mathematics

MDPI

Article

From the Intermediate Value Theorem to the Rank Theorem: An Elementary and Computable Path Through the Implicit and Inverse Function Theorems

TzeHoung Lee * and Xue-Ming Yuan

Centre for Maritime Studies, National University of Singapore (NUS), Singapore 119077, Singapore
* Correspondence: thlee@nus.edu.sg

Abstract

The Implicit Function Theorem (IFT) and the Inverse Function Theorem (InvFT) are foundational results of multivariable analysis, underpinning differential geometry, comparative statics in economics, and the local solvability of partial differential equations. We present self-contained elementary proofs of both theorems on $\mathbb{R}^n$, together with the companion Rank Theorem, following Dini's 1876 inductive strategy as simplified by de Oliveira: the IFT is established first by induction on the number of dependent variables m, using only the one-dimensional Intermediate Value Theorem (IVT) and Mean Value Theorem (MVT); the InvFT is then derived immediately from the IFT via an elementary injectivity estimate; and the Rank Theorem follows from the InvFT by an explicit block-rank argument. This approach requires no compactness arguments in $\mathbb{R}^n$ ($n \geq 2$), no contraction mappings, and no fixed-point theorems; its only analytic input is the completeness of $\mathbb{R}$, through the one-variable IVT and MVT. We then demonstrate that the resulting framework is not merely pedagogical but directly computable by deploying it end to end on four calibration and optimization primitives drawn from arbitrage-free interest-rate modeling and bank balance-sheet management: joint drift–local-volatility slice calibration; a minimum lending rate under an endogenous option-cost feedback, including an explicit fold bifurcation at which the IFT hypothesis fails; comparative statics of a mean–variance efficient frontier under a nonlinear risk limit; and reverse optimization via the InvFT and Rank Theorem, including an exact reachability/infeasibility criterion when a governing body specifies more targets than free decision variables. For each application, we state the map explicitly, verify the relevant regularity hypothesis at a numerically realized solution, derive the closed-form sensitivity the theorem provides, and verify that formula against central-difference approximation to at least five significant figures. The examples use synthetic data and serve as a proof of concept; all are reproducible from the expressions given in the text and from a supplementary script.

Keywords: Implicit Function Theorem; Inverse Function Theorem; Rank Theorem; Dini's theorem; elementary proof; Intermediate Value Theorem; comparative statics; local volatility calibration; arbitrage-free interest rate models; reverse optimization; efficient frontier; balance-sheet management; sensitivity analysis

MSC: 26B10; 26A24; 58C15; 91G30; 90C31; 91G70

Check for updates

Academic Editor: Patricia J. Y. Wong

Received: 14 September 2026
Revised: 18 October 2026
Accepted: 20 October 2026
Published: 5 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

1.1. The Elementary Path

The Implicit Function Theorem (IFT) is, with the Inverse Function Theorem (InvFT) and the Rank Theorem, one of the load-bearing results of multivariable analysis: it underlies the definition of an embedded submanifold in differential geometry [1–3], the local solvability of nonlinear elliptic and parabolic PDE [4], and the entire technique of comparative statics in economics [5–7]. Its standard modern proof, however, is not elementary: textbook treatments (e.g., [8]) typically invoke either the Banach contraction mapping principle or a compactness argument on a closed ball, tools that are themselves nontrivial and that obscure the theorem's genuinely simple one-dimensional content. Dini's original 1876 treatment [9] took a different, inductive route, and de Oliveira [10] showed that this route can be walked using only the one-dimensional Intermediate Value Theorem (IVT) and Mean Value Theorem (MVT), with no compactness argument in $\mathbb{R}^n$, no contraction mapping and no fixed-point theorem. (Completeness of $\mathbb{R}$ is still used, through the IVT and MVT, and cannot be avoided; see Section 3). We extend this elementary strategy to C^k regularity, to a Schur-complement form of the inductive step, to a quantitative injectivity estimate, and to the Rank Theorem, and we pair it with a systematic numerical check that each hypothesis invoked is actually met at a computed point. We then ask a question the purely analytic literature does not: *can this elementary, computable framework be carried through end to end on a quantitative problem from practice, rather than on textbook examples?*

1.2. The Applied Test Case

A series of industry white papers by Ho [11–17], which combine arbitrage-free interest-rate modeling (in the tradition of the peer-reviewed Ho–Lee model [18]) with balance-sheet performance attribution, propose describing a bank by what they call a canonical state vector u of economically meaningful coordinates—net interest income, economic value of equity, attributed spread income, duration, and the like—computed from balance-sheet actions w through a chain of pricing, attribution, and optimization operators. Two questions recur throughout that program and are left, in the applied literature, at the level of numerical procedure rather than theorem:

(Q1) **Calibration.** Given market data (a yield curve, a set of option premia), when does a model (a drift, a local-volatility surface) matching that data exist, when is it unique, and how does it move when the data is revised?

(Q2) **Reversal.** Given a target region of the canonical state space—a board-approved objective for margin, duration, or risk—when do balance-sheet actions exist that reach it, and how many independent targets can n free actions honor simultaneously?

Both questions are, formally, questions about a map $F : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^n$ (or $\Psi : \mathbb{R}^n \to \mathbb{R}^m$) near a point at which $F = 0$ (respectively at which Ψ takes a reference value), and locally both are answered, with an explicit derivative formula for the sensitivity the practitioner wants, by the three theorems of Section 3. The present paper therefore serves two purposes at once. It is, first, a self-contained analytic paper: every theorem it uses is proved from first principles in Section 3, not merely cited. It is, second, an applied paper: it asks whether the calibration and reversal problems of balance-sheet management are direct instances of these theorems or merely analogous to them, and shows, by explicit construction and numerical verification in four cases, that they are direct instances. Stating them this way gives explicit sensitivity (comparative-statics) formulas, each obtained from one linear solve at the computed point, that are not stated in the applied sources above, and it gives a precise mathematical counterpart—a rank condition— to informal industry requirements such as "reconciliation" and target feasibility.

The applications are a proof of concept. Their parameters are hand-chosen or synthetically generated (a stylized lattice slice, a synthetic skewed smile, three hypothetical assets) and are not calibrated to observed market quotes or to any institution's balance sheet. They are designed to show that the theorems can be deployed computationally and that the resulting closed-form sensitivities are correct; they make no empirical claim about actual markets or banks.

1.3. What the Elementary Route Contributes to the Applications

A calibration routine embedded in a risk system, or a pricing committee setting a minimum lending rate, cannot act on an abstract existence claim. It needs hypotheses it can check numerically at the point where it operates, a formula for how the solution moves when the inputs move, and some indication of how far the inputs may move before that formula stops being trustworthy. The elementary route is well suited to all three. Its hypotheses are finite-dimensional and checkable—the sign of a single partial derivative, a non-vanishing determinant, the rank of a minor— and in Sections 5–8, each one is verified at a computed solution. Its constants are produced at every step by one-variable estimates: the δ_1, ϵ_0 in the proof of Theorem 1 and the c, r of Lemma 3 are explicit.

We do not claim that these features are unavailable from the standard proofs. A contraction-mapping proof can also be made quantitative, and the estimate behind Lemma 3 is in fact one of its steps ([8], Theorem 9.24). Nor does the proof technique govern the numerical behavior reported below: the quadratic convergence of Newton's method follows from the standard local convergence theorem [19] (a C^2 map whose Jacobian is invertible at the root), and the fold point of Section 6 is located from the hypothesis $g_R \neq 0$ of the theorem, which every proof shares. What the elementary route offers is that the entire chain—from the IVT to the Rank Theorem—uses only one-variable tools and finite-dimensional linear algebra, so the same objects a practitioner computes (partial derivatives, determinants, minors) are the objects the proofs manipulate. Every sensitivity formula in Sections 5–8 is checked against central differences of re-solved systems (4 lists every check), and, where an exact solution exists (Section 7), against that solution as well.

Contributions and organization. The contributions are those listed as (N0)–(N4) in the Statement of Academic Novelty above; Section 2 surveys the literatures this paper connects and summarizes its position in Table 1. Section 3 contains the proofs (N0). Table 2 separates what is new in them from what is taken from de Oliveira [1] and the classical literature. Section 4 fixes notation, lists in Table 3 four recurring balance-sheet tasks and the theorem that answers each, and states the numerical conventions; every numerical check in the paper is collected in Table 4. Sections 5–8 contain the four worked applications (N1)–(N4). Section 9 records the common rank condition as a corollary of Section 3 and discusses scope and limitations, and Section 10 concludes.

2. Related Work

The material below draws on five largely separate studies that, to the authors' knowledge, have not previously been brought into direct contact. We survey each in turn and close by stating precisely the gap the present paper fills.

2.1. Classical Implicit and Inverse Function Theory

The Implicit Function Theorem originates with Dini's 1876 lecture notes on infinitesimal analysis [9], which give an inductive proof for C^1 systems reducing the n-equation case to n applications of the scalar case. In the early twentieth century, Goursat gave a proof by successive approximations, the forerunner of the contraction-mapping proof (see [20] for the history), and Bliss [21] gave a short existence proof whose mean-value injectivity

argument was later reused by de Oliveira [10]. The theorem entered the standard graduate analysis canon through Apostol [22], Rudin [8], and Spivak's coordinate-free treatment on manifolds [1]; Munkres [2] and Guillemin and Pollack [3] develop the companion Rank Theorem and its role in the definition of embedded submanifolds, which we invoke directly in Section 8. De Oliveira [10] revisited Dini's elementary strategy, proving the C^1 IFT and InvFT using only the Intermediate and Mean Value Theorems and avoiding compactness, contraction and fixed-point arguments; Section 3 below extends that strategy to C^k regularity, a Schur-complement induction, a quantitative injectivity estimate and the Rank Theorem (Table 2 gives the details). A parallel, more functional-analytic tradition treats the IFT in Banach space (the Hildebrandt–Graves theorem) and its set-valued/non-smooth generalizations [23], motivated originally by nonlinear elliptic PDE [4] and later by variational inequalities and optimization; Krantz and Parks [20] give the most complete single survey of the theorem's many proofs, generalizations, and historical variants. None of this literature addresses balance-sheet calibration or bank optimization; its relevance here is purely as the source of the theorems we apply.

2.2. Comparative Statics and the Envelope Theorem in Economics

Samuelson's *Foundations of Economic Analysis* [5] is the founding text of "comparative statics": the systematic derivation of $\partial x^\star/\partial\theta$ for an optimum or equilibrium $x^\star(\theta)$, explicitly via the IFT applied to first-order conditions. Silberberg and Suen [6] and, in a more general choice-theoretic setting, Milgrom and Segal's envelope theorems [7] extend the technique beyond smooth interior optima to constrained and non-smooth problems; our Section 7 nonlinear risk-limit application is, in this vocabulary, a comparative-statics exercise on a KKT system with a curved constraint manifold. This literature has not, to the authors' knowledge, been applied to bank balance-sheet policy limits, where the analogous constraint (a value-at-risk or volatility limit) is routinely nonlinear.

2.3. Arbitrage-Free Interest-Rate Modeling and Volatility Calibration

The arbitrage-free short-rate lattice underlying Section 5 descends from Ho and Lee's original binomial model [18], generalized to lognormal dynamics by Black, Derman, and Toy [24] and to a fully flexible forward-rate framework by Heath, Jarrow, and Morton [25]. The local-volatility idea used in Section 5 is Dupire's [26]: rather than assume a single volatility parameter, infer a volatility *surface* from the full grid of quoted option (equivalently swaption) prices, an approach developed contemporaneously for equity-index smiles by Derman and Kani [27] and Rubinstein's implied binomial trees [28]. The Black–Derman–Toy construction, in particular, requires solving jointly for a node's drift and volatility from two market prices at every step of the tree—precisely the two-parameter implicit system of Section 5—and is normally solved, in practice, by Newton iteration; Section 5 states the hypothesis under which that iteration solves a well-posed problem. Andreasen and Huge [29] address the closely related problem of building an arbitrage-free local-volatility surface directly from a full option chain via a sequence of such implicit solves. None of these sources states or uses the Implicit Function Theorem explicitly to certify well-posedness of the calibration; the calibration is presented as a numerical procedure whose success is verified empirically rather than guaranteed analytically.

2.4. Balance-Sheet Management, ALM, and Regulatory Context

The applied setting of this paper—a bank balance sheet managed jointly for interest-rate risk, credit risk, and profitability under one reconciled reporting framework—is the one described in Ho's IncomeRisk and Financial Representation program [11–17]. These are industry white papers and a monograph from a single firm; they are not peer-reviewed. We use them only as a source of motivating problems and terminology. Terms taken from them,

such as the "canonical state vector" and the "reconciliation" requirement, are informal industry concepts; where we use them (Sections 7–9), we give them a precise mathematical counterpart of our own, and no result of this paper is imported from those sources. The mathematics of Sections 5–8 is self-contained: each application is an explicit map that can be checked on its own terms. The regulatory backdrop is the Basel Committee's Interest Rate Risk in the Banking Book standard [30] and the current expected credit loss accounting standard [31], which require banks to measure risk on both an economic-value and an earnings basis. English, Van den Heuvel and Zakrajšek document empirically that interest-rate risk in the banking book is economically material and imperfectly hedged [32], and the supervisory review of the failure of Silicon Valley Bank in March 2023 [33] describes a case in which these risks were not managed in time.

2.5. Reverse and Inverse Optimization

The reverse-optimization problem of Section 8—given a target outcome, recover the decision variables—has a substantial history in portfolio theory under the name "reverse optimization" or "implied views." Sharpe [34] first proposed imputing expected returns from an observed efficient portfolio by inverting the mean–variance first-order conditions; Black and Litterman [35] popularized the technique as a way to blend market-implied and subjective views. A distinct operations-research literature on inverse optimization, surveyed by Ahuja and Orlin [36] and extended to conic and data-driven settings by Iyengar and Kang [37], studies the general problem of recovering a program's parameters from an observed or desired solution, typically via a projection onto a feasible parameter set. Our treatment differs from both bodies of work in using the Rank Theorem, rather than a projection or regression, to characterize *exact* reachability: Section 8.2 shows that with more targets than free decision variables, the natural object to report is not an approximate best fit but the lower-dimensional surface of attainable targets, together with an exact residual for any target off it (globally, in that example, because two of the targets are affine in the allocation).

2.6. Positioning of the Present Paper

Table 1 summarizes the gap. Each cell names a body of work that supplies either the analytic tool or the applied primitive; no existing source, so far as the authors are aware, connects a named row to a named column. To that extent, the paper offers what we believe is the first explicit cross-tabulation of the classical local-inversion theorems against the calibration and reverse-optimization primitives of balance-sheet management.

3. The Elementary Inductive Proofs: From the IVT to the Rank Theorem

We now give complete, self-contained proofs of the results used throughout the rest of the paper. The scalar Implicit Function Theorem is proved directly from the one-variable Intermediate Value Theorem and Mean Value Theorem (Section 3.1); the vector case follows by induction on the number of dependent variables, the inductive step being a Schur-complement reduction (Section 3.2); the Inverse Function Theorem follows from the vector IFT, together with a quantitative injectivity estimate (Section 3.3); and the Rank Theorem follows from the InvFT by a block-rank argument (Section 3.4).

