A Wallace Theorem for Nominally Scalable Economies*

Hiroki Fukai†

August 2, 2026

Abstract

This paper proves a Wallace theorem for pure-currency economies with nominal-scale invariance under anticipated permanent money creation. The Implementation Theorem shows that a proportional, welfare-neutral transfer can be added to any locally favorable policy direction to obtain a nonnegative, increasing policy that is truthfully implemented with private balances. The Genericity Theorem shows that, within a finite-dimensional family of nonproportional policy directions, failure of first-order welfare improvement is confined to a lower-dimensional exceptional set under a full-rank condition. A three-type Lagos–Wright economy verifies all assumptions from primitives and identifies an open parameter region with a strictly welfare-improving permanent transfer policy.

Keywords: Wallace theorem; pure-currency economy; money-financed transfers

JEL Classification: E40, E52, D50

*We would like to thank participants in various seminars and conferences for their helpful comments. Fukai acknowledges financial support from JSPS KAKENHI (26K00358 and 26H01961).

†Seinan Gakuin University. Address: 6-2-92 Nishijin, Sawara-ku, Fukuoka-shi, Fukuoka 814-8511, Japan. E-mail: fukai.hiroki@gmail.com

1 Introduction

Wallace (2014) conjectures that, in almost every pure-currency economy with a non-degenerate distribution of money holdings, appropriately designed money creation can raise ex ante welfare. The conjecture is distributional rather than merely inflationary: policy changes who receives liquidity, private portfolio choices determine the endogenous distribution of balances, and a permanent intervention affects trade and welfare over the entire equilibrium path. A satisfactory theorem must therefore allow households to adjust their money holdings in anticipation of policy and must account for the resulting changes in future prices, values, allocations, and balance distributions.

Existing work establishes important model-specific welfare-improvement results. Deviatov and Wallace (2001) construct an environment in which lump-sum money creation is beneficial, and Deviatov (2006) reports numerical examples of beneficial money creation in versions of the Trejos–Wright and Shi random-matching models under an alternative implementability criterion. In a divisible-money search economy with endogenous distributions of money and prices, Molico (2006) shows that proportional transfers are neutral and that, at low inflation, faster monetary expansion through lump-sum transfers tends to reduce dispersion and improve welfare. Rocheteau, Weill, and Wong (2018) derive an endogenous nondegenerate distribution of money holdings and show that lump-sum money creation can be welfare enhancing when labor productivity is low, whereas regressive transfers can approach the first best when productivity is high and agents are sufficiently patient. Amendola, Ferraris, and Mattesini (2021) provide a pure-currency example in which money-financed transfers improve ex ante welfare because of liquidity externalities. Nozawa and Yang (2021) numerically compute optimal transfer schemes in an extended Shi–Trejos–Wright model and find that money creation is optimal in more examples when the upper bound on individual money holdings is relaxed. Closely related implementation work by Bajaj et al. (2017) shows that an incentive-compatible transfer

scheme financed by money creation can decentralize the constrained-efficient allocation in a Lagos–Wright pure-currency economy with divisible money. These results establish important mechanisms and constructive examples, but not a general local theorem for a policy announced before date-0 portfolio choice and maintained permanently. A related paper, Fukai (2026), proves an implementation and finite-rank genericity result for a one-time, unanticipated transfer. In that setting the inherited balance distribution is fixed when policy is announced. In the present paper, the policy is anticipated and permanent, so households adjust their initial portfolios and the intervention changes the full future equilibrium sequence.

The paper develops the Implementation Theorem and the Genericity Theorem and then verifies both in a fully specified monetary economy. First, the Implementation Theorem uses exact neutrality of proportional transfers. In an economy with nominal-scale invariance, a proportional increase in all dated nominal balances, prices, payments, and ceilings changes only the unit of account, leaving the normalized equilibrium path and welfare unchanged. Consequently, a sufficiently large proportional component can be added to any policy direction with a positive welfare derivative. This modification makes the policy nonnegative and increasing without changing its first-order welfare effect. Because lower reported balances then never receive larger transfers, truthful reporting is compatible with private money holdings.

