A Simple Prudential-Effort Foundation for the Financial Trilemma*

Charles Nolan†
University of Glasgow

February 2, 2026

Abstract

The "financial trilemma" asserts that deep financial integration, purely national financial policies and financial stability cannot simultaneously be achieved. Existing formalizations employing ex post burden-sharing games imply the trilemma result hinges on equilibrium selection. We develop a minimal ex ante prudential-effort model where financial integration amplifies cross-border crisis risk and national regulators internalise only part of global losses. The unique symmetric Nash equilibrium underprovides prudential effort and cannot deliver first-best stability when both integration and national policy autonomy are high. That provides a unique-equilibrium foundation for the financial trilemma and clarifies when supranational prudential arrangements are needed.

JEL Classification: F33; G28; H41.
Keywords: Financial trilemma; Financial stability; Prudential coordination

†Tel +44(0)141 330 8693. E-mail: Charles.Nolan@glasgow.ac.uk

1 Introduction

The classic "monetary trilemma" states that, with high capital mobility, a country cannot simultaneously maintain a fixed exchange rate and an independent monetary policy (Obstfeld et al., 2005). A parallel "financial trilemma" has entered the policy debate (Schoenmaker, 2011; Obstfeld, 2015): in a world of internationally active banks and cross-border capital flows, it is impossible to achieve simultaneously (i) deep financial integration, (ii) purely national control over financial regulation and resolution, and (iii) a first-best level of financial stability (Obstfeld and Taylor, 2017).

Schoenmaker (2011) formalises the financial trilemma in a simple contribution game where national authorities decide ex post whether to participate in the rescue of internationally active banks. An inefficient equilibrium, where each country withholds support and bank failures are widespread, coexists with an efficient equilibrium in which countries coordinate on a joint rescue. Mathis (2023) shows that the inefficient outcome is never unique: the trilemma conclusion in that specific setup therefore rests on equilibrium selection rather than a fundamental incompatibility between objectives. Our aim is to provide a simple, plausible model that features a unique equilibrium and in which the financial trilemma follows directly from primitives.

We develop a minimal two-country model where national regulators choose prudential effort or intensity *ex ante*, rather than bailout contributions *ex post*. Intensity corresponds to a broad index of capital requirements, liquidity buffers and other constraints on banks' risk-taking. Financial integration increases the potential for cross-border spillovers; when banks are more integrated, weak regulation in one country raises crisis risk elsewhere. A global social planner internalises the full loss from a systemic crisis and chooses a level of prudential effort that trades off higher buffers against lower expected losses. In contrast, national regulators internalise only a fraction of the global loss; the rest is borne by foreign taxpayers, depositors or investors.

Within this simple set-up, the unique symmetric Nash equilibrium features systematically underprovided prudential effort relative to the global optimum whenever financial integration is positive and national regulators internalise only part of the global loss. The gap between the equilibrium and first-best prudential levels, and hence the degree of inefficiency in crisis risk, increases with financial integration. For low integration, national autonomy is more benign; for high integration, the combination of purely national prudential policy and deep integration is incompatible with first-best stability. Achieving first-best stability at high integration requires either centralising prudential authority (as in a banking union) or reducing integration.

The model is intentionally stripped down: it abstracts from bargaining, multiple bank types, asymmetric countries and dynamics. Its aim is to derive a genuine financial trilemma in a simple ex ante prudential-effort framework, based solely on cross-border externalities in crisis risk.

2 A prudential-effort model of the financial trilemma

Consider two countries, Home (H) and Foreign (F), whose banking systems are integrated to a degree measured by $\lambda \in [0,1]$. A higher λ means that a larger share of banks' activities are cross-border, so that distress in one country has stronger repercussions in the other. We take λ as given in this note, reflecting both policy choices (legal and regulatory barriers to cross-border activity) and market choices (banks' and investors' appetite for cross-border positions). More generally, λ combines a policy-driven component λ^P (access regimes, capital controls, passporting and similar arrangements) and a market-driven component λ^M (the endogenous intensity of cross-border activity chosen by intermediaries given those regimes). For present purposes, we treat λ as a reduced-form summary of both components–but the distinction is important in reality.

Each national regulator $i \in \{H, F\}$ chooses a prudential effort level $k_i \in [0,1]$. The variable k_i can be interpreted as a composite index of capital requirements, liquidity standards, supervisory intensity and other instruments that make the banking system more resilient. Higher values of k_i reduce crisis risk but are costly in terms of foregone lending and compliance costs. We denote the cost of prudential effort by $C(k_i)$ and assume throughout that C is twice continuously differentiable, strictly increasing and convex, with $C'(k) > 0$ and $C''(k) > 0$ on $[0,1]$.