What "elementary" means here.

The only analytic input to the arguments below is the completeness of $\mathbb{R}$, and it enters exclusively through the one-variable IVT and MVT, which we take as known. Completeness cannot be dispensed with, because the IFT is false over $\mathbb{Q}$: for $g(x,y) = y^2 - x$ at $(1,1)$, one has $g_y(1,1) = 2 \neq 0$, yet for the rationals $x_k = (k+1)/k \to 1$, the equation $y^2 = x_k$ has no rational solution (since $k(k+1)$ is never a perfect square for $k \geq 1$), so no implicit

function exists on any rational neighborhood of $x = 1$. What the arguments do avoid is the additional machinery of the standard proofs: there is no compactness argument in $\mathbb{R}^n$ with $n \geq 2$ (no extreme-value argument on balls or spheres and, in particular, no appeal to the equivalence of norms, because we use the Frobenius norm throughout), no contraction-mapping principle, no fixed-point theorem, and no limit of iterates. The only compactness anywhere is that of a closed bounded interval of $\mathbb{R}$, used implicitly in the textbook proof of the one-variable MVT via Rolle's theorem.

Table 1. The gap this paper addresses: analytic tool (rows) versus balance-sheet primitive (columns). Cell entries name the closest existing treatment; a dash indicates no located precedent applying that tool to that primitive.

	Rate Calibration	**Rate/Credit Pricing**	**Constrained Frontier**	**Reverse Optimization**
Scalar IFT	—	Merton put [38] (implicit)	—	—
General IFT	BDT/Dupire, numerical only [24,26,29]	—	Silberberg, envelope [6,7]	—
InvFT	—	—	—	Sharpe, Black–Litterman [34,35]
Rank Theorem	—	—	—	Ahuja–Orlin (approx. only) [36]

Notes on the non-empty cells, each of which is only approximately relevant. Merton [38]: in structural credit models, the spread, or an implied asset parameter, is defined only implicitly by an option-pricing equation, but the well-posedness and sensitivity of that equation are not analyzed through the IFT. BDT/Dupire/Andreasen–Huge [24,26,29]: calibrations are solved numerically by root-finding, without a stated IFT certificate of local uniqueness. Silberberg–Suen and Milgrom–Segal [6,7]: general comparative-statics and envelope results, not applied to volatility limits on a bank balance sheet; for the single-limit frontier of Section 7, the classical closed form [39] exists and serves as our benchmark. Sharpe and Black–Litterman [34,35]: the first-order conditions of a quadratic program are inverted; the map is linear, so the inversion is global linear algebra rather than the local nonlinear InvFT. Ahuja–Orlin [36]: inverse optimization by duality and projection, which yields best approximations rather than an exact reachability criterion.

Provenance.

Table 2 records, result by result, what is taken from de Oliveira [10], who proves the C^1 IFT and InvFT by this route, what is classical, and what is new here.

Table 2. Provenance of the results of Section 3.

Result	**Relation to the Literature**
Theorem 1 (scalar IFT)	The IVT/MVT argument of [10] (Theorem 1), which is stated for C^1. Added here: the C^k bootstrap, the two-sided uniqueness statement on $U \times V$, and explicit constants δ_1, ϵ_0.
Lemma 1	Standard linear algebra.
Theorem 2 (vector IFT)	[10] (Theorem 2) proves the C^1 case by first normalizing $\partial F/\partial y$ to the identity and then obtaining uniqueness from the injectivity of $(x, y) \mapsto (x, F(x, y))$ [10] (Lemma 4). Our proof keeps a general invertible Jacobian, uses the Schur complement (the determinant identity of the classical inductive proof attributed to Dini [9]; see [20]), obtains uniqueness directly from the choice of neighborhoods, and covers C^k.
Lemma 2	Standard (e.g., [8], Theorem 5.19).
Lemma 3	New in this form. [10] (Lemma 4), after Bliss [21], proves injectivity near p qualitatively; Lemma 3 gives a quantitative lower bound with explicit c and r. The same estimate is a step of the contraction-mapping proof (e.g., [8], Theorem 9.24); here it follows from Lemma 2 alone.
Theorem 3 (InvFT)	Derived from the vector IFT as in ([10], Theorem 3); here for C^k, with the explicit domain supplied by Lemma 3.
Lemmas 4–6, Theorem 4, Corollary 1	Not treated in [10]. The Rank Theorem and the shape of its proof are classical [1–3]; what is shown here is that the whole chain stays within the elementary toolkit described above.
Proposition 1	Standard (differentiate (3) once more).

Norms.

Throughout, $\|v\|$ is the Euclidean norm of $v \in \mathbb{R}^n$, and $\|A\|_F = (\sum_{i,j} a_{ij}^2)^{1/2}$ is the Frobenius (Hilbert–Schmidt) norm of a matrix $A = (a_{ij})$, as in [10], Lemma 1. Applying the Cauchy–Schwarz inequality to each row of A gives

$$\|Av\| \le \|A\|_F \, \|v\| \qquad \text{for every matrix } A \text{ and vector } v. \tag{1}$$

Because $\|A\|_F$ is the Euclidean norm of the entries of A, a matrix whose entries depend continuously on x is continuous in $\|\cdot\|_F$.

3.1. The Scalar Case

Theorem 1 (Scalar Implicit Function Theorem). *Let $\Omega \subset \mathbb{R}^n \times \mathbb{R}$ be open, $g \in C^k(\Omega;\mathbb{R})$, $k \ge 1$, and $(a,b) \in \Omega$ with $g(a,b) = 0$ and $\partial g/\partial y(a,b) \ne 0$. Then there exist an open ball $U \ni a$ in $\mathbb{R}^n$ and an open interval $V \ni b$ with $U \times V \subset \Omega$, and a function $\varphi \in C^k(U;V)$ with $\varphi(a) = b$, such that for all $(x,y) \in U \times V$*

$$g(x,y) = 0 \iff y = \varphi(x).$$

Moreover, $\partial g/\partial y \ne 0$ on $U \times V$, and

$$\frac{\partial \varphi}{\partial x_i}(x) = -\frac{\partial g/\partial x_i\,(x,\varphi(x))}{\partial g/\partial y\,(x,\varphi(x))}, \qquad i = 1,\dots,n, \quad x \in U. \tag{2}$$

Proof. Assume without loss of generality $g_y(a,b) > 0$ (else replace g by $-g$). Since g_y is continuous at (a,b), there are $\delta_0, \epsilon_0 > 0$ such that the closed box $\overline{B}(a,\delta_0) \times [b-\epsilon_0, b+\epsilon_0]$ lies in Ω and $g_y > 0$ throughout it.

Existence and uniqueness. For fixed $x \in B(a,\delta_0)$, the one-variable function $y \mapsto g(x,y)$ is strictly increasing on $[b-\epsilon_0, b+\epsilon_0]$: for $y < y'$ in this interval, the scalar MVT gives $g(x,y') - g(x,y) = g_y(x,\xi)(y'-y) > 0$ for some $\xi \in (y,y')$. Taking $x = a$ and using $g(a,b) = 0$ gives $g(a,b-\epsilon_0) < 0 < g(a,b+\epsilon_0)$. By continuity of $g(\cdot, b \pm \epsilon_0)$ at a, there is $\delta_1 \in (0,\delta_0]$ with $g(x,b-\epsilon_0) < 0 < g(x,b+\epsilon_0)$ for all $x \in B(a,\delta_1)$. Put $U := B(a,\delta_1)$ and $V := (b-\epsilon_0, b+\epsilon_0)$. For $x \in U$ the continuous function $y \mapsto g(x,y)$ changes sign on $[b-\epsilon_0, b+\epsilon_0]$, so by the IVT it has a root, which lies in V and is the only root in $[b-\epsilon_0, b+\epsilon_0]$ by strict monotonicity. Calling it $\varphi(x)$ gives $\varphi : U \to V$ with $\varphi(a) = b$ and, for $(x,y) \in U \times V$, $g(x,y) = 0 \iff y = \varphi(x)$.

Continuity. Fix $x_0 \in U$, $y_0 = \varphi(x_0)$, and $\eta > 0$ with $[y_0-\eta, y_0+\eta] \subset V$. As above, $g(x_0, y_0-\eta) < 0 < g(x_0, y_0+\eta)$, and by continuity there is $\delta_\eta > 0$ with $B(x_0,\delta_\eta) \subset U$ such that these strict inequalities persist for all $x \in B(x_0,\delta_\eta)$. The IVT and monotonicity, applied on $[y_0-\eta, y_0+\eta]$, then force $\varphi(x) \in (y_0-\eta, y_0+\eta)$ for such x. Hence, φ is continuous on U.

Differentiability and formula (2). Fix $x_0 \in U$ and a coordinate direction e_i. For small $t \ne 0$,

$$0 = [g(x_0+te_i, \varphi(x_0+te_i)) - g(x_0+te_i, \varphi(x_0))] + [g(x_0+te_i, \varphi(x_0)) - g(x_0, \varphi(x_0))].$$

The scalar MVT, applied to each bracket (in y, then in the single coordinate x_i), gives η_t between $\varphi(x_0)$ and $\varphi(x_0+te_i)$, and ξ_t on the segment from x_0 to x_0+te_i, with

$$\frac{\varphi(x_0+te_i) - \varphi(x_0)}{t} = -\frac{g_{x_i}(\xi_t, \varphi(x_0))}{g_y(x_0+te_i, \eta_t)},$$

the denominator being positive on the box. As $t \to 0$, $\xi_t \to x_0$ and, by continuity of φ, $\eta_t \to \varphi(x_0)$; continuity of g_{x_i} and g_y then gives (2).

C^k **regularity.** The right-hand side of (2) is a continuous function of x, so $\varphi \in C^1(U)$. If $\varphi \in C^j(U)$ for some $j < k$, then the right-hand side of (2) is a composition of the C^{k-1} functions g_{x_i}, g_y with the C^j map $x \mapsto (x, \varphi(x))$, hence of class C^j; so every $\partial\varphi/\partial x_i \in C^j$ and $\varphi \in C^{j+1}$. Induction on j gives $\varphi \in C^k(U)$. □

3.2. The General (Vector) Case

Lemma 1 (Schur complement). *Let* $M = \begin{pmatrix} A & \beta \\ \gamma^\top & \alpha \end{pmatrix} \in \mathbb{R}^{m\times m}$ *with* $A \in \mathbb{R}^{(m-1)\times(m-1)}$, $\beta, \gamma \in \mathbb{R}^{m-1}$ *and a scalar* $\alpha \neq 0$. *Then,* $\det M = \alpha \det S$, *where* $S := A - \alpha^{-1}\beta\gamma^\top$ *is the Schur complement of* α *in* M. *In particular, if* M *is invertible, then so is* S.

Proof. A direct multiplication gives $\begin{pmatrix} I_{m-1} & -\alpha^{-1}\beta \\ 0 & 1 \end{pmatrix} M = \begin{pmatrix} S & 0 \\ \gamma^\top & \alpha \end{pmatrix}$. The left factor has determinant 1, and the right-hand side is block lower-triangular, so $\det M = \alpha \det S$. □

Theorem 2 (General Implicit Function Theorem). *Let* $\Omega \subset \mathbb{R}^n \times \mathbb{R}^m$ *be open,* $F \in C^k(\Omega; \mathbb{R}^m)$, $k \geq 1$, *and* $(a, b) \in \Omega$ *with* $F(a, b) = 0$ *and* $\partial_y F(a, b)$ *invertible. Then there exist open sets* $U \ni a$ *and* $V \ni b$ *with* $U \times V \subset \Omega$, *and* $\varphi \in C^k(U; V)$ *with* $\varphi(a) = b$, *such that for all* $(x, y) \in U \times V$

$$F(x, y) = 0 \iff y = \varphi(x).$$

Moreover, $\partial_y F(x, \varphi(x))$ *is invertible for* $x \in U$, *and*

$$D\varphi(x) = -\big[\partial_y F(x, \varphi(x))\big]^{-1} \partial_x F(x, \varphi(x)) \in \mathbb{R}^{m\times n}. \tag{3}$$

Proof. We argue by induction on m; the case $m = 1$ is Theorem 1. Let $m \geq 2$ and assume the theorem for $m - 1$ dependent variables.

Step 1 (pivot). Because $\partial_y F(a, b)$ is invertible, its last row is not zero, so $\partial F_m/\partial y_j(a, b) \neq 0$ for some j. Relabeling the dependent variables (a permutation of the coordinates of y, which changes neither the hypotheses nor the conclusions) we may take $j = m$. Write $y = (y', y_m) \in \mathbb{R}^{m-1} \times \mathbb{R}$, $b = (b', b_m)$, $F = (F', F_m)$, and partition

$$\partial_y F(a, b) = \begin{pmatrix} A & \beta \\ \gamma^\top & \alpha \end{pmatrix}, \qquad \begin{aligned} A &= \partial_{y'} F'(a, b), & \beta &= \partial_{y_m} F'(a, b), \\ \gamma^\top &= \partial_{y'} F_m(a, b), & \alpha &= \partial_{y_m} F_m(a, b) \neq 0. \end{aligned}$$

Step 2 (solve the last equation). Apply Theorem 1 to the single equation $F_m(x, y', y_m) = 0$, with $(x, y') \in \mathbb{R}^{n+m-1}$ independent and y_m dependent, at the point $((a, b'), b_m)$. This gives an open set $W \ni (a, b')$, an open interval $I \ni b_m$ with $W \times I \subset \Omega$, and $\psi \in C^k(W; I)$ such that, for $((x, y'), y_m) \in W \times I$, $F_m(x, y', y_m) = 0 \iff y_m = \psi(x, y')$. By (2), $\partial_{y'}\psi(a, b') = -\alpha^{-1}\gamma^\top$.

Step 3 (reduced system; the Schur complement). Define $G(x, y') := F'\big(x, y', \psi(x, y')\big)$ on W, so $G \in C^k(W; \mathbb{R}^{m-1})$ and $G(a, b') = F'(a, b) = 0$. By the chain rule,

$$\partial_{y'} G(a, b') = A + \beta\, \partial_{y'}\psi(a, b') = A - \alpha^{-1}\beta\gamma^\top = S,$$

the Schur complement of α in $\partial_y F(a, b)$, which is invertible by Lemma 1.

Step 4 (inductive hypothesis). The theorem for $m - 1$ dependent variables, applied to G on the open set W at (a, b'), yields open sets $U \ni a$ and $V' \ni b'$ with $U \times V' \subset W$, and $\varphi' \in C^k(U; V')$ with $\varphi'(a) = b'$, such that for $(x, y') \in U \times V'$, $G(x, y') = 0 \iff y' = \varphi'(x)$.

Step 5 (assembly). Put $V := V' \times I$ and $\varphi(x) := \big(\varphi'(x), \psi(x, \varphi'(x))\big)$ for $x \in U$. Then, $U \times V \subset \Omega$ (because $U \times V' \subset W$ and $W \times I \subset \Omega$), $\varphi \in C^k(U; V)$, and $\varphi(a) = b$.