Second, the Genericity Theorem formalizes Wallace's phrase "almost every." Consider a finite-dimensional, economically specified family of nonproportional policy directions. Their welfare derivatives define a finite-dimensional map from primitive parameters into Euclidean space. If this map satisfies a full-rank condition whenever all selected directions are neutral, simultaneous neutrality occurs only on a lower-dimensional exceptional set. Outside that set, at least one direction has a nonzero welfare derivative; choosing its favorable orientation and applying the Implementation Theorem yields a feasible welfare-improving permanent policy. The result is deliberately finite dimensional and makes no unsupported claim about transversality

in the full infinite-dimensional policy space.

A three-type Lagos–Wright economy provides a complete primitive application. The model retains the centralized-market/decentralized-market structure of Lagos and Wright (2005), uses the proportional bargaining solution of Kalai (1977), and generates three endogenous stationary money positions. The application derives the welfare derivative of a permanent balance-contingent policy while allowing households to adjust their portfolios. It then constructs the smooth policy direction in (16) and proves that its welfare derivative is strictly positive on a nonempty open set of primitive parameters. Thus every assumption used by the Implementation Theorem and the Genericity Theorem is verified inside the model rather than imposed in reduced form. Locally, the application gives a stronger conclusion than the Genericity Theorem: the selected policy direction remains favorable throughout the maintained neighborhood, so the local exceptional set is empty.

The paper builds on the search-theoretic foundations of monetary exchange in Kiyotaki and Wright (1989, 1993), Trejos and Wright (1995) and Rocheteau and Wright (2005), and Lagos and Wright (2005). It is also related to research that endogenizes nondegenerate money distributions, including Zhou (1999) and Berentsen, Camera, and Waller (2005); Molico (2006) belongs to this distributional literature as well, but is cited above because it additionally establishes a model-specific welfare gain from lump-sum monetary expansion at low inflation. Exact neutrality of proportional transfers is related to nominal irrelevance results such as Wallace (1981), but it serves a different purpose here: it converts a locally favorable policy direction into a feasible and incentive-compatible permanent transfer rule without changing the first-order welfare gain.

Section 2 defines the private-balance environment, admissible permanent policies, and the local equilibrium derivative. Section 3 proves the Implementation Theorem. Section 4 proves the Genericity Theorem. Section 5 develops the complete three-type Lagos–Wright application. Section 6 concludes and clarifies the scope of the results.

2 Environment, Balances and Equilibrium

This section specifies the private-balance monetary environment, the admissible permanent policy class, and the local equilibrium and welfare objects used below.

2.1 Economic Environment

Time is discrete and infinite, $t = 0, 1, 2, \cdots$. There is a unit mass of atomless, infinitely lived households. Households may belong to finitely many publicly observable classes and may experience privately observed idiosyncratic histories. They discount future utility by a common factor $\beta \in (0, 1)$. The abstract results do not impose a particular within-period trading protocol: preferences, production technologies, matching probabilities, bargaining protocols, type distributions, and other structural features are primitives indexed by $\theta \in \Theta$. What is common across the admissible economies is the monetary, informational, and policy structure described below.

There are perishable consumption goods and intrinsically worthless fiat money. Goods cannot be carried across dates. Fiat money is the only privately held durable object relevant for decentralized exchange and intertemporal liquidity. It yields no direct utility, does not enter production, and is not redeemable for consumption goods, productive assets, or fiscal claims. Ordinary state-contingent taxation, enforceable private credit, and public record keeping sufficient to implement arbitrary transfers are absent. Money can nevertheless have positive value because limited commitment, anonymity, matching, and the absence of public monitoring make liquid balances useful in decentralized exchange. The monetary authority is the issuer of fiat money and can create additional nominal balances; the policies studied here do not confiscate or destroy money.

