹ Assets cannot be distributed to investors who cannot verify them, and verification has been too expensive to perform at the level of the individual small exposure. Heterogeneous, paper-based, operationally opaque assets do not securitize.

3. What changed: three converging shifts

3.1 Technology: the marginal cost of underwriting falls

The diligence that prices small-ticket finance out of existence consists of reading and cross-checking bills of lading, invoices and packing lists, screening counterparties, detecting trade-based money laundering typologies, assessing credit, and assembling an audit trail. This is structured, document-heavy, rule-governed work of the kind that contemporary agentic AI systems perform well, with human review reserved for exceptions. The economic significance is not that machines do this work faster, but that they convert a fixed labor cost per transaction into a much lower, scalable marginal cost. If the cost of compliant origination falls by an order of magnitude, the minimum viable ticket size falls with it, and a large fraction of the rejected 41 percent becomes commercially financeable at unchanged risk appetite.

An important qualification narrows this claim. The binding constraint banks report is not only the labor cost of compliance but the liability attached to compliance failure: sanctions and money-laundering regimes impose penalties that are largely invariant to how cheaply the screening was performed, and a bank declining a small shipment is often pricing the tail cost of a compliance error rather than the hours of an analyst. Automation therefore relaxes the constraint only if it reduces *error rates*, not merely labor input. The proposition that machine verification can match or exceed human screening accuracy at scale is plausible, the closest empirical analogue being the fintech-lending literature, where underwriting on non-traditional digital data has been shown to predict default as well as or better than traditional credit information,¹² an analogue that bears on the information content of non-traditional data rather than directly on the error rates of automated versus human compliance screening, but it has not been demonstrated in audited trade-finance production, and we do not assume it. Practitioner expectation points the same way: in the ADB's 2025 survey, over 85 percent of banks saw potential for AI in fraud prevention and risk management and a majority were assessing AI as a means of increasing financing capacity (n. 3 above, pp. 10-11), though this measures expectation rather than demonstrated performance. It is the first item on the measurement agenda of Section 7. Stated precisely, the technological claim of this paper is conditional: *if* automated origination can

¹⁰ Mortgages, auto loans and card receivables are funded substantially through capital markets. We are not aware of a rigorous comparative estimate of securitization intensity across these classes and trade receivables; the qualitative contrast is consistent with the small observed volume of trade-receivables securitization relative to the outstanding stock, and producing a rigorous estimate is part of the research agenda of Section 7.

¹¹ FCI (Factors Chain International), *World Factoring Statistics 2025* (Amsterdam: FCI, May 2026): global factoring and receivables-finance turnover of €4,039 billion in 2025, of which Europe accounted for approximately two-thirds (65.8 percent) and North America €160 billion, roughly 4 percent.

¹² Berg, T., V. Burg, A. Gombović and M. Puri (2020), "On the Rise of FinTechs: Credit Scoring Using Digital Footprints," *Review of Financial Studies* 33(7): 2845-2897, showing that simple digital-footprint variables equal or exceed the information content of credit-bureau scores for default prediction. The setting is itself short-term trade credit (goods shipped before payment, claims maturing in roughly 3.5 months): the digital-footprint model approximately matches the bureau score (AUC 69.6 versus 68.3 percent, significant only at the 10 percent level), and because the two information sets are almost uncorrelated, the combined model (AUC 73.6 percent) outperforms either alone. Their sample is itself conditional on approval (orders pre-screened to a predicted default rate below 10 percent), the same conditioning caveat Section 2 applies to the Trade Register.

achieve compliance and fraud-detection performance at parity with incumbent processes at an order of magnitude lower cost, *then* the minimum viable ticket size collapses. The conditional is testable within months on live transaction flow.