Let $(x, y) \in U \times V$. If $F(x, y) = 0$, then $F_m(x, y', y_m) = 0$ with $((x, y'), y_m) \in W \times I$, so $y_m = \psi(x, y')$; consequently, $0 = F'(x, y) = G(x, y')$ with $(x, y') \in U \times V'$, so $y' = \varphi'(x)$, and therefore $y = \varphi(x)$. Conversely, $F_m(x, \varphi(x)) = 0$ by the definition of ψ, and $F'(x, \varphi(x)) = G(x, \varphi'(x)) = 0$ by the definition of φ'.

Step 6 (Formula (3)). The function $x \mapsto \det \partial_y F(x, \varphi(x))$ is continuous and non-zero at a; shrinking U to a smaller open neighborhood of a (which preserves everything established above), it is non-zero on U. Differentiating $F(x, \varphi(x)) \equiv 0$ by the chain rule gives $\partial_x F + \partial_y F \, D\varphi = 0$, and solving yields (3). □

Figure 1 shows the logical dependency among the main results of this section and the applied section in which each is used.

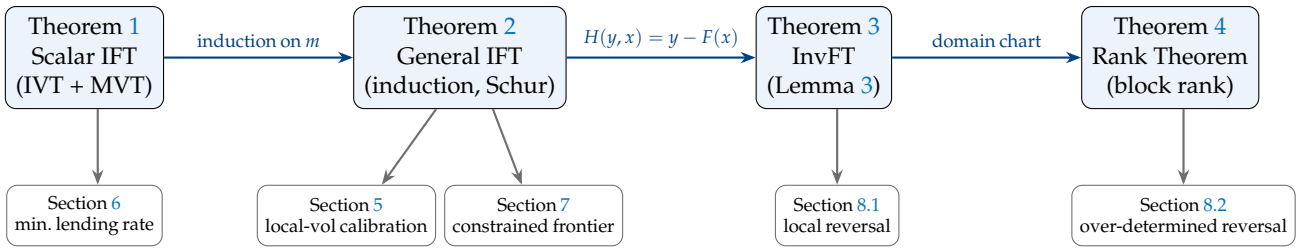


Figure 1. Logical dependency among Theorems 1–4 (top row; the key tool of each proof is noted in its box) and the four balance-sheet applications of Sections 5–8 that invoke them (bottom row).

3.3. The Inverse Function Theorem

Lemma 2 (Vector Mean Value Inequality). *If $\chi : [0, 1] \to \mathbb{R}^n$ is differentiable with $\|\chi'(t)\| \le M$ for all $t \in [0, 1]$, then $\|\chi(1) - \chi(0)\| \le M$.*

Proof. If $\chi(1) = \chi(0)$, there is nothing to prove. Otherwise, let $u := (\chi(1) - \chi(0))/\|\chi(1) - \chi(0)\|$ and $h(t) := \langle \chi(t), u \rangle$, a differentiable real function with $h'(t) = \langle \chi'(t), u \rangle$. By the scalar MVT, $h(1) - h(0) = h'(\xi)$ for some $\xi \in (0, 1)$, so $\|\chi(1) - \chi(0)\| = h(1) - h(0) = \langle \chi'(\xi), u \rangle \le \|\chi'(\xi)\| \le M$ by the Cauchy–Schwarz inequality and $\|u\| = 1$. □

Lemma 3 (Quantitative local injectivity). *Let $\Omega \subset \mathbb{R}^n$ be open, $F \in C^1(\Omega; \mathbb{R}^n)$, $p \in \Omega$, and suppose $DF(p)$ is invertible. Put*

$$c := \frac{1}{2\,\|DF(p)^{-1}\|_F},$$

and choose $r > 0$ such that $\overline{B}(p, r) \subset \Omega$ and $\|DF(x) - DF(p)\|_F \le c$ for all $x \in B(p, r)$. Then

$$\|F(x_1) - F(x_2)\| \;\ge\; c\,\|x_1 - x_2\| \qquad \text{for all } x_1, x_2 \in B(p, r). \tag{4}$$

In particular, F is injective on $B(p, r)$.

Proof. The radius r exists because the entries of DF are continuous at p, and $\|\cdot\|_F$ is the Euclidean norm of the entries. We use two estimates.

(i) *Lower bound for the linearization.* Since $DF(p)$ is invertible (the hypothesis of the lemma, and of Theorem 3), every vector $v \in \mathbb{R}^n$ satisfies, by (1),

$$\|v\| = \|DF(p)^{-1} DF(p) v\| \le \|DF(p)^{-1}\|_F \, \|DF(p) v\|, \qquad \text{that is,} \qquad \|DF(p) v\| \;\ge\; 2c\,\|v\|.$$

(ii) *Linearization error.* Fix $x_1, x_2 \in B(p, r)$, let $\Delta := x_2 - x_1$, and set $\chi(t) := F(x_1 + t\Delta) - t\,DF(p)\Delta$ for $t \in [0, 1]$; the segment $x_1 + t\Delta$ stays in the convex set $B(p, r)$.

Then, $\chi'(t) = \big[DF(x_1 + t\Delta) - DF(p)\big]\Delta$, so $\|\chi'(t)\| \le c\|\Delta\|$ by (1) and the choice of r. Lemma 2 gives

$$\big\|F(x_2) - F(x_1) - DF(p)\Delta\big\| = \|\chi(1) - \chi(0)\| \le c\,\|\Delta\|.$$

Conclusion. By the triangle inequality, then (i) and (ii),

$$\|F(x_2) - F(x_1)\| \;\ge\; \|DF(p)\Delta\| - \big\|DF(p)\Delta - \big(F(x_2) - F(x_1)\big)\big\| \;\ge\; 2c\|\Delta\| - c\|\Delta\| \;=\; c\,\|\Delta\|.$$

If $F(x_1) = F(x_2)$, this forces $\Delta = 0$. □

Theorem 3 (Inverse Function Theorem). *Let $\Omega \subset \mathbb{R}^n$ be open, $F \in C^k(\Omega;\mathbb{R}^n)$, $k \ge 1$, and $p \in \Omega$ with $DF(p)$ invertible. Then there exist open $U \ni p$ and open $V \ni F(p)$ such that $F|_U : U \to V$ is a C^k-diffeomorphism, and for $q = F(x) \in V$,*

$$D(F^{-1})(q) = [DF(x)]^{-1}. \tag{5}$$

Proof. Let $q := F(p)$ and apply Theorem 2 to $H(y,x) := y - F(x)$ on $\mathbb{R}^n \times \Omega$ at (q,p), with the roles of the letters interchanged: y is now the independent and x the dependent variable. Since $\partial_x H(q,p) = -DF(p)$ is invertible, there are an open set $V_0 \ni q$ and a C^k map $\phi : V_0 \to \mathbb{R}^n$ with $\phi(q) = p$ and $F(\phi(y)) = y$ for all $y \in V_0$.

Let r be as in Lemma 3, so that F is injective on $B(p,r)$. Since ϕ is continuous and $\phi(q) = p$, there is an open $V \ni q$, $V \subset V_0$, with $\phi(V) \subset B(p,r)$. Define

$$U := \{x \in B(p,r) : F(x) \in V\},$$

which is open (by continuity of F) and contains p.

For $x \in U$, the point $\phi(F(x))$ is defined, lies in $B(p,r)$, and satisfies $F(\phi(F(x))) = F(x)$; since x also lies in $B(p,r)$ and F is injective there, $\phi(F(x)) = x$. For $y \in V$, the point $\phi(y)$ lies in $B(p,r)$ and $F(\phi(y)) = y \in V$, so $\phi(y) \in U$; hence $F(U) = V$. Thus, $F|_U : U \to V$ is a bijection with inverse $\phi|_V \in C^k$, i.e., a C^k-diffeomorphism, and (5) follows by differentiating $F(F^{-1}(q)) = q$. □

3.4. The Rank Theorem

Lemma 4 (Block rank). *For $C \in \mathbb{R}^{(m-r)\times r}$ and $D \in \mathbb{R}^{(m-r)\times(n-r)}$, the $m \times n$ matrix $\begin{pmatrix} I_r & 0 \\ C & D \end{pmatrix}$ has rank $r + \operatorname{rank}(D)$.*

Proof. Left multiplication by the invertible $m \times m$ matrix $\begin{pmatrix} I_r & 0 \\ -C & I_{m-r} \end{pmatrix}$, which preserves rank, produces $\begin{pmatrix} I_r & 0 \\ 0 & D \end{pmatrix}$, whose rank is $r + \operatorname{rank}(D)$. □

Lemma 5 (Zero partial derivative on a convex domain). *If $h \in C^1(\Omega;\mathbb{R}^\ell)$ on a convex open $\Omega \subset \mathbb{R}^p \times \mathbb{R}^q$ has $\partial h/\partial v \equiv 0$ throughout Ω (writing points as (u,v)), then $h(u,v)$ does not depend on v.*

Proof. Fix u and v_1, v_2 with $(u,v_1), (u,v_2) \in \Omega$; by convexity the segment between them lies in Ω. The derivative of $t \mapsto h(u, v_1 + t(v_2 - v_1))$ is $\partial h/\partial v\,(u, v_1 + t(v_2 - v_1))\,(v_2 - v_1) = 0$, so Lemma 2 (with $M = 0$) gives $h(u,v_1) = h(u,v_2)$. □

Lemma 6 (Maximal rank is an open condition). *Let $\Omega \subset \mathbb{R}^n$ be open, $F \in C^1(\Omega;\mathbb{R}^m)$, $p \in \Omega$, and* $\operatorname{rank} DF(p) = s := \min(n,m)$. *Then* $\operatorname{rank} DF(x) = s$ *for all x in some open ball around p.*

Proof. Since $\operatorname{rank} DF(p) = s$, some $s \times s$ minor $\mu(x)$ of $DF(x)$ satisfies $\mu(p) \neq 0$. As a polynomial in the entries of $DF(x)$, which are continuous, μ is continuous, hence non-zero on a ball around p; there $\operatorname{rank} DF(x) \geq s$. And $\operatorname{rank} DF(x) \leq s$ always, because $DF(x)$ is $m \times n$. □

Lemma 6 fails at points of non-maximal rank (for $F(x) = x^3$ on $\mathbb{R}$, $\operatorname{rank} DF = 0$ at 0 but 1 elsewhere), which is why the Rank Theorem below assumes constant rank near p. At points of maximal rank—immersions ($r = n < m$) and submersions ($r = m < n$)—that hypothesis holds automatically.

Theorem 4 (Rank Theorem). *Let $\Omega \subset \mathbb{R}^n$ be open, $F \in C^k(\Omega;\mathbb{R}^m)$ with $k \geq 1$, and $p \in \Omega$. Suppose* $\operatorname{rank} DF(x) = r$ *for all x in some open neighborhood $\Omega_0 \subset \Omega$ of p. Then there exist open sets $U \ni p$ and $V \ni F(p)$ with $F(U) \subset V$, and C^k-diffeomorphisms $\phi : U \to \widetilde{U} \subset \mathbb{R}^n$ and $\psi : V \to \widetilde{V} \subset \mathbb{R}^m$, such that, in the new coordinates $z = \phi(x)$,*

$$(\psi \circ F \circ \phi^{-1})(z) = \pi_r(z) := (z_1, \ldots, z_r, 0, \ldots, 0) \qquad \text{for all } z \in \widetilde{U}; \tag{6}$$

equivalently, in the original coordinates, $F(x) = \psi^{-1}\big(\pi_r(\phi(x))\big)$ for all $x \in U$.

Proof. If $r = 0$, then $DF \equiv 0$ on a ball $U \subset \Omega_0$ around p, so $F \equiv F(p)$ on U by Lemma 5 (with no u variables), and $\phi = \mathrm{id}$, $\psi(y) = y - F(p)$ work. Assume $r \geq 1$.

Step 0 (permutations). Some $r \times r$ minor of $DF(p)$ is non-zero. Choose permutation matrices $P \in \mathbb{R}^{n\times n}$ and $Q \in \mathbb{R}^{m\times m}$ that move the corresponding columns and rows to the top-left corner, and put $\widehat{F}(\xi) := QF(P\xi)$. If (6) holds for $\widehat{F}$ at $P^{-1}p$ with charts $\widehat{\phi}, \widehat{\psi}$, then it holds for F at p with $\phi := \widehat{\phi} \circ P^{-1}$ and $\psi := \widehat{\psi} \circ Q$, because $\psi \circ F \circ \phi^{-1} = \widehat{\psi} \circ \widehat{F} \circ \widehat{\phi}^{-1}$; and the rank of the Jacobian is unchanged. So we may assume that $\partial(F_1, \ldots, F_r)/\partial(x_1, \ldots, x_r)$ is invertible at p.

Step 1 (domain chart). Write $x = (x', x'') \in \mathbb{R}^r \times \mathbb{R}^{n-r}$ and $F = (F', F'')$ with $F' = (F_1, \ldots, F_r)$, and define $\Phi(x) := \big(F'(x), x''\big)$. Its Jacobian at p is block upper-triangular, $\begin{pmatrix} \partial_{x'}F'(p) & \partial_{x''}F'(p) \\ 0 & I_{n-r} \end{pmatrix}$, with determinant $\det \partial_{x'}F'(p) \neq 0$. By Theorem 3, Φ maps an open set $U_0 \ni p$, which we may take inside Ω_0, diffeomorphically onto an open set containing $\Phi(p) =: (u_0, v_0)$. Choose open balls $B_1 \subset \mathbb{R}^r$ about u_0 and $B_2 \subset \mathbb{R}^{n-r}$ about v_0 with $B_1 \times B_2 \subset \Phi(U_0)$, and put $\widetilde{U} := B_1 \times B_2$ (convex), $U := (\Phi|_{U_0})^{-1}(\widetilde{U})$ and $\phi := \Phi|_U$.

Step 2 (F in the new coordinates). For $(u, v) \in \widetilde{U}$ the first r components of $F(\phi^{-1}(u, v))$ equal u, because the first r components of Φ are F'. Hence, $F(\phi^{-1}(u, v)) = \big(u, h(u, v)\big)$ for some $h \in C^k(\widetilde{U};\mathbb{R}^{m-r})$.

Step 3 (where constant rank is used). For $z \in \widetilde{U}$, $D(F \circ \phi^{-1})(z) = DF(\phi^{-1}(z))\, D\phi^{-1}(z)$ with $D\phi^{-1}(z)$ invertible, so its rank equals $\operatorname{rank} DF(\phi^{-1}(z)) = r$, because $\phi^{-1}(z) \in U \subset \Omega_0$. This is the only use of the constant-rank hypothesis. On the other hand, $D(F \circ \phi^{-1}) = \begin{pmatrix} I_r & 0 \\ \partial_u h & \partial_v h \end{pmatrix}$, whose rank is $r + \operatorname{rank} \partial_v h$ by Lemma 4. Hence, $\partial_v h \equiv 0$ on $\widetilde{U}$, and since $\widetilde{U}$ is convex, Lemma 5 gives $h(u, v) = \bar{h}(u) := h(u, v_0)$.