Decentralized trade is not publicly monitored. An individual's meetings, production, consumption, payments, receipts, and other actions in decentralized trade are not observable to the economy as a whole, and there is no public record from which

they can later be reconstructed. Consequently, decentralized-trade histories cannot support enforceable unsecured credit, public-history punishments, or arbitrary history-contingent transfers. This informational friction makes fiat money essential for decentralized exchange.

At each date, trade and portfolio choice, together with any policy-induced transfers, determine the cross-sectional distribution of nominal balances. A household observes its own balance and private history. The monetary authority observes public aggregate states, including aggregate nominal money, aggregate prices, publicly observable classes, and any other public state specified by the model, but it does not observe individual decentralized-trade actions, private histories, or true individual balances. It observes only the balance voluntarily presented when the permanent transfer policy is applied. The authority therefore cannot condition transfers directly on an unobserved true balance or private trading history. A household may conceal part of its money but cannot present more money than it owns, because presented units must be physically tendered or otherwise verified. This one-sided reporting technology is the source of the monotonicity restriction imposed below.

Predetermined date-0 inherited wealth and public states are fixed when the monetary authority announces the policy. Unlike the one-time, unanticipated intervention in Fukai (2026), the present policy is announced before the date-0 portfolio choice, is anticipated by households, and is applied at every date. Agents therefore choose date-0 and subsequent money portfolios knowing the permanent balance-contingent policy. The policy may change the entire future sequence of portfolios, distributions, nominal scales, prices, and allocations.

The policy experiment is local: the paper compares the zero-policy valued-money equilibrium with nearby valued-money equilibria generated by sufficiently small permanent policies. An equilibrium specifies household decisions, decentralized-trade outcomes, nominal prices and scales, continuation values, public aggregate states, and the full path of money-balance distributions, subject to individual optimiza-

tion, feasibility, market clearing, compliance with the policy rule, and distributional consistency. Ex ante utilitarian welfare is the population average of expected discounted utility evaluated before the policy is announced and before date-0 portfolio choice. The abstract theorem treats the equilibrium equations and welfare functional as model-dependent objects; the Lagos–Wright application in Section 5 verifies them from explicit primitives.

2.2 Concealment and Feasible Permanent Policies

Let $(M_t)_{t\geq 0}$ be the sequence of dated nominal scales, with $M_t > 0$ at every date. Fix a date t and the normalized monetary unit used by the selected equilibrium. Let $x \in X = [0, \bar{x}]$ be an agent's true normalized pre-policy money balance. Its corresponding nominal balance is $m = M_t x$. The monetary authority does not observe x. An agent may present any amount $\widehat{x} \in [0, x]$ and retain the concealed balance $x - \widehat{x}$. An agent may report less than her actual money balance, but she cannot report a balance greater than the amount she holds.

The authority announces a permanent policy a, applied at every date with intensity $\varepsilon \geq 0$. If an agent presents $\widehat{x}$, she receives $M_t \varepsilon a(\widehat{x})$ newly issued nominal units. The concealed balance remains hers, so the transfer is added to, rather than substituted for, her existing money. Aggregate issuance determines the scale factor Γ_t and hence $M_{t+1} = \Gamma_t M_t$. Nominal balances measured before the policy stage use M_t, whereas transfers and balances determined at the policy stage are expressed relative to M_{t+1}. The resulting normalized post-policy balance is denoted by $T_{\varepsilon a}(x)$. Its exact formula is model specific.

All nominal balances are expressed relative to the dated nominal scale M_t. Thus an agent with nominal balance m has normalized balance $x = m/M_t$. A *local equilibrium branch* is the collection of nearby equilibria selected as the policy intensity ε varies around zero, with the policy a and primitives θ held fixed.[1] The normalized

1. Mathematically, this branch is simply a path indexed by the policy intensity ε. We use the

domain $X = [0, \bar{x}]$ is fixed along this selected local equilibrium branch, while the corresponding nominal holding domain is $M_t X = [0, M_t \bar{x}]$ and therefore moves with M_t. The endpoint $\bar{x}$ is also the normalized money-holding ceiling.