3.2 Law: digital documents become enforceable collateral

Cheap underwriting is worthless if the resulting claim cannot be enforced. Until recently, the legal architecture of trade ran on paper: possession of a paper bill of lading was constructive possession of the goods, and electronic substitutes had uncertain status. That has changed since 2021 across the major common-law trading jurisdictions, with the European Union’s contribution phasing in on a longer timetable. The United Kingdom’s Electronic Trade Documents Act 2023 gives electronic bills of lading and bills of exchange the same legal effect as paper,¹³ implementing the principles of the UNCITRAL Model Law on Electronic Transferable Records that Singapore and a growing list of jurisdictions have also adopted.¹⁴ In the United States, control of an electronic bill of lading has been recognized since the 2003 revision of UCC Article 7 (§ 7-106), and the 2022 amendments to the Uniform Commercial Code, now enacted in 33 jurisdictions including the District of Columbia, refine that control framework and extend perfection by control to a wider class of digital assets,¹⁵ while UCC § 9-313 permits perfection by possession through a bailee, so that the party controlling the bill of lading holds collateral that physically travels with the shipment.¹⁶ In the European Union, the ViDA package adopted in March 2025 will make structured e-invoicing and near-real-time digital reporting mandatory for intra-EU B2B transactions from July 2030, turning the invoice itself into a structured data object reported to tax authorities in near-real time, on a known statutory timetable.¹⁷ Collectively these reforms do for trade claims what land registries did for mortgages: they make the collateral legible, transferable and enforceable at low cost. The UK and Singapore reforms are in force today; the direction of travel elsewhere is dated and legislated rather than speculative.

3.3 Data: underwriting on observed flows rather than reported financials

The third shift is informational. Transport and freight management systems now generate continuous, third-party-verifiable telemetry on the physical movement of goods. A lender that underwrites on shipment data observes the trade itself, from booking and vessel to milestone events and proof of delivery, rather than relying on the borrower’s self-reported financial statements. This matters most for fraud: the classic failure mode of receivables finance is the fictitious invoice, a claim on a trade that never happened.¹⁸ An invoice anchored to an observed, independently documented shipment is

¹³ Electronic Trade Documents Act 2023 (UK), in force 20 September 2023. The Act gives qualifying electronic trade documents, including electronic bills of lading and bills of exchange, the same legal treatment, effect and functionality as their paper equivalents. See Law Commission of England and Wales, *Electronic Trade Documents: Report and Bill* (Law Com No 405, 2022).

¹⁴ UNCITRAL, *Model Law on Electronic Transferable Records* (2017). Singapore implemented MLETR via the Electronic Transactions (Amendment) Act 2021; a growing list of jurisdictions has followed.

¹⁵ Uniform Law Commission and American Law Institute, *Amendments to the Uniform Commercial Code* (2022), creating Article 12 (controllable electronic records), refining the § 7-106 control rules for electronic documents of title, and conforming Article 9 amendments on perfection in digital assets. With New York’s December 2025 enactment (effective June 2026), the amendments are adopted in more than 30 jurisdictions, including Delaware and the District of Columbia. An electronic bill of lading is an electronic document of title governed by § 7-106, in place since the 2003 revision of Article 7 and enacted in virtually every state, and is expressly excluded from the Article 12 definition (§ 12-102(a)(1)); the 2022 amendments refine rather than create the control framework for trade documents.

¹⁶ Where the bill of lading is electronic, UCC § 7-106 defines control of an electronic document of title; a security interest in goods covered by a negotiable document may be perfected by perfecting a security interest in the document (§ 9-312(c)), and perfection in an electronic document is achieved by control (§ 9-314). Where goods are held by a bailee outside a negotiable document, § 9-313 permits perfection by possession, including through the bailee’s authenticated acknowledgment that it holds for the secured party. In either case the financier that controls the document, or holds the acknowledgment, has a perfected interest in collateral that travels with the shipment.