Step 4 (target chart). Let $V := B_1 \times \mathbb{R}^{m-r}$, an open set containing $F(p) = (u_0, \bar{h}(u_0))$, and $\psi(s, w) := \big(s, w - \bar{h}(s)\big)$. Then, ψ is a C^k bijection of V onto $\widetilde{V} := V$ with C^k inverse $(s, w) \mapsto (s, w + \bar{h}(s))$. By Step 2, $F(U) \subset V$, and for $(u, v) \in \widetilde{U}$,

$$\psi\big(F(\phi^{-1}(u, v))\big) = \psi\big(u, \bar{h}(u)\big) = (u, 0) = \pi_r(u, v).$$

(If $r = m$, there is no h and ψ is the identity on $V = B_1$.) □

Corollary 1 (Immersion case)**.** *Let $\Omega \subset \mathbb{R}^n$ be open, $F \in C^k(\Omega; \mathbb{R}^m)$ with $k \geq 1$ and $n < m$, and* $\operatorname{rank} DF(p) = n$ *at a single point $p \in \Omega$. Then there are open sets $U \ni p$ and $V \ni F(p)$ with $F(U) \subset V$ such that F is injective on U and*

$$F(U) = \{y \in V : \psi(y) \in \mathbb{R}^n \times \{0\}\}$$

is an n-dimensional embedded C^k-submanifold of $\mathbb{R}^m$; after a permutation of the components of F it is the graph of a C^k function $\bar{h} : B_1 \to \mathbb{R}^{m-n}$ over an open ball $B_1 \subset \mathbb{R}^n$. In particular, no $x \in U$ is mapped to a point of $V \setminus F(U)$.

Proof. By Lemma 6, $\operatorname{rank} DF \equiv n$ near p, so Theorem 4 applies with $r = n$. There are then no x'' variables, $\widetilde{U} = B_1$, and the proof of Theorem 4 gives $F(\phi^{-1}(u)) = (u, \bar{h}(u))$ for $u \in B_1$ (after the permutation Q). Hence, $F(U)$ is the graph of $\bar{h}$, which equals $\{(s, w) \in V : w = \bar{h}(s)\} = \{y \in V : \psi(y) \in \mathbb{R}^n \times \{0\}\}$, and $F(x_1) = F(x_2)$ forces $\phi(x_1) = \phi(x_2)$, i.e., $x_1 = x_2$. □

Corollary 1 is local: it says nothing about points $x \notin U$. A global conclusion—that a target is attained by *no* $x \in \Omega$—needs additional global information; in Section 8.2, that information is available because the relevant projection of the map is an affine bijection.

We use Theorems 1–3 to produce comparative statics (Sections 5–7), Theorem 3 for local reversal (Section 8.1), and Corollary 1 for over-determined reversal (Section 8.2). We also need second derivatives of implicit functions (Sections 5 and 6).

Proposition 1 (Second-order implicit derivative)**.** *Under the hypotheses of Theorem 2 with $k \geq 2$ and a single independent variable ($n = 1$),*

$$\varphi''(x) = -\left[\partial_y F\right]^{-1}\Big(F_{xx} + 2F_{xy}\,\varphi'(x) + F_{yy}\left[\varphi'(x), \varphi'(x)\right]\Big), \tag{7}$$

where $F_{xy}v := \sum_j \partial^2 F/\partial x \partial y_j\, v_j$, $F_{yy}[v, v] := \sum_{j,l} \partial^2 F/\partial y_j \partial y_l\, v_j v_l$, and all derivatives are evaluated at $(x, \varphi(x))$. For $m = 1$, this reads $\varphi'' = -(g_{xx} + 2g_{xy}\varphi' + g_{yy}\varphi'^2)/g_y$.

Proof. Differentiating $F_x + \partial_y F\,\varphi' = 0$ once more (using $F_{xy} = F_{yx}$, as $F \in C^2$) gives $F_{xx} + 2F_{xy}\varphi' + F_{yy}[\varphi', \varphi'] + \partial_y F\,\varphi'' = 0$. □

4. Balance-Sheet Notation and the General Well-Posedness Question

Throughout, $w \in \mathbb{R}^n$ denotes a vector of balance-sheet actions (portfolio weights, offer-rate components, or calibration parameters, as context dictates), and $u = \Psi(w) \in \mathbb{R}^m$ denotes a vector of attributed, economically interpretable coordinates computed from w by the pricing and attribution operators of arbitrage-free interest-rate theory: economic value, net interest income, duration, attributed spread income, and so on. Four tasks recur throughout balance-sheet management. Table 3 lists them. Each row is read from left to right: the task, the equation whose solution defines it, and the result of Section 3 that gives local existence and uniqueness of that solution together with a formula for its derivative. Which result applies is fixed purely by a dimension count — how many unknowns are being solved for, and whether the count of unknowns matches the count of targets:

Table 3. Four recurring balance-sheet tasks, each written as an equation to be solved, and the result of Section 3 that gives local existence, uniqueness and the derivative of its solution.

Task	As a Root/Fixed-Point Problem	Theorem Invoked
Calibration	solve $F(\theta;$ market data$) = 0$ for model parameters θ; ask how $\theta^\star$ moves with the data	Theorem 1 (scalar) or 2 (vector)
Hurdle/rate solve with feedback	solve $g(R) = 0$ for a rate R that enters its own right-hand side	Theorem 1
Constrained optimization	solve KKT stationarity $F(w, \lambda; \theta) = 0$ for (w, λ); ask how the optimum moves with a policy parameter θ	Theorem 2
Reverse optimization	solve $\Psi(w) = u^\star$ for w given a target $u^\star$	Theorem 3 ($m = n$) or Theorem 4 ($m \neq n$)

The pattern in the right-hand column of Table 3 is not a coincidence of exposition: a single unknown pinned down by its own defining equation (Section 6) is the scalar case; several unknowns pinned down jointly by as many market quotes or first-order conditions (Sections 5 and 7) is the vector case; and asking the question in reverse—given the outputs, recover the inputs (Section 8)—is the InvFT when the count of free actions matches the count of targets and the Rank Theorem the moment it does not. In the applied sources of Section 2, the numerical solve in each row (typically by Newton's method or a KKT linear system) is usually performed without stating the theorem that certifies *that the solve is meaningful*—that the solution is locally unique, depends smoothly on the data, and has a computable derivative. The next four sections state, in each case, which theorem applies, verify its hypothesis numerically at a computed solution, derive the resulting sensitivity formula, and check that formula against central differences. They are worked applications of Section 3, not new theorems.

Numerical conventions.

All computations use IEEE double precision (NumPy/SciPy). Nonlinear systems are solved by Newton's method with analytic Jacobians, iterated until the residual stops decreasing (machine precision). A closed-form sensitivity $\partial y^\star/\partial p$ is checked against the central difference $[y^\star(p+h) - y^\star(p-h)]/(2h)$ of two re-solved systems. Its error is approximately $\frac{h^2}{6}|y^{\star\prime\prime\prime}| + \varepsilon_{\text{mach}}|y^\star|/h$ (truncation plus round-off), which is minimized at $h_{\text{opt}} = (3\varepsilon_{\text{mach}}|y^\star|/|y^{\star\prime\prime\prime}|)^{1/3}$; when derivatives scale naturally with p this is $h \approx \varepsilon_{\text{mach}}^{1/3}|p| \approx 6 \times 10^{-6}|p|$ ([40], Section 8.1). For each check, we evaluated the discrepancy at $h = 10^{-3}, 10^{-4}, \ldots, 10^{-8}$. In every case, it falls by a factor of about 100 per decade of h (the h^2 truncation regime) until round-off takes over, below $h \approx 10^{-6}$ to 10^{-7}.

Table 4 collects every such check made in Sections 5–8. Each row names the closed-form sensitivity being checked, the input p that is perturbed and its magnitude, the step h, and the largest absolute discrepancy between the closed form and the central difference over all components. The discrepancies lie between 2.4×10^{-12} and 1×10^{-6} so every closed form agrees with an independent computation to at least six significant figures; the largest discrepancy is for the second derivative in Section 5, which is checked by a second difference. The step in each row lies in or next to the optimal decade; In Section 7, it is smaller than elsewhere because the third derivative of the solution in $\sigma_{\max}$ is large, so truncation error dominates down to $h = 10^{-7} \approx \varepsilon_{\text{mach}}^{1/3}\sigma_{\max}$. All numbers in the paper are reproduced by the script supplied as Supplementary Material.

Table 4. Every closed-form sensitivity in Sections 5–8 against a central difference of re-solved systems: the perturbed input p, its value, the step h, and the largest absolute discrepancy over all components.

Section	Closed Form	Input p	$\lvert p \rvert$	h	Max. Discrepancy
Section 5	J_{cal}^{-1}, Equation (10)	P_{t+1}, $C(K)$	0.90, 0.0068	10^{-5}	1.2×10^{-9}
Section 5	$\sigma^{\star\prime\prime}(C)$, Propposition 1	$C(K)$	0.0068	10^{-4} (2nd diff.)	1×10^{-6}
Section 6	$1/(1-\theta')$, Equation (12)	κ	0.0075	10^{-6}	2.4×10^{-12}
Section 7	$J_{\text{KKT}}^{-1} e_5$, Equation (17)	$\sigma_{\max}$	0.012	10^{-7}	2.3×10^{-9}
Section 8	Equation (21)	$m^\star$, $D^\star$	0.014, 3.0	10^{-6}	1.0×10^{-8}

5. Application I: Joint Drift–Local-Volatility Slice Calibration

This first application is the first row of Table 3 in its purest vector form: two free parameters, two market quotes, and the question of whether they pin each other down is a direct instance of Theorem 2. It is also the application in which the gap between an existence claim and a usable result is sharpest, since the object under calibration—a local-volatility surface—is built in practice strike by strike by Newton iteration (Section 2.3), without a stated check that each step solves a well-posed problem.

5.1. The Map

Consider one time slice $[t, t+1]$ of a two-state short-rate lattice with state price φ_t already known from the previous slice (bootstrapped by no-arbitrage as in the classical Ho–Lee construction). The two child rates are

$$r_d(a,\sigma) = a - \sigma\sqrt{\Delta}, \qquad r_u(a,\sigma) = a + \sigma\sqrt{\Delta},$$

where a is the local drift and σ the local volatility for this slice, both free. The slice is calibrated jointly to two market quotes: the zero-coupon bond price P_{t+1} (the no-arbitrage condition) and a cap price $C(K)$ struck at K maturing at $t+1$ (the smile/skew condition). Define $F : \mathbb{R}^2 \to \mathbb{R}^2$ by

$$F(a,\sigma) = \begin{pmatrix} \varphi_t\left[\frac{1}{2}e^{-r_d} + \frac{1}{2}e^{-r_u}\right] - P_{t+1} \\ \varphi_t\left[\frac{1}{2}(r_d - K)^+ e^{-r_d} + \frac{1}{2}(r_u - K)^+ e^{-r_u}\right] - C(K) \end{pmatrix}. \tag{8}$$

The positive parts $(r-K)^+$ make F non-differentiable on the two lines $r_d = K$ and $r_u = K$, so the IFT can only be applied away from them. At every solution reported below, $r_d < K < r_u$, so both kinks are inactive in a neighborhood of the solution, where F is C^∞:

$$F_2(a,\sigma) = \tfrac{1}{2}\varphi_t\,(r_u - K)e^{-r_u} - C(K) \quad \text{near the solution.}$$

Calibration is the root $F(a,\sigma) = 0$; Theorem 2, applied with $x = (P_{t+1}, C(K))$ as the independent variable and $y = (a,\sigma)$ as the dependent one, shows that whenever $J_{\text{cal}} := \partial F/\partial(a,\sigma)$ is non-singular at such a solution, the solution is locally unique and depends C^∞ on the two market quotes, with derivative given by (3). The same linear system, written $\mathbf{M}_t\,\delta\sigma_t = \delta\mathbf{C}_t$, is used informally in local-volatility bootstrapping [12]; (3) identifies it as the exact derivative of the calibration map. Figure 2 depicts the slice.

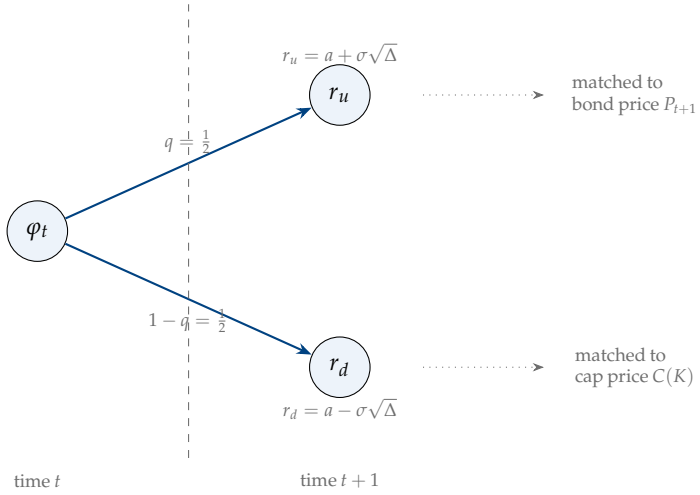


Figure 2. The two-state slice calibrated in Section 5: the two free parameters (a, σ) generate two child rates whose risk-neutral discounted payoffs are matched jointly to a bond price and a cap price, giving the 2×2 system (8).

5.2. Numerical Verification

Take $\varphi_t = 0.936131$ (consistent with a five-year Ho–Lee lattice calibrated to a 3.00–3.90% par curve), $\Delta = 1$, $K = 4.00\%$, $P_{t+1} = 0.898975$, $C(K) = 0.0068$. Newton's method for (8) from $(a_0, \sigma_0) = (4.40\%, 1.10\%)$, with the analytic Jacobian, produces the iterates listed in Table 5: the residual falls from 3×10^{-3} to zero in three steps.

Example 1 (Newton iteration). *The ratio* $\|F(a_{k+1}, \sigma_{k+1})\| / \|F(a_k, \sigma_k)\|^2$ *is* 1.22 *at* $k = 1$ *and* 1.09 *at* $k = 2$*, the quadratic rate predicted by the standard local convergence theorem for Newton's method [19]. The calibrated slice is* $a^\star = 4.0609\%$, $\sigma^\star = 1.4746\%$*, with child rates* $r_d = 2.586\% < K < r_u = 5.535\%$.

k	(a_k, σ_k)	$\|F(a_k, \sigma_k)\|$	*Step Norm*
0	(0.044000, 0.011000)	3.090×10^{-3}	5.071×10^{-3}
1	(0.040596, 0.014758)	1.167×10^{-5}	1.801×10^{-5}
2	(0.040609, 0.014746)	1.479×10^{-10}	2.293×10^{-10}
3	(0.040609, 0.014746)	0	0

At the solution, the analytic Jacobian is

$$J_{\text{cal}} = \tfrac{1}{2}\varphi_t \begin{pmatrix} -e^{-r_d} - e^{-r_u} & \sqrt{\Delta}\,(e^{-r_d} - e^{-r_u}) \\ e^{-r_u}\big(1 - (r_u - K)\big) & \sqrt{\Delta}\,e^{-r_u}\big(1 - (r_u - K)\big) \end{pmatrix} = \begin{pmatrix} -0.898975 & 0.013255 \\ 0.436060 & 0.436060 \end{pmatrix}, \tag{9}$$

with $\det J_{\text{cal}} = -0.39779 \neq 0$ (the entry -0.898975 equals $-P_{t+1}$, as it must at a solution). So Theorem 2 applies. Since $\partial F / \partial (P, C) = -I_2$, Formula (3) gives the sensitivity of the calibrated parameters to a revision of either quote as $-J_{\text{cal}}^{-1}(-I_2) = J_{\text{cal}}^{-1}$:

$$\frac{\partial (a^\star, \sigma^\star)}{\partial (P_{t+1}, C(K))} = \begin{pmatrix} -1.0962146 & 0.0333223 \\ 1.0962146 & 2.2599406 \end{pmatrix}. \tag{10}$$

Example 2 (verification against central differences). *Re-solving* (8) *with* $P_{t+1} \to P_{t+1} \pm h$ *and* $C(K) \to C(K) \pm h$*,* $h = 10^{-5}$*, gives* $\partial(a, \sigma)/\partial P = (-1.0962146, 1.0962146)$ *and* $\partial(a, \sigma)/\partial C = (0.0333223, 2.2599406)$*, agreeing with* (10) *to within* 1.2×10^{-9}*. In particular* $\partial \sigma^\star / \partial C(K) = 2.260$*: a* 1 *basis-point upward revision of the quoted cap premium moves the calibrated local volatility by* 2.26 *basis points, a figure that* (3) *delivers without re-calibration.*

Figure 3 (blue circles) plots the residual norm against iteration count for this example on a logarithmic scale.