Before the policy stage, the nominal upper bound on money holdings is $M_t \bar{x}$. After the transfer policy is applied, aggregate money issuance changes the nominal scale to $M_{t+1} = \Gamma_t M_t$, so the corresponding nominal upper bound becomes $M_{t+1}\bar{x} = \Gamma_t M_t \bar{x}$. If the transfer policy would otherwise leave an agent with a balance above this updated ceiling, the policy rule replaces that balance by the ceiling itself and retires only the excess nominal money. Accordingly, for sufficiently small changes in the policy parameter ε, nominal balances, the nominal support endpoint, and the realized normalized support may all vary, while every normalized balance continues to lie in the fixed state space X. Because the ceiling scales proportionally with the nominal unit, a purely proportional change in the nominal scale does not make the ceiling bind for real reasons.

The following assumptions impose the corresponding local equilibrium and distributional requirements. Define the proportional policy $p \in \mathcal{B}$ by $p(x) := x$. Thus εp gives each reported normalized balance the same proportional increase.

Assumption 1. *For $\theta \in \Theta_0$:*

(i) after all dated nominal variables are divided by M_t, every nearby equilibrium and its welfare are described on the same normalized domain X and the same Banach spaces; the corresponding nominal variables are recovered using M_t;

(ii) common scaling of dated nominal objects and ceiling indexed to the nominal scales leaves the normalized system unchanged;

(iii) every sufficiently small proportional permanent policy εp leaves $T_{\varepsilon p}(x) = x$ and only rescales M_t;

(iv) continuation utility is weakly increasing in post-policy money, and an atomless

term *equilibrium branch* rather than *equilibrium path* because the latter term could be confused with the date-indexed path of equilibrium variables over time t.

individual's report does not affect aggregate prices or public continuation states.

Part (*iii*) is stronger than neutrality at one date: the proportional policy must reproduce the same normalized allocation, portfolio distribution, and continuation values along the entire infinite path. Part (*iv*) supplies the monotonicity conditions used below to establish truthful presentation.

For each type j, let $\nu_{j,t}$ denote the distribution of normalized inherited real balances entering date t. We represent the sequence $(\nu_{j,t})_{t\geq 0}$ in a fixed, policy-independent Banach space $\mathcal{D}_\omega$ of exponentially weighted bounded sequences of finite signed measures on X. Let $\mathbf{z}$ denote the complete normalized equilibrium path other than the current distribution argument written separately below. It includes normalized allocations, prices, continuation values, portfolio choices, the distribution paths of all types, and public aggregate states. Conditional on these equilibrium objects, the permanent policy a, and the primitives θ, the distribution evolves according to

$$\nu_{j,t+1} = \mathcal{Q}_{j,t}(\nu_{j,t}; \mathbf{z}, a, \theta),$$

so that $\mathcal{Q}_{j,t}$ maps the current distribution into the next-period distribution.

Assumption 2. *For every j and t, $\mathcal{Q}_{j,t}$ maps $\mathcal{D}_\omega$ into itself and is C^1 in the current distribution and the nearby equilibrium-policy variables. From the predetermined date-0 distribution, every sufficiently close permanent equilibrium generates a uniformly bounded sequence $(\nu_{j,t})_{t\geq 0} \in \mathcal{D}_\omega$, including when support points move or atoms merge. The date-0 nominal distribution is converted into real units only after P_0 is determined.*

This assumption is not an observability restriction. It requires all nearby distribution paths to be represented in the same fixed space $\mathcal{D}_\omega$. A permanent policy changes all future wealth distributions, but the equilibrium derivative remains well defined because the representation remains valid even when support points move or

distinct atoms merge. The distinction between continuous individual choice and possibly atomic endogenous money distributions is standard in the distributional search literature (Zhou 1999; Berentsen, Camera, and Waller 2005; Molico 2006).

A fixed nominal cap would violate proportional neutrality because a pure nominal expansion could make an agent hit the unchanged cap. A ceiling indexed to the nominal scale is instead expressed in normalized units, or multiplied by the current nominal scale.