¹⁷ Council of the European Union, VAT in the Digital Age (ViDA) package, formally adopted 11 March 2025 and entered into force 14 April 2025; structured e-invoicing and near-real-time digital reporting become mandatory for intra-EU B2B transactions from 1 July 2030, with national systems converging by 2035. European Commission Newsroom, 4 April 2025.

¹⁸ On fictitious and “prospective” receivables as the central fraud typology in receivables finance, see House of Commons Treasury Committee, *Lessons from Greensill Capital*, n. 35 below. On trade credit generally: Petersen, M. A., and R. G. Rajan (1997), “Trade Credit: Theories and Evidence,” *Review of Financial Studies* 10(3): 661-691.

structurally harder to fabricate. The precedent for underwriting on observed behavior rather than reported financials is by now well established empirically: digital-footprint and cash-flow data have been shown to predict consumer and small-business default at least as well as bureau data.¹⁹ Logistics data extends the same principle to trade credit, with the added advantage that the underlying events are physical and externally observable.

4. The mechanism: converting operational risk into investable credit risk

Why would these shifts liquefy supply chains rather than merely cheapen bank lending? The answer lies in decomposing the risk a financier of small-ticket trade actually bears. Only part of it is credit risk, the probability-weighted loss from non-payment, which the ICC data show to be small and which investors can price. The remainder is operational tail risk: fraud, document failure, legal unenforceability, servicing error. This component is not merely large; it is unquantifiable, because it depends on the integrity of processes the investor cannot observe. A first-loss investor cannot price what it cannot quantify, and so declines, whatever the headline default statistics say.

Automation and digital collateral attack exactly this component. When every document is machine-verified at source, every decision is logged in an auditable trail, every exception is escalated to human review under a defined protocol, and the financier controls an enforceable electronic bill of lading so that cargo releases only on repayment, the operational tail is engineered down rather than priced in. It is not engineered away; Section 6 discusses the adversarial response. What remains, if the engineering works, is predominantly quantifiable credit risk on a short-duration, self-liquidating asset, which is the raw material of securitization. The sequence is mechanical: standardized origination produces a seasoned performance record; the record supports forward-flow purchases by institutional investors; sufficient seasoning supports rated asset-backed issuance. Trade exposures then trade the way mortgage exposures trade.

It is worth being concrete about the legal architecture the mechanism presupposes, because the claim that operational risk can be engineered down stands or falls on enforceability. A reference implementation in the United States, where the deepest pool of institutional credit capital sits, runs as follows. A bankruptcy-remote special-purpose vehicle finances the shipment or purchases the receivable; the electronic bill of lading governing the goods is held to the order of that vehicle as an electronic document of title under UCC § 7-106, control of which perfects a security interest in both the document and the goods it covers (§§ 9-312(c), 9-314), with a bailee's authenticated acknowledgment under § 9-313 available where goods are held outside a negotiable document (nn. 15 and 16 above), so that the financier's interest is perfected by control or possession rather than by filing alone and the collateral travels with the shipment; release of the goods at destination is conditioned on payment or acceptance; the originator retains a first-loss position; and decision-level audit trails are available to funders. Because § 7-106 dates from the 2003 revision of Article 7 and is enacted in virtually every state, the structure does not depend on the still-incomplete rollout of the 2022 amendments. Functionally equivalent structures are available today under the UK Electronic Trade Documents Act and Singapore's implementation of the Model Law, where the electronic document itself is now capable of possession. The point of specifying the structure is not novelty, since each element is settled law or settled securitization practice, but executability: the combination has only recently become available at the level of an individual small shipment, which is the unit-cost claim of Section 3 restated in legal form.