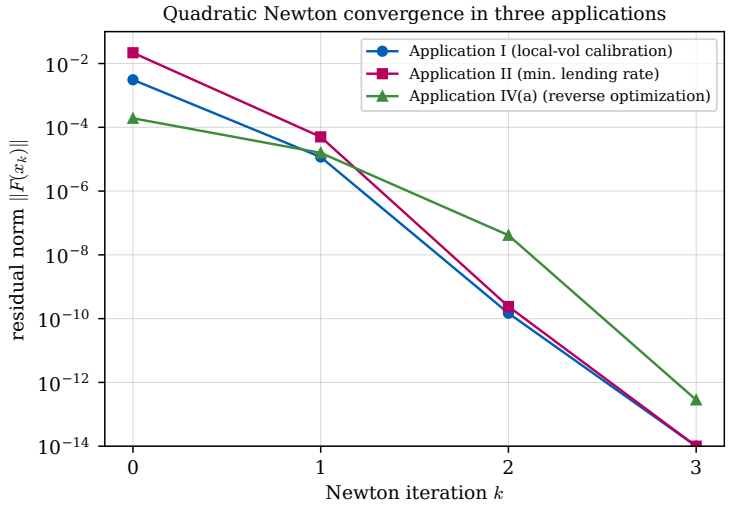


Figure 3. Newton residual norm $\|F(x_k)\|$ versus iteration k (logarithmic vertical scale) for the calibration of Section 5 (blue circles), the minimum-lending-rate solve of Section 6 (magenta squares), and the local reversal of Section 8.1 (green triangles). All three exhibit the characteristic doubling of correct digits per iteration of quadratic convergence; residuals below 10^{-14} are plotted at 10^{-14} (machine-precision floor).

5.3. Second-Order Sensitivity and Robustness Across Strikes

Proposition 1 gives the curvature of the calibration map—how much the first-order sensitivity (10) itself changes as the cap quote moves by more than an infinitesimal amount. Write $\phi(C) := \sigma^\star(P_{t+1}, C)$ for fixed P_{t+1}. Here, $x = C$, $y = (a, \sigma)$, $F_{xx} = 0$ and $F_{xy} = 0$ (the quote enters (8) additively), so (7) reduces to $\partial_C^2(a^\star, \sigma^\star) = -J_{\text{cal}}^{-1} F_{yy}[v, v]$ with $v = J_{\text{cal}}^{-1} e_2$ the second column of (10). Evaluating the second derivatives of F analytically gives $\phi''(C(K)) = 5.414069$.

Example 3 (second-order sensitivity)**.** *Example 3 lists central second differences of ϕ at three step sizes, each value $\phi(C \pm h)$ obtained by re-solving* (8) *to machine precision, give*

h	$\phi(C+h)$	$\phi(C-h)$	$[\phi(C+h) - 2\phi(C) + \phi(C-h)]/h^2$
10^{-3}	0.0170085	0.0124886	5.414162
10^{-4}	0.0149718	0.0145199	5.414070
10^{-5}	0.0147684	0.0147232	5.414072

At $h = 10^{-3}$, the $O(h^2)$ truncation error is visible (9×10^{-5}); at $h = 10^{-4}$, the estimate agrees with the analytic value to 1×10^{-6}; at $h = 10^{-5}$, round-off in the re-solved values, of order $\varepsilon_{\text{mach}} \|J_{\text{cal}}^{-1}\|_F / h^2 \approx 6 \times 10^{-6}$, begins to dominate. The positive sign says the calibration map is convex in the quoted cap price near $C(K)$, so the linear sensitivity (10) understates *the local-volatility response to a large upward revision of $C(K)$ and* overstates *it for a large downward one.*

Table 5 tests robustness against strike choice using a genuinely skewed synthetic market: "true" cap prices are generated from a skewed two-state model with up-probability $q_{\text{true}} = 0.35$ (rather than the symmetric $q = \frac{1}{2}$ our fitting model (8) assumes) and $\sigma_{\text{true}} = 1.60\%$, with its own drift recalibrated so that it too reprices the same $P_{t+1} = 0.898975$ exactly. Our symmetric two-parameter model (8) is then re-solved against $(P_{t+1}, C_{\text{true}}(K))$ independently at each strike.

Table 5. Calibrated slice and Jacobian determinant across a range of cap strikes K; market cap prices $C(K)$ generated by a skewed ($q_{\text{true}} = 0.35$) true model and refit, strike by strike, with the symmetric two-parameter model (8), $P_{t+1} = 0.898975$ fixed throughout.

K	$C_{\text{true}}(K)$	$a^\star$	$\sigma^\star$	$\det J_{\text{cal}}$
3.00%	0.009680	4.0563%	1.1217%	−0.3952
3.50%	0.008140	4.0582%	1.2761%	−0.3966
4.00%	0.006599	4.0602%	1.4292%	−0.3980
4.50%	0.005058	4.0625%	1.5809%	−0.3994
5.00%	0.003518	4.0650%	1.7314%	−0.4007

At every strike $r_d < K < r_u$ (so F is smooth near the fit) and the determinant stays well away from zero ($|\det J_{\text{cal}}| \in [0.395, 0.401]$), so Theorem 2 gives a locally unique, smoothly varying fit at each one; but $\sigma^\star$ rises monotonically from 1.12% to 1.73% as K rises from 3% to 5% —the model is correctly signaling, through the instability of its own fitted parameter across strikes, that a single symmetric two-state slice cannot represent the skewed truth it is being asked to match. This is the mechanism, made numerically explicit, behind the informal practitioner's observation that "a smile means the model is misspecified": with only two free parameters, a skewed generating process forces a strike-dependent fit, and the L-parameter generalization of Section 2.3 (matching L strikes with L state-indexed local parameters simultaneously) is the natural remedy.

Remark 1. *Where the slice carries $L > 2$ free local-volatility parameters (one per grid point, matched to L quoted strikes), $F : \mathbb{R}^L \to \mathbb{R}^L$ and Theorem 2 applies verbatim with $J_{\text{cal}} \in \mathbb{R}^{L\times L}$ in the role of the matrix $\mathbf{M}_t$ used in local-volatility bootstrapping [12]. Near-singularity of $\mathbf{M}_t$ (thin quotation in the deep tail) is the failure of the hypothesis of Theorem 2, which is why the tail of $\sigma(\cdot, t)$ is governed by extrapolation rather than calibration.*

6. Application II: Minimum Lending Rate Under Option-Cost Feedback

Where Application I used the vector IFT, this application isolates the scalar case, and for that reason, is the one in which the fold-bifurcation phenomenon of Section 1.3 can be exhibited with full transparency: a single equation in a single unknown either has $g_R \neq 0$ at its root, in which case Theorem 1 applies outright, or it does not, in which case the theorem's hypothesis and the economic concept it certifies fail at the very same instant.

6.1. The Map

The standard risk-based pricing decomposition treats the offer rate as an additive sum $R = f + \theta + \kappa + v + x + m$ of funding cost, option cost, expected credit loss, unexpected-loss capital charge, operating cost, and target margin, each computed independently of R. In reality, the prepayment option cost θ is itself a function of the coupon rate being offered: a higher offer rate increases the borrower's future refinancing incentive and hence the value of the option being granted. Model this feedback by $\theta = \theta(R)$ and write the minimum lending rate $R^\star$ as the root of the scalar equation

$$g(R) := R - \big[f + \theta(R) + \kappa + v + x + m\big] = 0. \tag{11}$$

Provided $g_R(R^\star) = 1 - \theta'(R^\star) \neq 0$, Theorem 1 (applied with the pricing components as the exogenous parameter vector) gives $R^\star$ as a locally unique, C^k function of (f, κ, v, x, m) with, by (2),

$$\frac{\partial R^\star}{\partial \kappa} = \frac{1}{1 - \theta'(R^\star)} > 1 \quad \text{whenever } \theta' > 0, \tag{12}$$

strictly larger than the naive value 1 obtained by treating θ as fixed. Every other pricing component (funding cost, capital charge, operating cost, target margin) inherits the identical amplification factor, since each enters (11) exactly as κ does.

Proposition 2 (Fold point and breakdown of the minimum lending rate)**.** *Let $\theta = \theta(R;\eta)$ depend on a feedback-strength parameter η, let $g(R;\eta)$ be (11) with this θ, and assume $g \in C^2$ near a point (R_c, η_c) at which*

$$g = 0, \qquad g_R = 0, \qquad g_{RR} \neq 0, \qquad g_\eta \neq 0.$$

Then there are $\delta, \varepsilon > 0$ and a C^2 function $H : (R_c - \delta, R_c + \delta) \to \mathbb{R}$ with $H(R_c) = \eta_c$ such that, for $|R - R_c| < \delta$ and $|\eta - \eta_c| < \varepsilon$,

$$g(R;\eta) = 0 \iff \eta = H(R), \qquad H'(R_c) = 0, \qquad H''(R_c) = -\frac{g_{RR}}{g_\eta}. \tag{13}$$

If $H''(R_c) < 0$, then $\eta_c - H(R) = \frac{1}{2}|H''(R_c)|(R - R_c)^2 + o\big((R - R_c)^2\big)$, so in this neighborhood, (11) has two roots for η slightly below η_c, one at $\eta = \eta_c$, and none for $\eta > \eta_c$. Along the lower branch, $R < R_c$ the amplification factor $1/g_R = 1/(1 - \theta')$ of (12) and the slope $dR^\star/d\eta = -g_\eta/g_R$ both diverge as $\eta \to \eta_c^-$.

Proof. Since $g_\eta(R_c, \eta_c) \neq 0$, Theorem 1 applies with the roles of the variables exchanged (R independent, η dependent) and gives H together with the equivalence in (13). Differentiating $g(R, H(R)) \equiv 0$ once gives $g_R + g_\eta H' = 0$, so $H'(R_c) = 0$; differentiating again and evaluating at R_c, where $H' = 0$, leaves $g_{RR} + g_\eta H'' = 0$ (this is Proposition 1 with $m = 1$). Taylor's theorem then gives the quadratic expansion, and the root count follows because H has a strict local maximum η_c at R_c. Finally, $g_R(R_c, \eta_c) = 0$ and $g_{RR} \neq 0$ give $g_R(R, H(R)) = g_{RR}(R_c, \eta_c)(R - R_c) + o(R - R_c)$, which tends to 0 along the branch, while g_η stays bounded away from 0. □

Equivalently, near (R_c, η_c), $g(R;\eta) = \frac{1}{2} g_{RR}(R - R_c)^2 + g_\eta(\eta - \eta_c) +$ higher-order terms, which after rescaling R and η is the normal form $(R - R_c)^2 - (\eta - \eta_c)$ of a saddle-node (fold) bifurcation; the conditions $g_{RR} \neq 0$ and $g_\eta \neq 0$ are its standard non-degeneracy and transversality conditions ([41], Chapter 3). The proof above uses only Theorem 1.

6.2. Numerical Verification

Take $f = 3.00\%$, $\kappa = 0.75\%$, $v = 0.60\%$, $x = 0.40\%$, $m = 1.50\%$, and an exponential feedback $\theta(R) = \theta_0 e^{\eta(R - R_0)}$ with $\theta_0 = 0.45\%$, $\eta = 6.0$, $R_0 = 5.00\%$ (a locally calibrated proxy for the convex rise in prepayment-option value with coupon). Newton's method for (11) from the starting value $R = 9\%$ produces the iterates listed in Example 4: Newton iterates for the minimum-lending-rate equation (11): iterate, residual and Newton step.

Example 4 (Newton iteration)**.** *$R^\star = 6.7498\%$, $\theta(R^\star) = 0.4998\%$, $\theta'(R^\star) = \eta\,\theta(R^\star) = 0.029989$. Here, g is strictly concave in R ($g_{RR} = -\eta^2\theta < 0$), so it has at most two roots; $R^\star$ is the lower one, where $g_R > 0$.*

k	R_k	$g(R_k)$	Step
0	9.0000%	2.178×10^{-2}	-2.255×10^{-2}
1	6.7447%	-5.009×10^{-5}	5.164×10^{-5}
2	6.7498%	-2.399×10^{-10}	2.473×10^{-10}
3	6.7498%	0	0

By (12), $\partial R^\star/\partial\kappa = 1/(1-0.029989) = 1.030916$.

Example 5 (verification against central differences). *Re-solving* (11) *at* $\kappa \pm h$, $h = 10^{-6}$, *by Newton's method to machine precision and forming the central difference gives* $\partial R^\star/\partial\kappa = 1.0309159871$, *agreeing with the closed form* 1.0309159871 *to* 2.4×10^{-12}. *The feedback loop raises the sensitivity of the minimum lending rate to expected credit loss by* 3.09% *relative to the feedback-blind value* 1 *implied by an additive decomposition of the kind described in [11]; if* κ *moves by* 50 *basis points through a credit cycle, the additive decomposition understates the required repricing by* 1.5 *basis points. The effect is small here because* $\eta\theta(R^\star)$ *is small, but it grows without bound as* $\theta'(R^\star) \to 1^-$.

Figure 4 plots $\theta(R)$ and the affine function $R - (f + \kappa + v + x + m)$ together; their unique intersection is $R^\star$ (the second, upper intersection lies off the plotted range). A steeper $\theta(R)$ (larger η) raises $R^\star$ and moves the two intersections towards each other until they merge at a tangency, which is the geometric content of Proposition 2.