What matters for aggregate financial stability is the average prudential stance,

$$\kappa \equiv \frac{k_H + k_F}{2}, \tag{1}$$

and the degree of financial integration λ. We assume that the probability of a systemic crisis is given by a function $\pi(\kappa, \lambda)$ which is continuously differentiable in κ, strictly decreasing in κ and increasing in λ. For simplicity we occasionally refer to the linear benchmark

$$\pi(\kappa, \lambda) = \lambda(1 - \kappa), \tag{2}$$

but the main proposition is stated under more general conditions.

If a systemic crisis occurs, it creates a global welfare loss of magnitude $L > 0$, which can be interpreted as an output loss, the fiscal cost of bank recapitalisation, or a broader measure of social costs. A *global social planner* internalises the full loss L and chooses (k_H, k_F) to minimise the expected global cost,

$$J^{\mathrm{SP}}(k_H, k_F) = L\pi(\kappa, \lambda) + C(k_H) + C(k_F). \tag{3}$$

Using $\kappa = (k_H + k_F)/2$ and differentiating with respect to k_H yields

$$\frac{\partial J^{\mathrm{SP}}}{\partial k_H} = L\pi_\kappa(\kappa, \lambda)\frac{\partial \kappa}{\partial k_H} + C'(k_H) = \frac{L}{2}\pi_\kappa(\kappa, \lambda) + C'(k_H), \tag{4}$$

where π_κ denotes the partial derivative of π with respect to κ. In a symmetric optimum, $k_H^* = k_F^* = k^*(\lambda)$ and $\kappa^* = k^*(\lambda)$, and the first-order condition becomes

$$C'\big(k^*(\lambda)\big) = -\frac{L}{2}\,\pi_\kappa\big(k^*(\lambda), \lambda\big). \tag{5}$$

Because C' is increasing and $\pi_\kappa < 0$, equation (5) pins down a unique $k^*(\lambda)$ under standard convexity assumptions.

Now replace the social planner with two national regulators. Each regulator cares about domestic costs and internalises only a fraction $\theta \in (0, 1]$ of the global loss L in the event of a crisis. The remaining $(1 - \theta)L$ is borne by foreign taxpayers, depositors and investors.

The objective function of regulator i is

$$J_i(k_i; k_j) = \theta L \pi(\kappa, \lambda) + C(k_i), \tag{6}$$

where k_j is the prudential effort chosen by the other regulator and $\kappa = (k_H + k_F)/2$ as before. For given k_j, regulator i chooses k_i to minimise J_i. Differentiating with respect to k_i yields

$$\frac{\partial J_i}{\partial k_i} = \theta L \pi_\kappa(\kappa, \lambda) \frac{\partial \kappa}{\partial k_i} + C'(k_i) = \frac{\theta L}{2} \pi_\kappa(\kappa, \lambda) + C'(k_i). \tag{7}$$

Before continuing, it is worth remarking that in practice, the degree of financial integration and the degree of cross-border internalisation are unlikely to be independent. Greater integration (a higher λ) typically enlarges the cross-border footprint of financial institutions and the expected incidence of spillovers, while leaving national regulators politically and fiscally accountable primarily to domestic constituencies. In that case, effective internalisation may fall as integration rises, i.e. $\theta = \theta(\lambda)$ with $\theta'(\lambda) < 0$. This does not change any of the qualitative results below; if anything it reinforces them. The planner's optimality condition still implies $k^*(\lambda)$ is increasing in λ when integration raises the marginal social benefit of effort. By contrast, the national regulator's condition $C'(k_i) = -(\theta(\lambda)L/2)\pi_\kappa(\kappa, \lambda)$ scales that benefit down by $\theta(\lambda)$, so a declining $\theta(\lambda)$ widens the wedge between the planner and decentralised choices as integration deepens.[1]

Game-theoretic formulation and main result

For a given degree of financial integration $\lambda \in (0, 1]$ and internalisation parameter $\theta \in (0, 1]$, define the game

$$\Gamma(\lambda, \theta) = \big(\{H, F\}, \{K_i\}_{i=H,F}, \{u_i\}_{i=H,F}\big),$$

where the players are the two national regulators, each strategy set is $K_i = [0, 1]$ (prudential effort k_i), and regulator i's payoff is $u_i(k_i, k_j; \lambda, \theta) = -J_i(k_i; k_j)$, with J_i as above.