Two historical episodes suggest the magnitude of what standardization of this kind can do, and each carries a warning as well as a promise. The first is the standardization of the American mortgage

¹⁹ Berg et al. (2020), n. 12 above. Berg et al. concern consumer digital footprints only; for small-business default predicted on platform transaction data, see Frost, J., L. Gambacorta, Y. Huang, H. S. Shin and P. Zbinden (2019), "BigTech and the Changing Structure of Financial Intermediation," *Economic Policy* 34(100): 761-799, finding that a big-tech lender's internal rating built on platform data predicts loan losses more accurately than the local credit bureau's score in Argentina; on bank-account cash-flow data for consumers and small businesses, see FinRegLab (2019), *The Use of Cash-Flow Data in Underwriting Credit: Empirical Research Findings* (Washington, DC: FinRegLab), an independent evaluation of cash-flow variables from six non-bank lenders, found predictive of credit risk alone and in combination with traditional scores.

from the 1970s, which converted a local, heterogeneous bank asset into a globally funded one and permanently changed the cost and availability of housing finance. The analogy must be handled honestly: mortgage liquidity was built not on standardization alone but on agency guarantees that absorbed credit risk at taxpayer expense, and no public guarantor stands behind trade receivables. Private substitutes, retained first-loss positions, decision-level transparency to investors, and the discipline of short tenor, must do the work that the government-sponsored enterprises did for mortgages, which is one reason the screening problem discussed in Section 6 is first-order rather than incidental. The second episode is physical rather than financial: containerization, which by standardizing the unit of freight collapsed handling costs and is estimated to have done more for trade among industrialized countries than trade agreements did over the same period.²⁰ The liquid supply chain is, on this reading, the completion of the container revolution: the same standardization, applied to the financing of the box rather than to the box itself. The comparison is not ours alone: the ADB describes the digitalization of trade documents as potentially “transformational on a scale not seen since the size of container ships was standardized” (n. 3 above, p. 13).

5. Economic consequences

5.1 Firms: growth uncapped from cash

The first-order effect operates on credit-constrained firms. The best causal evidence on what happens when constrained firms receive credit comes from directed-lending settings: Banerjee and Duflo show that Indian firms gaining access to expanded directed credit, in their setting working-capital credit lines extended to formally registered manufacturers well above microenterprise scale, used it to expand production rather than to substitute for other borrowing. Sales and costs rose nearly one-for-one with the credit, and the implied return on a marginal rupee of lending was roughly 100 percent before interest, several multiples of even the 30 to 60 percent informal market rates they document, which is the signature of binding constraints.²¹ Trade-exposed SMEs are constrained in a specific, mechanical way: a cash conversion cycle in the region of 100 to 135 days means each order locks up capital for a quarter or more,²² so a fixed bank line caps the number of orders a firm can run concurrently regardless of demand. Financing that scales with each shipment converts the constraint from a stock (the credit line) into a flow (the trade itself). For a firm earning double-digit gross margins per order cycle against financing costs of 1 to 2 percent per month, the arithmetic of relaxation is large.

5.2 Trade: the credit-supply channel runs in both directions

A substantial empirical literature establishes that trade finance is not a veil: shocks to its supply move real trade flows, and the choice of financing instrument itself responds to risk and enforcement environments, exactly the margins that digital documents and observed flows shift.²³ Amiti and Weinstein show that deteriorating bank health transmitted directly to the exports of the firms those

²⁰ Levinson, M. (2006), *The Box: How the Shipping Container Made the World Smaller and the World Economy Bigger* (Princeton University Press); Bernhofen, D. M., Z. El-Sahli and R. Kneller (2016), “Estimating the Effects of the Container Revolution on World Trade,” *Journal of International Economics* 98: 36-50, who find containerization’s effect on trade among industrialized countries exceeded that of trade agreements over the period studied.

²¹ Banerjee, A. V., and E. Duflo (2014), “Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program,” *Review of Economic Studies* 81(2): 572-607. Instrumental-variable elasticities of sales and costs with respect to bank credit of 0.75 and 0.70; a marginal rupee of credit raises profit by Rs 0.89 net of interest, a return of roughly 105 percent before interest against a 16 percent bank rate and informal rates of 30 to 60 percent. The firms (plant and machinery of Rs 6.5 to 30 million) sit squarely within the SME band at issue here, and the credit took the form of working-capital limits, the asset class of this paper.