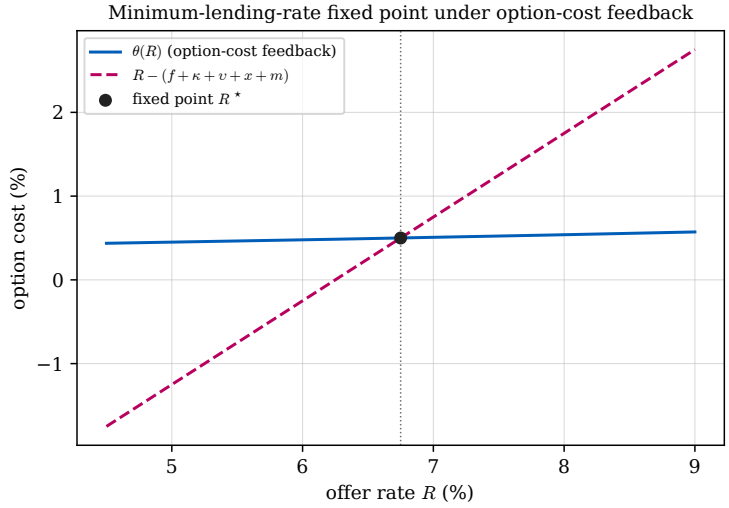


Figure 4. The minimum lending rate as the intersection of the option-cost feedback curve $\theta(R)$ (solid) and the affine residual $R - (f + \kappa + v + x + m)$ (dashed), base case $\eta = 6.0$; the marked point is $R^\star = 6.7498\%$ from Example 4.

Example 6 (locating the fold). *Holding* θ_0, R_0 *fixed and increasing* η *along the exponential family of Example 4, solving* $g(R;\eta) = 0$ *and* $g_R(R;\eta) = 1 - \eta\theta_0 e^{\eta(R-R_0)} = 0$ *simultaneously by Newton's method locates the fold at*

$$\eta_c = 46.001, \qquad R_c = 8.424\%, \qquad \theta(R_c) = 1/\eta_c = 2.174\%.$$

There $g_{RR} = -\eta_c^2\theta(R_c) = -\eta_c = -46.001$ *and* $g_\eta = -(R_c - R_0)\theta(R_c) = -7.443 \times 10^{-4}$, *both non-zero, so Proposition 2 applies with* $H''(R_c) = -g_{RR}/g_\eta = -61{,}804 < 0$. *Its quadratic expansion predicts that at* $\eta = 45.9$, *the two roots lie at* $R_c \pm \sqrt{2(\eta_c - \eta)/|H''(R_c)|} = R_c \pm 0.1808\%$; *solving* (11) *directly gives* 8.2523% *and* 8.6150%, *a half-gap of* 0.1813%. *Figure 5 traces the amplification factor* $1/(1-\theta'(R^\star(\eta)))$ *along the lower branch for* $\eta \in (0, \eta_c)$: *it rises from* 1.03 *at the base case* $\eta = 6.0$ *to a vertical asymptote at* η_c. *Beyond the fold, the absence of roots is global, not only local: for each* η, *the maximum of the concave function* $g(\cdot;\eta)$ *is attained where* $\theta = 1/\eta$ *and equals* $R_0 - c + \big(-\ln(\eta\theta_0) - 1\big)/\eta$, *with* $c = f + \kappa + v + x + m$. *This is decreasing in* η *for* $\eta < 1/\theta_0$, *is negative for* $\eta > 1/\theta_0$ *(where it is below* $R_0 - c = -1.25\%$*), and vanishes at* η_c. *Hence,* $g(R;\eta) < 0$ *for every* R *when* $\eta > \eta_c$: *no minimum lending rate exists at all past the fold.*

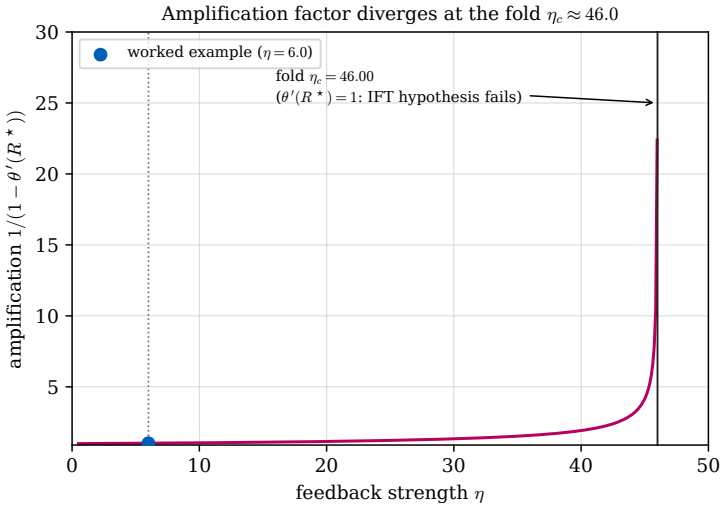


Figure 5. Amplification factor $1/(1-\theta'(R^\star))$ as a function of feedback strength η, diverging at the fold bifurcation $\eta_c \approx 46.0$ identified in Proposition 2 and Example 6. The marked point is the base case ($\eta = 6.0$), well inside the region where $R^\star$ exists.

Remark 2. *Example 6 suggests a computable early-warning quantity for a loan pricing committee: instead of $\theta'(R^\star)$, which has no natural policy threshold, it can compare the estimated pass-through η of coupon into prepayment-option value (estimable in principle from historical refinancing data) with the computed fold value η_c. The ratio η/η_c then measures proximity to a point at which the risk-based pricing equation has no solution, not merely to one at which its forecasts become inaccurate. The exponential feedback here is illustrative; for another specification of θ, the fold location must be recomputed, but Proposition 2 applies unchanged.*

7. Application III: Comparative Statics Under a Nonlinear Risk Limit

This application marks where the machinery of Section 3 is and is not needed. With a quadratic risk *penalty*, the first-order conditions are linear and comparative statics is linear algebra. With a hard risk *limit*, the multiplier of the limit becomes an unknown, and the first-order conditions form a nonlinear system. For the particular problem studied here (one budget constraint and one risk limit), the classical two-fund solution [39] is still available in closed form, and we use it as an exact benchmark. The IFT computation never uses that closed form, so it applies unchanged to KKT systems that have none.

7.1. The Penalty Benchmark, and What a Hard Limit Changes

With a risk penalty $\gamma > 0$, the problem $\max_w \ \mu'w - \frac{\gamma}{2}w'\Sigma w$ subject to $\mathbf{1}'w = 1$ has first-order conditions $\mu - \lambda_1\mathbf{1} - \gamma\Sigma w = 0$, $\mathbf{1}'w = 1$, which are linear in (w, λ_1) for given γ. With $a = \mathbf{1}'\Sigma^{-1}\mathbf{1}$, $b = \mathbf{1}'\Sigma^{-1}\mu$, $c = \mu'\Sigma^{-1}\mu$ and $d = ac - b^2 > 0$, the solution is $w(\gamma) = \Sigma^{-1}(\mu - \lambda_1\mathbf{1})/\gamma$ with $\lambda_1 = (b-\gamma)/a$, and its dependence on γ follows by differentiating an explicit formula.

Bank policy, however, sets a hard limit, usually stated in volatility units: $\sqrt{w'\Sigma w} \leq \sigma_{\max}$. Since Σ is positive definite (its eigenvalues are 1.96×10^{-5}, 1.92×10^{-4} and 1.09×10^{-3}), the limit binds at the optimum whenever $\sigma_{\max}$ exceeds the minimum attainable volatility $a^{-1/2}$, and the problem is $\max_w \mu'w$ subject to $\mathbf{1}'w = 1$ and $\sqrt{w'\Sigma w} = \sigma_{\max}$. Stating the same limit in variance units, $w'\Sigma w = \sigma^2_{\max}$, describes the same feasible set and therefore gives the same optimal weights; only the multiplier is rescaled (the variance multiplier equals $\lambda_2/(2\sigma_{\max})$, where λ_2 is the volatility multiplier). In either form, the multiplier λ_2 is now an unknown, and the stationarity condition $\mu - \lambda_1\mathbf{1} - \lambda_2\Sigma w/\sqrt{w'\Sigma w} = 0$ is

nonlinear in (w, λ_2) jointly. For this problem, the solution is known [39]: the largest margin attainable at volatility σ, and the weights that attain margin m at least variance, are

$$m^\star(\sigma) = \frac{b + \sqrt{d\,(a\sigma^2 - 1)}}{a}, \qquad w(m) = \frac{\Sigma^{-1}\left[(c - bm)\mathbf{1} + (am - b)\mu\right]}{d}, \tag{14}$$

and by the envelope theorem, the multiplier of the volatility limit is $\lambda_2^\star = dm^\star/d\sigma = \sqrt{d}\,\sigma/\sqrt{a\sigma^2 - 1}$.

What the IFT adds is a method that does not depend on (14). A second nonlinear limit (for example, a separate volatility limit on a sub-book) or a nonlinear funding-cost term changes the KKT system and removes the closed form, but the computation below—solve the KKT system by Newton's method, check that its Jacobian is non-singular, and solve one linear system for the sensitivities—goes through unchanged.

7.2. The Map

With $n = 3$ assets, let $z = (w, \lambda_1, \lambda_2) \in \mathbb{R}^5$ collect the weights and the two Lagrange multipliers (budget, risk limit) and define $F : \mathbb{R}^5 \times \mathbb{R} \to \mathbb{R}^5$ by the KKT stationarity and constraint system

$$F(z; \sigma_{\max}) = \begin{pmatrix} \mu - \lambda_1 \mathbf{1} - \lambda_2 \dfrac{\Sigma w}{s(w)} \\ \mathbf{1}'w - 1 \\ s(w) - \sigma_{\max} \end{pmatrix} = 0, \qquad s(w) := \sqrt{w'\Sigma w}. \tag{15}$$

Because Σ is positive definite, $s(w) > 0$ for $w \neq 0$, and every w with $\mathbf{1}'w = 1$ is non-zero, so F is C^∞ near every solution. Its Jacobian with respect to z is

$$J_{\mathrm{KKT}} = \begin{pmatrix} -\lambda_2\left(\dfrac{\Sigma}{s} - \dfrac{\Sigma w w'\Sigma}{s^3}\right) & -\mathbf{1} & -\dfrac{\Sigma w}{s} \\ \mathbf{1}' & 0 & 0 \\ \dfrac{w'\Sigma}{s} & 0 & 0 \end{pmatrix} \in \mathbb{R}^{5\times 5}. \tag{16}$$

Whenever J_{KKT} is non-singular at a solution $z^\star(\sigma_{\max})$, Theorem 2 gives $z^\star$ as a locally unique C^∞ function of the policy parameter $\sigma_{\max}$, with (since $\partial F/\partial\sigma_{\max} = -e_5$)

$$\frac{dz^\star}{d\sigma_{\max}} = -J_{\mathrm{KKT}}^{-1}\frac{\partial F}{\partial \sigma_{\max}} = J_{\mathrm{KKT}}^{-1} e_5. \tag{17}$$

7.3. Numerical Verification

Take three assets with (excess) margins $\mu = (0.45, 1.35, 2.40)\%$, volatilities $(0.45, 1.60, 3.20)\%$, and correlations $\rho_{12} = \rho_{13} = 0.15$, $\rho_{23} = 0.45$ (a low-risk floating instrument, a spread-duration holding, and a credit asset). Newton's method for (15) at $\sigma_{\max} = 1.20\%$, with the analytic Jacobian (16), converges to

$$w^\star = (0.3681,\ 0.4126,\ 0.2192), \qquad \lambda_1^\star = 0.003456, \qquad \lambda_2^\star = 0.752740,$$

with residual $\|F\| = 1.7 \times 10^{-18}$ and $\det J_{\mathrm{KKT}} = -1.12 \times 10^{-5} \neq 0$, so Theorem 2 applies. (The determinant is small because the entries of Σ are of order 10^{-4}; non-singularity, not the size of the determinant, is what the theorem requires.)

Example 7 (comparative statics verified)**.** *The IFT sensitivity* (17) *evaluates to*

$$\frac{d(w_1, w_2, w_3, \lambda_1, \lambda_2)}{d\sigma_{\max}} = (-57.6202,\ 35.3194,\ 22.3008,\ 0.11993,\ -9.99410).$$

Two independent checks: (i) the closed form (14) *gives the same* $w^\star$ *and* $\lambda_2^\star = dm^\star/d\sigma$*, and its derivative* $dw/d\sigma = w'(m^\star)\, dm^\star/d\sigma = \Sigma^{-1}(a\mu - b\mathbf{1})/d \cdot dm^\star/d\sigma$ *agrees with the first three components above to* 7×10^{-15}*, while* $d\lambda_2^\star/d\sigma = d^2m^\star/d\sigma^2 = -\sqrt{d}/(a\sigma^2 - 1)^{3/2} = -9.99410$*; (ii) central differences of the re-solved system at* $\sigma_{\max} \pm h$*,* $h = 10^{-7}$*, agree with* (17) *to* 2.3×10^{-9} *in every component. A* 10 *basis-point increase in the limit moves about 5.8 percentage points of weight out of the floating instrument, into the other two.*

Figure 6 plots the frontier obtained by re-solving (15) at sixty values of $\sigma_{\max}$, together with the closed form (14); Figure 7 shows the multiplier $\lambda_2^\star$ as a function of the limit.

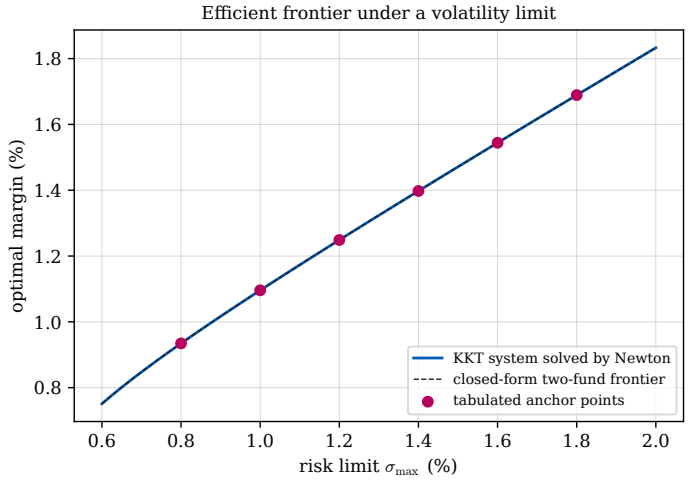


Figure 6. Efficient frontier under the volatility limit (15): optimal margin against the policy limit $\sigma_{\max}$. Solid: Newton solutions of the KKT system; dashed: the closed form $m^\star(\sigma)$ of (14); points: the rows of Table 6. The two curves coincide to plotting accuracy.

Table 6 repeats the computation at six policy limits between 0.80% *and* 1.80%. As the limit is relaxed, weight moves out of the floating instrument into the two riskier assets, the optimal margin rises, and the multiplier $\lambda \star 2$ falls; at every limit the margin and the multiplier agree with the closed forms of (14) to 10^{-15}.

Table 6. Optimal allocation and multiplier of the volatility limit across policy limits, from Newton's method for (15). The margin and $\lambda_2^\star$ columns agree with the closed forms $m^\star(\sigma_{\max})$ and $dm^\star/d\sigma_{\max}$ of (14) to 10^{-15}.