Let b denote an agent's normalized balance immediately before the ceiling is applied, and let $\Gamma > 0$ be the contemporaneous scaling factor for nominal balances. The post-policy ceiling is therefore $\Gamma\bar{x}$. The post-policy balance is therefore $\min\{b, \Gamma\bar{x}\}$.

Proposition 1. *For the proportional policy $a = p$ and $\Gamma = 1 + \varepsilon$,*

$$\min\{(1+\varepsilon)x, (1+\varepsilon)\bar{x}\} = (1+\varepsilon)x$$

for all $x \in [0, \bar{x}]$. Hence proportional scaling never triggers retirement of excess money at the ceiling.

Proof of Proposition 1. Since $x \leq \bar{x}$, $(1+\varepsilon)x \leq (1+\varepsilon)\bar{x}$, so the first argument of the minimum is selected. Equality also holds at $x = \bar{x}$. □

The proposition shows that exact proportional neutrality respects, rather than ignores, the upper boundary: the ceiling scales with the monetary unit, so a pure redenomination destroys no money.

We take $\mathcal{B} := C^2(X) := \{a : X \to \mathbb{R} : a,\ a',$ and a'' are continuous on $X\}$, the space of twice continuously differentiable real-valued functions on X, equipped with the norm $\|a\|_{\mathcal{B}} := \max_{x\in X} |a(x)| + \max_{x\in X} |a'(x)| + \max_{x\in X} |a''(x)|$. Thus a policy is a twice continuously differentiable transfer shape on the common normalized balance domain. It must be defined on all of X, including balances not realized in the baseline distribution, because an agent with true balance x compares truthful presentation with every lower report $\widehat{x} \leq x$ and because nearby anticipated-policy equilibria may

change the support. The role of the first two derivatives will be stated when individual portfolio choices are introduced below.

Let $\mathcal{K} := \{a \in \mathcal{B} : a(0) \geq 0,\ a'(x) \geq 0 \text{ for all } x \in X\}$. The condition $a(0) \geq 0$ rules out confiscation from a zero-balance agent, while $a' \geq 0$ ensures that presenting more money never lowers the transfer. Thus, $\mathcal{K}$ is the set of nonnegative, nondecreasing policies; such policies are financed by gross money creation and do not reward concealment. Signed directions outside $\mathcal{K}$ are used only as mathematical directions for welfare differentiation.

Lemma 1. *Under Assumption 1 (iv), every $a \in \mathcal{K}$ makes truthful presentation weakly optimal against all reports $\widehat{x} \leq x$.*

Proof. For every $\widehat{x} \leq x$, monotonicity gives $a(\widehat{x}) \leq a(x)$. Since the concealed balance remains with the agent, truthful presentation yields weakly more total post-policy money. An atomless report does not affect aggregate prices or public continuation states, and continuation utility is weakly increasing in money. Hence $\widehat{x} = x$ weakly dominates every downward-concealment deviation. □

The lemma permits ties. Strict truth telling requires a strictly increasing policy or an appropriate tie-breaking convention. If presentation affects payoff-relevant objects other than post-policy money, those effects must be added to the incentive constraint. The theorem below uses this lemma only when an individual's report affects utility through post-policy money alone and, because the individual is atomless, does not change aggregate prices or public continuation states.

2.3 Local Permanent Equilibrium and Welfare

Let $\mathcal{Z}_\omega$ be an exponentially weighted sequence space for equilibrium paths. Let $\mathcal{Y}_\omega$ be the corresponding exponentially weighted space of residual sequences, where each residual records the dated deviation of an equilibrium condition from zero. Thus, for a candidate path $\mathbf{z}$ under a permanent policy εa, the element $\mathcal{F}(\mathbf{z}, \varepsilon a; \theta) \in \mathcal{Y}_\omega$ measures

the failure of that path to satisfy the equilibrium conditions. The equilibrium system is therefore $\mathcal{F}(\mathbf{z}, \varepsilon a; \theta) = 0$. Thus, a solution $\mathbf{z}$ to this equation is an equilibrium under the permanent policy εa, given the primitives θ. The vector $\mathbf{z}$ contains normalized allocations, prices, value functions, portfolio and inherited-wealth distributions, and any public states. The date-0 inherited datum is fixed across policy comparisons.