²² See n. 2 above.

²³ On the theory and empirics of payment terms and instrument choice in international trade: Schmidt-Eisenlohr, T. (2013), “Towards a Theory of Trade Finance,” *Journal of International Economics* 91(1): 96-112; Antràs, P., and C. F. Foley (2015), “Poultry in Motion: A Study of International Trade Finance Practices,” *Journal of Political Economy* 123(4): 853-901; Niepmann, F., and T. Schmidt-Eisenlohr (2017), “No Guarantees, No Trade: How Banks Affect Export Patterns,” *Journal of International Economics* 108: 338-350.

banks financed, and hit exports far harder than the same firms' domestic sales;²⁴ Paravisini and co-authors, using matched credit-export data from Peru, estimate the elasticity of exports to bank credit supply;²⁵ and Chor and Manova show that tighter credit conditions significantly deepened the trade collapse of 2008 and 2009.²⁶ Manova's work further demonstrates that credit constraints distort not just the volume but the composition of trade, suppressing entry by smaller firms into finance-intensive sectors.²⁷

If contractions in trade-finance supply contract trade, an expansion that closes part of a US\$2.5 trillion gap should expand it. A deliberately rough calculation indicates the order of magnitude at stake, and it rests on four assumptions we state explicitly. First, the financeable share: suppose only one-fifth of the measured gap proved commercially financeable once unit costs fell, roughly US\$500 billion of additional annual working-capital supply. Second, the elasticity: the matched-data estimates of the intensive-margin elasticity of exports to credit are in the region of 0.2.²⁸ Third, the margins: that elasticity captures only the intensive margin; effects at the extensive margin, entry by previously excluded firms, operate in addition and are precisely where a reduction in minimum ticket size bites, although the evidence on the relative size of the two margins is mixed. Fourth, symmetry: the elasticity is identified from a credit contraction, and applying it to an expansion assumes the response is symmetric, an extrapolation rather than a finding, with the easing counterfactual in Chor and Manova the closest direct evidence in the expansion direction. On those assumptions, the implied trade effects are measured in the hundreds of billions of dollars annually, concentrated among the small firms and emerging-market exporters where rejection currently falls; relax any one of them and the number moves accordingly. We present this as an order of magnitude, not a forecast; turning it into an estimate is part of the research agenda. The ADB expects demand for trade finance to rise as trade diversification and supply-chain reconfiguration accelerate, and warns that if supply does not increase commensurately, the resulting gaps will impede the opportunities arising from that reorganization, making the supply response more, not less, consequential.²⁹

5.3 Capital markets: a new short-duration asset class

For investors, liquid supply chains create something scarce: a large-volume, short-duration, self-liquidating asset class whose performance is anchored in observable flows of goods and short maturities, even though it remains exposed to trade and macroeconomic cycles, with historical credit

²⁴ Amiti, M., and D. E. Weinstein (2011), "Exports and Financial Shocks," *Quarterly Journal of Economics* 126(4): 1841-1877. Using matched bank-firm data on Japanese listed manufacturers (1987-1999, with a 2008-2010 check), they estimate that a 30 percent fall in a main bank's market-to-book value reduces client exports by about 2.7 percent relative to clients of unaffected banks, with no significant effect on domestic sales, and that financial shocks account for roughly 20 to 50 percent of the shortfall of export growth from trend in crisis years. Two limits on extrapolation: the effects are driven almost entirely by declines in bank health, with increases statistically indistinguishable from zero, so inferring a supply expansion from contraction evidence is an extrapolation rather than a finding of theirs; and the sample comprises large listed firms rather than SMEs.

²⁵ Paravisini, D., V. Rappoport, P. Schnabl and D. Wolfenzon (2015), "Dissecting the Effect of Credit Supply on Trade: Evidence from Matched Credit-Export Data," *Review of Economic Studies* 82(1): 333-359.