$\sigma_{\max}$	w_1	w_2	w_3	Margin	$\lambda_2^\star$ (Shadow Price)
0.80%	0.6087	0.2652	0.1261	0.9346%	0.84114
1.00%	0.4853	0.3408	0.1739	1.0958%	0.78060
1.20%	0.3681	0.4126	0.2192	1.2489%	0.75274
1.40%	0.2542	0.4825	0.2633	1.3977%	0.73732
1.60%	0.1421	0.5512	0.3067	1.5442%	0.72780
1.80%	0.0312	0.6192	0.3496	1.6891%	0.72149

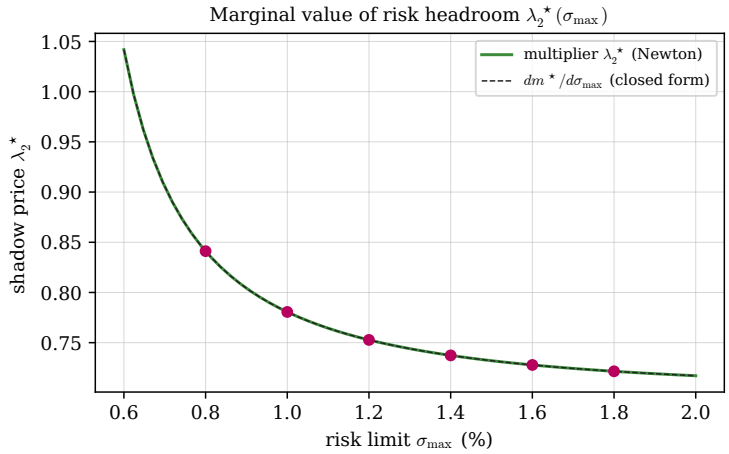


Figure 7. Multiplier $\lambda_2^\star(\sigma_{\max})$ of the volatility limit: the additional margin per unit of additional volatility headroom. Solid: Newton solutions; dashed: $dm^\star/d\sigma$ from (14). It decreases because the frontier is concave, $d^2m^\star/d\sigma^2 = -\sqrt{d}/(a\sigma^2 - 1)^{3/2} < 0$; the multiplier of the equivalent variance limit, $\lambda_2^\star/(2\sigma_{\max})$, decreases as well.

Remark 3. *The multiplier $\lambda_2^\star$ is the marginal margin per unit of volatility headroom. A natural consistency check for an institution is to compare it with the hurdle return per unit of risk used in its loan pricing (Section 6): if pricing and allocation used one cost of risk, the two would coincide. Industry white papers describe this kind of requirement informally as internal consistency [16]; (17) gives one precise counterpart, since it makes $\lambda_2^\star$ and its slope available from the Jacobian already formed by Newton's method, without re-optimizing at each candidate hurdle.*

8. Application IV: Reverse Optimization via the Inverse Function and Rank Theorems

The first three applications ask "does a solution exist, and how does it move"; this last one asks the question a board actually poses, which is the same question read backward: "here is where we want to be—what must we do." That reversal is precisely the InvFT when the number of levers matches the number of targets, and precisely the Rank Theorem the moment a board asks for more than the balance sheet has degrees of freedom to give it. In that case, part of the question about priorities becomes a dimension count that can be settled before the meeting.

Reverse optimization inverts the usual planning problem: a governing body specifies a target vector $u^\star$ of state coordinates (margin, volatility, duration) and asks for balance-sheet weights w achieving it, rather than proposing w and evaluating $\Psi(w)$. We treat the two-target case (InvFT, part (a)) and the three-target, two-parameter case (Rank Theorem, part (b)) with the same three-asset universe as Section 7, parametrized after the budget constraint by the two free weights (w_1, w_2), with $w_3 = 1 - w_1 - w_2$; that is, $w = e_3 + E\,(w_1, w_2)'$ with the constant selector $E = \begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \end{pmatrix}'$. We write $\|w\|_\Sigma := \sqrt{w'\Sigma w}$ for the portfolio volatility; it is positive and C^∞ on $\{\mathbf{1}'w = 1\}$ because Σ is positive definite and such w are non-zero.

8.1. Part (a): Local Invertibility of a Two-Target Map

Let $\Psi_2(w_1, w_2) = (\mu'w,\ \|w\|_\Sigma)$: margin and volatility as functions of the two free weights. At the base allocation $w_0 = (0.10, 0.85, 0.05)$,

$$\Psi_2(w_0) = (1.3125\%,\ 1.4469\%),$$

$$D\Psi_2(w_0) = \begin{pmatrix} -0.0195 & -0.0105 \\ -0.01637 & -0.00131 \end{pmatrix}, \quad \det D\Psi_2(w_0) = -1.463 \times 10^{-4} \neq 0.$$

By Theorem 3, Ψ_2 is a local C^∞-diffeomorphism near w_0: for any target $(m^\star, v^\star)$ sufficiently close to $\Psi_2(w_0)$, a unique nearby w with $\Psi_2(w) = (m^\star, v^\star)$ exists. It can be computed by Newton's method, whose iterates converge to it quadratically from any sufficiently close start [19]; Theorem 3 is what guarantees that the limit is the *unique* nearby preimage.

Example 8 (local inversion)**.** *Target* $(m^\star, v^\star) = \Psi_2(w_0) + (0.0150\%, -0.0120\%)$. *Newton iterates on* $w = (w_1, w_2)$ *are listed in Example 8: Newton iterates for the local inversion of* Ψ_2 *: allocation and distance of* $\Psi_2(w_k)$ *from the target.*

k	(w_1, w_2, w_3)	$\|\Psi_2(w_k) - target\|$
0	(0.1000, 0.8500, 0.0500)	1.921×10^{-4}
1	(0.1100, 0.8172, 0.0728)	1.581×10^{-5}
2	(0.1112, 0.8149, 0.0739)	4.120×10^{-8}
3	(0.1112, 0.8149, 0.0739)	0

The residual ratios $\|r_{k+1}\| / \|r_k\|^2$ *are* 428 *and* 165*, bounded as quadratic convergence requires (green triangles in Figure 3). The allocation reaching the target exactly is* $w^\star = (11.12\%, 81.49\%, 7.39\%)$.

8.2. Part (b): Exact Reachability with Three Targets and Two Free Weights

Add a third, affine target—portfolio duration $D'w$ with $D = (0.30, 4.20, 3.10)$—giving $\Psi_3 = (\Psi_2, D'w) : \mathbb{R}^2 \to \mathbb{R}^3$. Because margin and duration are *both* affine in (w_1, w_2), their joint map

$$\begin{pmatrix} \text{margin} \\ \text{duration} \end{pmatrix} = \begin{pmatrix} \mu_3 \\ D_3 \end{pmatrix} + \underbrace{\begin{pmatrix} \mu_1 - \mu_3 & \mu_2 - \mu_3 \\ D_1 - D_3 & D_2 - D_3 \end{pmatrix}}_{=:A} \begin{pmatrix} w_1 \\ w_2 \end{pmatrix} \tag{18}$$

is an affine bijection of $\mathbb{R}^2$, since

$$A = \begin{pmatrix} -0.0195 & -0.0105 \\ -2.8 & 1.1 \end{pmatrix}, \qquad \det A = -0.05085 \neq 0. \tag{19}$$

(This is linear algebra; no inverse function theorem is needed for an affine map.) Every margin/duration pair is therefore attained by exactly one (w_1, w_2), and that pair determines w and hence the third coordinate, volatility, which cannot be assigned independently.

The rank structure is the same at every point, not only at w_0. The margin and duration rows of $D\Psi_3(w)$ are the constant rows of A, so $\operatorname{rank} D\Psi_3(w) = 2$ for every $w \in \mathbb{R}^2$, and Ψ_3 is C^∞ on all of $\mathbb{R}^2$. Corollary 1 (with $n = 2 < m = 3$) therefore applies at every point: near each allocation, the image of Ψ_3 is a smooth two-dimensional surface. In this example, the conclusion is also global, because (m, D) is an affine bijection of (w_1, w_2): with $v_{\text{impl}}(m, D) := \|w(m, D)\|_\Sigma$, where $w(m, D)$ is the unique allocation with margin m and duration D,

$$\Psi_3(\mathbb{R}^2) = \{(m,\ v_{\text{impl}}(m, D),\ D) : (m, D) \in \mathbb{R}^2\}, \tag{20}$$

the graph of a C^∞ function over the whole margin–duration plane. Hence, a three-target request $(m^\star, v^\star, D^\star)$ is attained by some w if and only if $v^\star = v_{\text{impl}}(m^\star, D^\star)$; otherwise it is attained by no w at all. Figure 8 shows this surface.

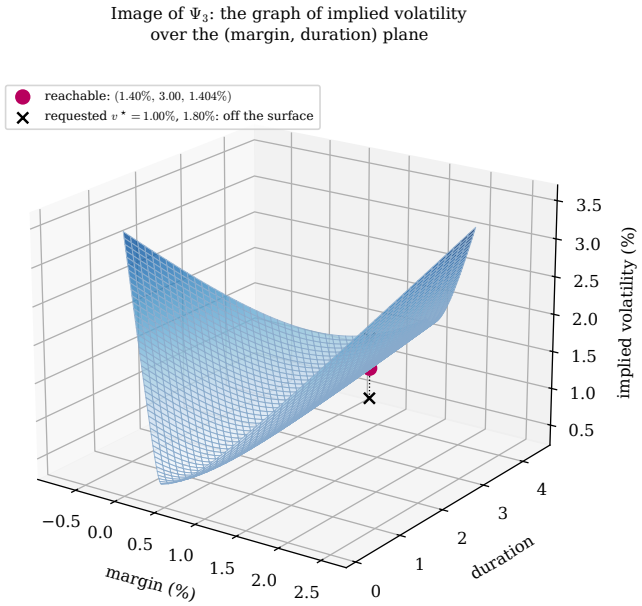


Figure 8. The image of Ψ_3 for (w_1, w_2) in the square $-0.05 \le w_1, w_2 \le 1.05$ (with $w_3 = 1 - w_1 - w_2$): implied volatility as a function of margin and duration. By (20), no allocation places the balance sheet off this surface. Magenta: the reachable point $(m^\star, D^\star, v_{\text{impl}}) = (1.40\%, 3.00, 1.40385\%)$ of Example 9. Black crosses: the requests $v^\star = 1.00\%$ and $v^\star = 1.80\%$ at the same margin and duration, which lie 0.4039 percentage points below and 0.3961 points above the surface and are attained by no allocation.

Example 9 (exact reachability and closed-form infeasibility)**.** *Target margin* $m^\star = 1.40\%$*, target duration* $D^\star = 3.00$*. Solving* (18) *exactly gives* $w^\star = (23.70\%, 51.23\%, 25.07\%)$*, with implied volatility* $v_{impl} = \sqrt{w^{\star\prime}\Sigma w^\star} = 1.40385\%$*. A request for volatility* 1.40385% *in addition is met exactly. Requests for volatility* 1.40%*,* 1.00% *or* 1.80% *leave residuals* $v^\star - v_{\text{impl}} = -0.0039, -0.4039$ *and* $+0.3961$ *percentage points respectively, which no choice of* w *can close, since* w *is already determined by the first two targets. This is an exact statement, not a least-squares approximation.*

The sensitivity of the pinned-down implied volatility to the two exact targets follows by the chain rule through (18):

$$\frac{\partial v_{\text{impl}}}{\partial (m^\star, D^\star)} = \frac{(\Sigma w^\star)'}{\|w^\star\|_\Sigma} E\, A^{-1}, \tag{21}$$

with E the selector defined above. It evaluates to $(1.300228,\ 0.000209)$: holding target duration fixed, a 1 percentage point increase in the target margin raises the volatility that comes with it by 1.30 percentage points, while a change in target duration alone has almost no effect. Central differences of v_{impl} at $m^\star \pm h$ and $D^\star \pm h$, $h = 10^{-6}$, reproduce these values to 1.0×10^{-8}. Table 7 shows the same pattern at other target pairs: the implied volatility and its sensitivity vary smoothly with the targets, as (20) requires.

The fourth row already requires a small short position in asset 1 to hit the stated targets exactly, and the fifth requires a larger one; a board restricting $w \ge 0$ would find those pairs infeasible for a second reason, independent of the one in Example 9: a sign constraint rather than a rank obstruction. The weights in this section are sign-unrestricted, and (20) describes reachability without further admissibility constraints; it does not by itself guarantee that the reaching allocation is otherwise admissible.

Table 7. Exact reverse optimization at five margin/duration target pairs: uniquely determined weights, implied (pinned-down) volatility, and its closed-form sensitivity (21) to the two exact targets.

$m^\star$	$D^\star$	$w^\star = (w_1, w_2, w_3)$	v_{impl}	$\partial v_{\text{impl}}/\partial(m^\star, D^\star)$
1.00%	2.00	(53.00%, 34.91%, 12.09%)	0.8836%	(1.0643, 0.000615)
1.20%	2.50	(38.35%, 43.07%, 18.58%)	1.1373%	(1.2159, 0.000368)
1.40%	3.00	(23.70%, 51.23%, 25.07%)	1.4039%	(1.3002, 0.000209)
1.60%	3.50	(9.05%, 59.39%, 31.56%)	1.6771%	(1.3522, 0.000101)
1.80%	4.00	(−5.60%, 67.55%, 38.05%)	1.9542%	(1.3869, 0.000022)

Remark 4 (Governance reading)**.** *Corollary 1 and* (20) *make precise a fact balance-sheet managers meet informally: a balance sheet with two degrees of freedom cannot meet three independent targets except by coincidence. The useful response is not to search for a w that does not exist but to report the residual* $v^\star - v_{\text{impl}}(m^\star, D^\star)$ *and the sensitivity* (21)*, which state what meeting the first two targets implies for the third, and to let the board choose which two coordinates bind. Industry white papers on reverse optimization [17] recommend reporting such a shortfall when targets conflict; the formulas above are one precise counterpart of that informal recommendation.*

9. Discussion: A Rank-Based Classification of Balance-Sheet Well-Posedness

Sections 5–8 apply the results of Section 3 to four problems that look unrelated: a lattice calibration, a pricing equation with feedback, a constrained portfolio, and a set of governance targets. What they share is a rank condition on the derivative of a map from actions to state coordinates. The following corollary states that condition once. It adds no analysis to Section 3; it is the Rank Theorem specialized to maximal rank, organized by the sign of $m - n$.