We impose the following regularity conditions.

Assumption 1. (a) The crisis probability $\pi(\kappa, \lambda)$ is continuously differentiable in κ and strictly decreasing in κ for all $\lambda > 0$, with $\pi_\kappa(\kappa, \lambda) < 0$. (b) The effort cost $C(k)$ is twice continuously differentiable, strictly increasing and strictly convex, with $C'(k) > 0$ and $C''(k) > 0$ on $[0, 1]$. (c) For all (λ, θ) of interest, the planner's objective $J^{\mathrm{SP}}(k_H, k_F)$ is strictly convex in (k_H, k_F) and each regulator's objective $J_i(k_i; k_j)$ is strictly convex in k_i.

Assumption 1 ensures that both the planner and each national regulator face a well-behaved trade-off, so that optimal choices and best responses are unique.

Proposition 1. *Fix $\lambda \in (0, 1]$ and $\theta \in (0, 1]$, and suppose Assumption 1 holds. Then:*

[1] I thank the referee for emphasising this point.

1. *(Planner) There exists a unique symmetric solution $k^*(\lambda)$ to the global planner's problem, with $k_H^* = k_F^* = k^*(\lambda)$. It is characterised by the first-order condition*

$$C'\big(k^*(\lambda)\big) = -\frac{L}{2}\,\pi_\kappa\big(k^*(\lambda), \lambda\big).$$

2. *(National regulators) The game $\Gamma(\lambda, \theta)$ has a unique symmetric Nash equilibrium $k^N(\lambda, \theta)$, with $k_H^N = k_F^N = k^N(\lambda, \theta)$. It is characterised by*

$$C'\big(k^N(\lambda, \theta)\big) = -\frac{\theta L}{2}\,\pi_\kappa\big(k^N(\lambda, \theta), \lambda\big).$$

3. *(Underprovision of prudential effort) If $\theta \in (0, 1)$, then for any $\lambda > 0$ the Nash-equilibrium prudential effort is strictly below the planner's choice,*

$$k^N(\lambda, \theta) < k^*(\lambda),$$

and the associated crisis probability is strictly higher,

$$\pi^N(\lambda, \theta) \equiv \pi\big(k^N(\lambda, \theta), \lambda\big) \ > \ \pi^*(\lambda) \equiv \pi\big(k^*(\lambda), \lambda\big).$$

Thus, the three 'goals' – high financial integration (λ large), purely national control ($\theta < 1$), and first-best stability ($k = k^*(\lambda)$) – cannot be jointly attained; with $\theta < 1$ and $\lambda > 0$ the unique equilibrium necessarily yields $k^N(\lambda, \theta) < k^*(\lambda)$.

Proof. Parts (1) and (2) follow from standard arguments. Under Assumption 1(c), the planner's objective J^{SP} is strictly convex in (k_H, k_F), so there is a unique minimiser. Symmetry of the environment implies that the unique solution is symmetric, $k_H^* = k_F^* = k^*(\lambda)$, and the stated first-order condition follows from differentiating J^{SP} with respect to k_H and imposing symmetry. Analogously, for each fixed k_j the regulator's objective $J_i(k_i; k_j)$ is strictly convex in k_i, so there is a unique best response. The game $\Gamma(\lambda, \theta)$ is symmetric, and the pair of best-response functions has a unique fixed point in $[0, 1]^2$; under symmetry this fixed point must be of the form (k^N, k^N) and satisfies the stated characterisation.

For part (3), consider the first-order condition of the planner at $k^*(\lambda)$,

$$C'\big(k^*(\lambda)\big) + \frac{L}{2}\,\pi_\kappa\big(k^*(\lambda), \lambda\big) = 0,$$

and evaluate instead the Nash-regulator first-order condition at the same point:

$$C'\big(k^*(\lambda)\big) + \frac{\theta L}{2}\,\pi_\kappa\big(k^*(\lambda), \lambda\big) = \frac{(\theta - 1)L}{2}\,\pi_\kappa\big(k^*(\lambda), \lambda\big).$$

Because $\lambda > 0$, $\theta - 1 < 0$ and $\pi_\kappa < 0$, the right-hand side is strictly positive, so the left-hand side is positive as well. By strict convexity of J_i in k_i and continuity of the first-order condition, this implies that at $k^*(\lambda)$ the regulator's marginal cost of effort is below the perceived marginal benefit, so the unique minimiser of J_i—and hence $k^N(\lambda, \theta)$—must lie strictly below $k^*(\lambda)$. Intuitively, because each regulator internalises only a fraction $\theta < 1$

of the global loss L, the perceived marginal benefit of additional prudential effort is scaled down relative to the planner's, leading to lower effort.