²⁶ Chor, D., and K. Manova (2012), "Off the Cliff and Back? Credit Conditions and International Trade during the Global Financial Crisis," *Journal of International Economics* 87(1): 117-133. Identification is from exporting-country interbank rates interacted with sector-level financial vulnerability in monthly US import data. Conservative within-country estimates imply US imports would have fallen about 2.5 percent more had interbank rates stayed at their September 2008 peaks, and about 5.5 percent less had rates dropped immediately to their August 2009 lows; the easing counterfactual provides some direct evidence on the expansion direction. They also find the collapse in US final demand more decisive than tighter credit supply, so exporter-side credit conditions act as an amplifier alongside demand.

²⁷ Manova, K. (2013), "Credit Constraints, Heterogeneous Firms, and International Trade," *Review of Economic Studies* 80(2): 711-744.

²⁸ Paravisini et al. (2015), n. 25 above: intensive-margin elasticity of export volumes to total bank credit of 0.195 (0.217 in value), operating through fewer and smaller shipments; no significant effect on entry or exit at one year (a negligible entry effect at two years), consistent with credit financing working capital. Identified from a crisis contraction among continuing exporters; applying it to an expansion assumes symmetry. Extensive-margin effects (Manova (2013), n. 27 above) operate in addition. The calculation in the text is illustrative only.

²⁹ ADB Briefs No. 378, n. 3 above, pp. 1-2 and 6: 87.5 percent of surveyed banks expect demand for trade finance to rise as trade diversification accelerates, and the brief warns that if higher demand materializes without a material increase in supply, large gaps would impede the opportunities arising from trade diversification and supply-chain reconfiguration.

losses below one percent.³⁰ The macro-financial consequences of opening such a channel are familiar from the mortgage precedent: capital deepening, spread compression as investor competition replaces bank-balance-sheet pricing, and a partial decoupling of trade credit supply from the health of any individual banking system. The last point is worth dwelling on. When banks contracted in 2008 and 2009, trade finance contracted with them.³¹ A distribution channel into diversified institutional capital would make working-capital supply more resilient to banking-sector stress, although, as Section 6 discusses, it imports risks of its own.

5.4 Distribution: who gains

Because the financing gap is concentrated among SMEs and in developing regions, the incidence of closing it is progressive across firms. SMEs employ the majority of the world's workers, and World Bank evidence indicates that financing-induced job creation is strongest among precisely the small, high-growth firms most likely to be rejected today.³² The mechanism also has a leveling property. It extends to small firms the embedded, transaction-level financing that large firms already enjoy through captive finance arms and supply-chain finance programs. Liquidity, in this sense, is not a new privilege but the generalization of an existing one.

6. Risks, limits and failure modes

We have stated the strongest version of the case so far. Five failure modes deserve equal prominence, and one recent collapse binds several of them together.

Model monoculture and correlated error. If a small number of AI underwriting systems originate a large share of exposures, their shared blind spots become systemic. An error class that a human underwriter commits idiosyncratically, a model commits at portfolio scale. Mitigations such as challenger models, human exception review and deliberate diversity of underwriting systems across the market are available but not automatic, and there is as yet no empirical record on which to assess them.

An adversarial fraud equilibrium. Automated verification raises the cost of today's fraud typologies; it also creates incentives to develop tomorrow's, including synthetic shipment data and compromised telemetry. The claim that logistics-anchored underwriting reduces fraud is an equilibrium claim, and the equilibrium will be contested. Physical control of cargo through the bill of lading is the strongest available backstop precisely because it does not depend on data integrity alone.³³

Securitization's screening problem. The 2008 crisis demonstrated that originate-to-distribute models can erode screening incentives: Keys and co-authors show that mortgage loans just easy enough to securitize defaulted significantly more often than observably similar loans that stayed on balance sheet.³⁴ Any architecture that distributes trade exposures to capital markets must answer this directly, through retained first-loss positions, full decision-level audit trails available to investors, and origination economics linked to performance. The fact that trade assets are short-tenor helps

³⁰ ICC Trade Register, n. 7 above, with the product-mix caveat of Section 2: realized losses on the open-account products that would dominate a liquid supply chain are low but not as low as on letters of credit, and a new origination channel must earn its own loss record.