Corollary 2 (Maximal-rank cases of the state map)**.** *Let* $\Omega \subset \mathbb{R}^n$ *be open,* $\Psi \in C^k(\Omega; \mathbb{R}^m)$ *with* $k \geq 1$ *the map from balance-sheet actions to state coordinates,* $w_0 \in \Omega$ *and* $u_0 = \Psi(w_0)$*. Suppose* $\operatorname{rank} D\Psi(w_0) = \min(n, m)$ *at this single point. Then,* $\operatorname{rank} D\Psi = \min(n, m)$ *on a neighborhood of* w_0 *(Lemma 6), and there are open sets* $U \ni w_0$ *and* $V \ni u_0$ *such that:*

(a) *if* $m = n$, $\Psi : U \to V$ *is a* C^k*-diffeomorphism; for each* $u \in V$ *the equation* $\Psi(w) = u$ *has exactly one solution in* U*, and* $D(\Psi^{-1})(u_0) = D\Psi(w_0)^{-1}$ *(Theorem 3);*

(b) *if* $m < n$, $\Psi(U)$ *is an open neighborhood of* u_0*, and for each* $u \in \Psi(U)$*, the solution set* $\{w \in U : \Psi(w) = u\}$ *is an* $(n - m)$*-dimensional* C^k*-submanifold (Theorem 4 with* $r = m$*: in the coordinates* $z = \phi(w)$ *it is the slice* $\{z \in \widetilde{U} : (z_1, \ldots, z_m) = \psi(u)\}$*);*

(c) *if* $m > n$, Ψ *is injective on* U *and* $\Psi(U)$ *is an* n*-dimensional* C^k*-submanifold of* $\mathbb{R}^m$*, locally the graph of a* C^k *function of* n *of the coordinates; no* $w \in U$ *is mapped to a point of* $V \setminus \Psi(U)$ *(Corollary 1).*

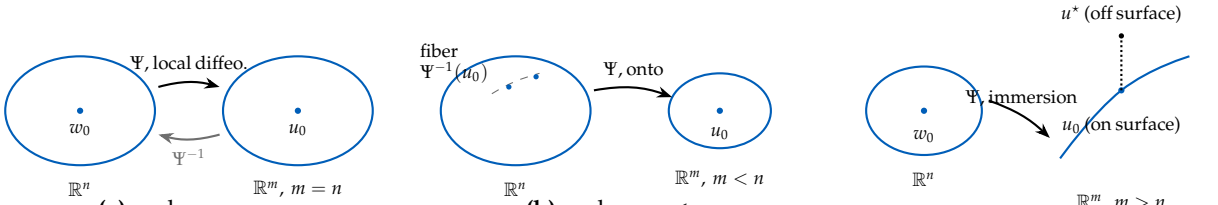
All three statements are local. A statement about every $w \in \Omega$*, such as* (20)*, needs additional global information.*

Proof. Lemma 6 gives constant rank $\min(n, m)$ near w_0, so Theorem 4 applies with $r = \min(n, m)$. Case (a) is Theorem 3, and case (c) is Corollary 1. In case (b), (6) with $r = m$ reads, $\psi \circ \Psi \circ \phi^{-1}(z) = (z_1, \ldots, z_m)$ on $\widetilde{U}$. Coordinate projection maps open sets to open sets, so $\Psi(U) = \psi^{-1}(\pi_m(\widetilde{U}))$ is open, and the solution set of $\Psi(w) = u$ in U is the image under the diffeomorphism ϕ^{-1} of the slice stated in (b), an open subset of an $(n - m)$-dimensional affine subspace. □

Remark 5 (Where the hypothesis fails, and an informal reading)**.** *Two of the failures met in the applications are failures of the maximal-rank hypothesis in the square case (a): near-singularity of*

the calibration matrix $\mathbf{M}_t$ *in Section 5, and the fold of Section 6, where* $g_R = 1 - \theta'$ *drops from rank* 1 *to rank* 0*. The infeasible three-target request of Section 8.2 is of a different kind: there,* $D\Psi_3$ *has maximal rank* $n = 2 < m = 3$ *everywhere, and unreachability of off-surface targets is exactly what case (c) predicts, not a loss of rank. Industry white papers on balance-sheet management [14,17] describe an informal "reconciliation" requirement: that bottom-up and top-down descriptions of the balance sheet agree. Our reading, which is an interpretation and not a result taken from those papers, is that when the top-down description is a target* $u \in \mathbb{R}^m$ *with* $m < n$*, case (b) is one precise counterpart of that requirement: full row rank of* $D\Psi(w_0)$ *guarantees that every nearby target is met by an* $(n - m)$*-parameter family of allocations.*

Figure 9 depicts the three cases schematically.



(a) rank $= n = m$
bijective near w_0

(b) rank $= m < n$
$(n-m)$-param. family reaches u_0

(c) rank $= n < m$
$u^\star$ off the image: unreachable

Figure 9. The three maximal-rank cases of Corollary 2. (**a**) $m = n$: the InvFT gives a local bijection (Section 8.1). (**b**) $m < n$: a submersion, so every nearby target is reached by an $(n - m)$-dimensional family of allocations (our reading of the informal reconciliation requirement, Remark 5). (**c**) $m > n$: an immersion (Corollary 1); the image is an n-dimensional surface, and a target off it, such as $u^\star$, is not attained near w_0; in Section 8.2 it is not attained at all (Figure 8).

Corollary 2 contains no analysis beyond Section 3. We state it because it reduces the well-posedness questions of the applications to one check—the rank of $D\Psi$ at the computed point—that tells a practitioner which of the three local pictures applies, and because Sections 5–8 show what that check costs and what it delivers in each case.

Scope and Limitations

Four caveats bound the results above. First, all four theorems are *local*: the sensitivity formulas and the reachability statements hold in a neighborhood of the reference point whose radius the theorem statements do not give. The proofs do: the explicit δ_1, ϵ_0 in the proof of Theorem 1, and the explicit r, c of Lemma 3 give a radius of validity, typically a conservative one. Large policy moves (a large change in $\sigma_{\max}$ in Section 7, a large revision of a market quote in Section 5) may leave the region where the linear sensitivity is a good approximation, and the nonlinear system must then be re-solved rather than the derivative extrapolated; Table 6 re-solves at each row for this reason. Second, the results presume that Ψ is genuinely C^1. Behavioral assumptions (deposit decay, prepayment functions) calibrated from historical data carry estimation error that the implicit-function machinery does not model, and the exponential feedback of Section 6 is an illustrative specification of a more complex and imperfectly known feedback function. Third, Corollary 2 assumes maximal rank at w_0. Where the rank is lower but locally constant, Theorem 4 still applies; where it is not locally constant, the local picture needs a case analysis beyond this paper, of which the fold of Proposition 2 is the simplest instance. Fourth, all data in Sections 5–8 are synthetic or hand-chosen: the lattice slice and its skewed test market, the pricing components and feedback parameter, the three-asset universe, and the board targets. The applications are a proof of concept for the method—they show that each hypothesis can be checked, and each sensitivity computed and independently verified, at realistic magnitudes—and are not empirical findings about particular markets or institutions.

10. Conclusions

This paper has made two linked contributions. Analytically, we have given a fully self-contained, elementary proof of the Implicit Function Theorem, the Inverse Function Theorem, and the Rank Theorem on $\mathbb{R}^n$ (Section 3), following Dini's inductive strategy: the scalar case from the IVT and MVT alone, the vector case by induction via an explicit Schur-complement reduction, the InvFT via a short bi-Lipschitz injectivity estimate, and the Rank Theorem via a block-rank argument and a connected-domain zero-derivative lemma—with no compactness argument in $\mathbb{R}^n$, no contraction mapping, and no fixed-point theorem anywhere in the development; the one property of $\mathbb{R}$ used beyond algebra is its completeness, through the IVT. Applied, we have shown, by explicit construction and numerical verification in four settings, that standard calibration, pricing, allocation and reverse-optimization problems of balance-sheet management are direct instances of this circle of results. In each case, writing the map down and checking the rank of its Jacobian at the computed solution gave a sensitivity or reachability formula, which we verified against central differences of re-solved systems and, in Section 7, against the classical two-fund closed form. Corollary 2 collects the cases by the sign of $m - n$, and Remark 5 offers our reading of an informal "reconciliation" requirement from the industry literature as an instance of its maximal-rank hypothesis. The elementary proofs are constructive, so their constants (δ_1, ϵ_0 in the proof of Theorem 1, r, c in Lemma 3) give explicit, if conservative, radii of validity; contraction-mapping proofs also yield explicit constants, so the advantage claimed here is transparency of the argument rather than quantities unavailable by other routes. The applications are proofs of concept on synthetic data.

Future work could extend the calibration analysis of Section 5 to the full L-parameter local-volatility slice with a genuine strike ladder and market quotes; turn the explicit constants of Section 3 into usable a priori radii of validity for the sensitivity formulas; study how far the elementary strategy extends towards infinite-dimensional (Banach-space) settings, where compactness or contraction arguments appear to become unavoidable; and analyze balance-sheet maps at points where the rank of $D\Psi$ is neither maximal nor locally constant, beyond the fold of Section 6.

Author Contributions: Conceptualization, T.L.; methodology, T.L. and X.-M.Y.; formal analysis, T.L. and X.-M.Y.; software, T.L.; validation, T.L. and X.-M.Y.; writing—original draft preparation, T.L.; writing—review and editing, T.L. and X.-M.Y. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Acknowledgments: The authors thank the two anonymous reviewers for their careful reading of the manuscript and for comments that substantially improved the presentation of Section 3 and of the applications.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Spivak, M. *Calculus on Manifolds*; Perseus Books: Cambridge, MA, USA, 1965.
2. Munkres, J.R. *Analysis on Manifolds*; Addison-Wesley: Redwood City, CA, USA, 1991.
3. Guillemin, V.; Pollack, A. *Differential Topology*; Prentice-Hall: Englewood Cliffs, NJ, USA, 1974.

4. Gilbarg, D.; Trudinger, N.S. *Elliptic Partial Differential Equations of Second Order*, 3rd ed.; Springer: Berlin, Germany, 2001.
5. Samuelson, P.A. *Foundations of Economic Analysis*; Harvard University Press: Cambridge, MA, USA, 1947.
6. Silberberg, E.; Suen, W. *The Structure of Economics: A Mathematical Analysis*, 3rd ed.; McGraw-Hill: Boston, MA, USA, 2001.
7. Milgrom, P.; Segal, I. Envelope theorems for arbitrary choice sets. *Econometrica* **2002**, *70*, 583–601. [CrossRef]
8. Rudin, W. *Principles of Mathematical Analysis*, 3rd ed.; McGraw-Hill: New York, NY, USA, 1976.
9. Dini, U. *Lezione di Analisi Infinitesimale*; Nistri: Pisa, Italy, 1907; Volume 1, pp. 197–241.
10. de Oliveira, O.R.B. The Implicit and the Inverse Function Theorems: Easy Proofs. *Real Anal. Exch.* **2013**, *39*, 207–218. [CrossRef]
11. Ho, T.S.Y. pOAS: A New Profitability Measure for Loan Acquisition & Loan Portfolio Management. In *THC White Paper*; Series One #1; Thomas Ho Company: New York, NY, USA, 2019.
12. Ho, T.S.Y. Manage Option Risk: Local Volatilities Model. In *THC White Paper*; Series One #1; Thomas Ho Company: New York, NY, USA, 2020.
13. Ho, T.S.Y. Market Interest Rate Forecast (MIRF): An Arbitrage-Free Interest Rate Modeling Approach. In *THC White Paper*; Series One #3; Thomas Ho Company: New York, NY, USA, 2020.
14. Ho, T.S.Y. Banking Paradigm Shift: Challenges and Solutions. In *THC White Paper*; Series One #4; Thomas Ho Company: New York, NY, USA, 2020.
15. Ho, T.S.Y. Balance Sheet Performance Attribution. In *THC White Paper*; Thomas Ho Company: New York, NY, USA, 2024.
16. Ho, T.S.Y. The IncomeRisk™ Model: The Blueprint for the Next Generation of Operating Banking Systems (v1). In *THC White Paper*; Thomas Ho Company: New York, NY, USA, 2026.
17. Ho, T.S.Y. Financial Representation: Foundations of Balance Sheet Intelligence. In *THC Monograph*; Thomas Ho Company: New York, NY, USA, 2026.
18. Ho, T.S.Y.; Lee, S.B. Term structure movements and pricing interest rate contingent claims. *J. Finance* **1986**, *41*, 1011–1029. [CrossRef]
19. Ortega, J.M.; Rheinboldt, W.C. *Iterative Solution of Nonlinear Equations in Several Variables*; Academic Press: New York, NY, USA, 1970.
20. Krantz, S.G.; Parks, H.R. *The Implicit Function Theorem: History, Theory, and Applications*; Birkhäuser: Boston, MA, USA, 2002.
21. Bliss, G.A. A new proof of the existence theorem for implicit functions. *Bull. Am. Math. Soc.* **1912**, *18*, 175–179. [CrossRef]
22. Apostol, T.M. *Mathematical Analysis: A Modern Approach to Advanced Calculus*; Addison-Wesley: Reading, MA, USA, 1957.
23. Dontchev, A.L.; Rockafellar, R.T. *Implicit Functions and Solution Mappings*; Springer: New York, NY, USA, 2009.
24. Black, F.; Derman, E.; Toy, W. A one-factor model of interest rates and its application to treasury bond options. *Financ. Anal. J.* **1990**, *46*, 33–39. [CrossRef]
25. Heath, D.; Jarrow, R.; Morton, A. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica* **1992**, *60*, 77–105. [CrossRef]
26. Dupire, B. Pricing with a smile. *Risk* **1994**, *7*, 18–20.
27. Derman, E.; Kani, I. Riding on a smile. *Risk* **1994**, *7*, 32–39.
28. Rubinstein, M. Implied binomial trees. *J. Finance* **1994**, *49*, 771–818. [CrossRef]
29. Andreasen, J.; Huge, B. Volatility interpolation. *Risk* **2011**, *24*, 76–79.
30. Basel Committee on Banking Supervision. *Standards: Interest Rate Risk in the Banking Book*; Bank for International Settlements: Basel, Switzerland, 2016.
31. Financial Accounting Standards Board. *Accounting Standards Update No. 2016-13: Financial Instruments—Credit Losses (Topic 326)*; FASB: Norwalk, CT, USA, 2016.
32. English, W.B.; Van den Heuvel, S.J.; Zakrajšek, E. Interest rate risk and bank equity valuations. *J. Monet. Econ.* **2018**, *98*, 80–97. [CrossRef]
33. Barr, M.S. *Review of the Federal Reserve's Supervision and Regulation of Silicon Valley Bank*; Board of Governors of the Federal Reserve System: Washington, DC, USA, 2023.
34. Sharpe, W.F. Imputing expected security returns from portfolio composition. *J. Financ. Quant. Anal.* **1974**, *9*, 463–472. [CrossRef]
35. Black, F.; Litterman, R. Global portfolio optimization. *Financ. Anal. J.* **1992**, *48*, 28–43. [CrossRef]
36. Ahuja, R.K.; Orlin, J.B. Inverse optimization. *Oper. Res.* **2001**, *49*, 771–783. [CrossRef]
37. Iyengar, G.; Kang, W. Inverse conic programming with applications. *Oper. Res. Lett.* **2005**, *33*, 319–330. [CrossRef]
38. Merton, R.C. On the pricing of corporate debt: The risk structure of interest rates. *J. Finance* **1974**, *29*, 449–470. [CrossRef]
39. Merton, R.C. An analytic derivation of the efficient portfolio frontier. *J. Financ. Quant. Anal.* **1972**, *7*, 1851–1872. [CrossRef]
40. Nocedal, J.; Wright, S.J. *Numerical Optimization*, 2nd ed.; Springer: New York, NY, USA, 2006.
41. Kuznetsov, Y.A. *Elements of Applied Bifurcation Theory*, 3rd ed.; Springer: New York, NY, USA, 2004.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.