Since $\pi(\cdot, \lambda)$ is strictly decreasing in its first argument by Assumption 1(a), $k^N(\lambda, \theta) < k^*(\lambda)$ implies $\pi^N(\lambda, \theta) = \pi(k^N(\lambda, \theta), \lambda) > \pi(k^*(\lambda), \lambda) = \pi^*(\lambda)$. This establishes strict underprovision of prudential effort and strictly higher crisis risk in the unique symmetric Nash equilibrium whenever $\theta \in (0, 1)$ and $\lambda > 0$. □

In a linear-quadratic benchmark (2) with $C(k) = \frac{c}{2}k^2$, equations (5) and (7) yield closed-form solutions

$$k^*(\lambda) = \frac{L\lambda}{2c}, \qquad k^N(\lambda, \theta) = \frac{\theta L\lambda}{2c},$$

so that $k^N(\lambda, \theta) = \theta k^*(\lambda)$ and the wedge between the planner's and Nash prudential levels scales linearly with $(1 - \theta)\lambda$.

3 Relation to existing work and monetary considerations

Two stylised applications help to illustrate how the model maps into real-world cases. First, consider a monetary union with high financial integration but initially national prudential authority, such as the pre-crisis euro area. Membership in the union raises λ by encouraging cross-border banking within a single currency area, while prudential and resolution decisions remain largely national, with $\theta < 1$ as national authorities internalise domestic costs but not the full cross-border fallout of bank distress. Proposition 1 then implies that, at high λ and $\theta < 1$, equilibrium prudential effort will fall short of the union-wide optimum, and the system will be more crisis-prone than a planner would choose. The subsequent creation of a banking union and common resolution framework can be interpreted as a move towards $\theta \approx 1$ at the union level, bringing $k^N(\lambda, \theta)$ closer to $k^*(\lambda)$ for a given degree of integration.

Second, consider a small open economy with a floating exchange rate, open capital markets and a relatively conservative national macroprudential regime. Here capital-account openness and participation in global markets imply a non-trivial λ, but the absence of a hard peg and the presence of currency risk tend to limit the most extreme cross-border exposures. At the same time, a strong domestic prudential framework—operationalised in the model as a relatively high effective θ and ambitious k—keeps the probability of systemic crisis, and hence risk premia, within bounds. Here, the financial trilemma is present but less stark as the country maintains some national autonomy and financial integration while containing crisis risk through robust prudential policy, and monetary-policy autonomy is preserved through the floating rate. This configuration corresponds to regimes often recommended for financially open small economies in practice.

4 Conclusion

The model focuses on regulators' ex ante choice of prudential effort. Cross-border externalities operate through the crisis probability. Under our assumptions, the best-response functions are single-valued and there is a unique symmetric Nash equilibrium. Underprovision of prudential effort, and thus elevated crisis risk, are not artefacts of equilibrium selection, but a consequence of national regulators internalising only part of the global loss. This provides a simple, unique-equilibrium foundation for the financial trilemma. Achieving first-best stability at high integration requires either supranational prudential arrangements or a retreat from deep financial integration. Because prudential choices and capital-flow measures feed into risk premia and wedges in standard uncovered interest parity conditions, the financial trilemma and the monetary trilemma are intertwined. Extending this simple framework to a richer setting with endogenous integration and explicit monetary policy is important for future work.

References

[1] Obstfeld, M (2015) Trilemmas and tradeoffs: living with financial globalisation BIS Working Papers No 480

[2] Obstfeld, M, J C Shambaugh and A M Taylor (2005) The trilemma in history: tradeoffs among exchange rates, monetary policies, and capital mobility, *Review of Economics and Statistics*, vol 87, no 3, August, pp 423–38.

[3] Obstfeld, M, and A M Taylor (2017) International Monetary Relations: Taking Finance Seriously. CEPR Discussion Paper No. DP12079

[4] Mathis, J (2023) Corrigendum to ["The financial trilemma" Economics letters 111 (2011) 57–59], *Economics Letters*, Volume 224.

[5] Schoenmaker, D (2011), The financial trilemma, *Economics Letters*, Volume 111, Issue 1.