³¹ Chor and Manova, n. 26 above, document the effect of tightened bank credit conditions on trade flows; Amiti and Weinstein, n. 24 above, document the transmission from bank health to client exports. For estimates of the contraction in trade-finance quantities itself, a worldwide shortfall estimated at US\$25-500 billion in the second half of 2008, see Auboin, M. (2009), "Boosting the Availability of Trade Finance in the Current Crisis: Background Analysis for a Substantial G20 Package," CEPR Policy Insight No. 35, and the IMF-BAFT survey evidence summarized in Chor and Manova, Section 3.

³² World Bank, n. 6 above; see also World Bank, "Quest to better understand the relationship between SME finance and job creation," World Bank Blogs (2022): in relative terms, employment gains from financing among firms of 10 to 49 employees are roughly three times those of firms with more than 100 employees.

³³ See n. 16 above.

³⁴ Keys, B. J., T. Mukherjee, A. Seru and V. Vig (2010), "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans," *Quarterly Journal of Economics* 125(1): 307-362.

(originators face the consequences of bad screening within months, not decades), but the lesson stands: liquidity without aligned screening incentives is how safe asset classes stop being safe.

The cautionary case is recent: Greensill. Supply-chain finance has already produced its own distribution failure, and any paper proposing to send trade assets to capital markets must confront it. Greensill Capital, which collapsed in March 2021, financed claims marketed as supply-chain assets that included “prospective receivables” anchored to no completed trade and in some cases to no existing customer relationship; concentrated its book heavily in a handful of related obligors; distributed the exposures through fund structures whose end investors had no ability to verify the underlying claims; and depended on trade credit insurance whose withdrawal proved fatal.³⁵ Greensill is sometimes read as an indictment of supply-chain finance as an asset class. We read it as this paper’s mechanism stated in the negative: claims that cannot be independently verified at the level of the individual trade should not be distributed, and when they are, the operational tail that Section 4 describes is transferred, unpriced, to investors who cannot see it. Each design requirement proposed here, claims anchored to observed shipments rather than projections, machine verification of documents at source, decision-level audit trails accessible to funders, control of the underlying collateral, and retained first loss, is the direct negation of a specific Greensill failure mode. The episode is therefore both the strongest argument for skepticism about this market and the clearest specification of what an architecture must demonstrate before it deserves institutional funding.

Procyclicality. Credit that scales automatically with trade flows will also contract automatically when trade contracts, and capital-markets funding can withdraw faster than relationship bank lending.

Governance and accountability. Delegating credit decisions to AI systems raises unsettled questions of explainability, fairness across borrower populations and clear human responsibility for final decisions. These questions are regulatory as much as technical, and they remain open in every major jurisdiction.

7. Conclusion: a measurable research agenda, and an invitation

The case for liquid supply chains rests on a single economic claim: that the trade finance gap is substantially a unit-cost phenomenon, and that the unit cost is now falling for technological, legal and informational reasons that are observable and dated. If the claim is right, the consequences run from the cash position of a small exporter to the composition of global capital markets. If it is wrong, the gap will persist even as underwriting costs fall, and the explanation must lie elsewhere, in risk, in regulation or in demand. Section 2 set out why the existing public data cannot decide the question; only measurement on live origination can.

The research design. The claim is testable, and the ADB itself has called for exactly this measurement.³⁶ A credible pilot evaluation would track, against matched controls: underwriting and compliance cost per shipment; compliance and fraud-detection error rates relative to incumbent human processes, the condition on which the entire technological argument of Section 3.1 rests; time from application to decision; approval rates for previously rejected SME segments; realized credit losses in newly approved segments against both the ICC Trade Register benchmarks and the originator’s incumbent book, which is the discriminating statistic between the unit-cost and adverse-selection explanations of the gap; the willingness of unaffiliated institutional investors to fund the resulting assets, and at what spread; and downstream firm outcomes such as order volume, employment and survival. Each is observable within the life of a single financed exposure, whose

³⁵ House of Commons Treasury Committee, *Lessons from Greensill Capital*, Sixth Report of Session 2021-22 (HC 151, July 2021), documenting the financing of “prospective receivables” not arising from existing trading relationships, the concentration of exposure to the GFG Alliance, and the role of withdrawn trade credit insurance in the March 2021 collapse. On the wind-down of the Credit Suisse supply-chain finance funds through which the exposures reached end investors: FINMA, “FINMA concludes proceedings against Credit Suisse in the Greensill case,” 28 February 2023; Credit Suisse Asset Management, suspension and liquidation announcements, March 2021.

³⁶ ADB Briefs No. 378, n. 3 above, p. 5 (and n. 19 therein): the narrowing of SME rejection rates may reflect support from nonbank institutions including fintech platforms, “warrants continued monitoring and research,” and would benefit from analysis with industry partners, the measurement agenda proposed here. On AI specifically, pp. 10-11.

tenor runs from roughly 45 days for in-transit finance to 60 to 90 days for post-shipment receivables, which makes trade finance an unusually fast laboratory for a question that extends well beyond it: whether AI expands access to scarce institutional capability, or merely makes existing institutions cheaper to run. That question is now being asked, and funded, by the institutions building the underlying systems, and trade finance offers their measurement programs something most settings cannot: the output of the automated work is a priced, short-maturity, independently auditable financial asset, so the value of the automation can be read off realized losses and market spreads rather than inferred from surveys or task-level benchmarks.

The offer to researchers: the design window. The author's firm is at the beginning of origination and has, as yet, no performance record; we state this plainly, because it defines what can honestly be offered. What a new channel offers that an established one cannot is the design window: almost every evaluation in empirical finance is constrained by decisions taken before the researchers arrived, variables never logged, declined applicants never tracked, rollouts never randomized, and every one of those decisions becomes irreversible here at launch. The engagement sought is specific: a small number of structured research partnerships, with priority for proposals arriving with independent third-party funding, whether from research councils, multilateral institutions or the economic-research programs of major AI laboratories, since external funding both preserves independence and ensures the work is resourced to completion.

Governance commitments. To such partnerships the firm commits the following, under agreements that protect counterparty confidentiality while guaranteeing independence of analysis and publication, including of unfavorable results, with no right of approval over findings. The pipeline will be instrumented for research from the first transaction: decision-level logs, shipment-level telemetry, and retention of declined as well as approved applications, without which the selection question at the heart of Section 2 cannot be answered. Where operationally feasible, rollout will be structured to support causal identification; any randomization would apply in one direction only, extending credit to applicants whose counterfactual is rejection and never withholding it from those who would otherwise be approved. Pre-registration of the evaluation design before the data exist, including the sample and seasoning thresholds at which performance results are first reported, is welcomed.

Because individual exposures mature in roughly 45 to 90 days, a usable panel forms at pilot volumes and within quarters, on a timetable compatible with grant cycles; the bet a researcher makes is not on the firm reaching scale, but on it reaching its first cohorts of transactions. An evaluation designed before the first transaction, by a funded examiner with no stake in the result, is worth more than any retrospective dataset: to the literature, to the institutional investors who will in time be asked to fund the resulting assets, and to the discipline of the firm itself. Researchers and research funders interested in the design window while it remains open are invited to make contact.