Assumption 3. *For each $\theta \in \Theta_0$, $\mathcal{F}(\mathbf{z}^0, 0; \theta) = 0$, $\mathcal{F}$ is C^1 nearby, and $D_{\mathbf{z}}\mathcal{F}(\mathbf{z}^0, 0; \theta) : \mathcal{Z}_\omega \to \mathcal{Y}_\omega$ is an isomorphism. The local branch $\mathbf{z}(\varepsilon a; \theta)$ remains feasible, remains a valued-money equilibrium, and stays in the same nominal-scale-invariant regime.*

The assumption is the dynamic counterpart of an isolated regular equilibrium. It excludes local regime switches and lets the Banach-space implicit-function theorem, stated as Theorem A.1 in Appendix A, convert the infinite equilibrium system into a differentiable reduced welfare map.

Let $W(\varepsilon, a; \theta)$ be ex ante welfare on the selected branch. The zero policy $a = 0$ lies on the boundary of the cone of feasible policies. To evaluate derivatives in both positive and negative directions around this boundary point, we use a smooth extension of the equilibrium and welfare maps to an open neighborhood of $a = 0$ in the ambient Banach space.

A *nonproportional policy direction* is a policy perturbation that changes the transfer schedule relative to the proportional, welfare-neutral direction $p(x) = x$. For the genericity analysis in Section 4, we select finitely many economically meaningful nonproportional policy directions $e^1, \ldots, e^\ell$ and define their span $E := \operatorname{span}\{e^1, \ldots, e^\ell\}$. We call E a *finite-dimensional space of nonproportional policy directions*. Its linearity and finite dimension follow from this construction; they are not consequences of the welfare assumption below. The substantive restriction imposed below is only the regularity of the welfare derivative as primitives vary.

Assumption 4. *For some $L_\theta \in \mathcal{B}^*$, locally uniformly over bounded directions, $W(\varepsilon, a; \theta) = W(0; \theta) + \varepsilon L_\theta a + o(\varepsilon)$ as $\varepsilon \downarrow 0$. The map $\theta \mapsto L_\theta$ is operator-norm*

continuous. For every fixed finite-dimensional space of nonproportional policy directions E used in Section 4, the map obtained by restricting L_θ to E is C^1 in θ.

The scalar $L_\theta a$ includes the direct redistribution of liquidity and every induced first-order response of portfolios, distributions, prices, and continuation values. By Assumption 1 (*iii*),

$$L_\theta p = 0. \tag{1}$$

This single neutral direction is the implementation instrument used below.

The abstract derivative can be expressed directly in terms of the equilibrium operator. This formula is useful because it separates the direct welfare effect of a policy from the indirect effect transmitted through the entire permanent equilibrium path.

Proposition 2. *Under Assumptions 3–4,*

$$D_a \mathbf{z}(0;\theta)[h] = -[D_{\mathbf{z}}\mathcal{F}]^{-1} D_a \mathcal{F}[h], \tag{2}$$

$$L_\theta h = D_a W[h] - D_{\mathbf{z}} W [D_{\mathbf{z}}\mathcal{F}]^{-1} D_a \mathcal{F}[h]. \tag{3}$$

Thus $L_\theta \in \mathcal{B}^*$ is bounded and linear. For p, nominal scales may move although normalized allocations and welfare do not.

Proof of Proposition 2. Differentiate the equilibrium identity $\mathcal{F}(\mathbf{z}(a;\theta), a; \theta) = 0$ at $a = 0$ and use invertibility of $D_{\mathbf{z}}\mathcal{F}$ to obtain (2). The chain rule for reduced welfare gives (3). Boundedness follows from boundedness of the derivatives and of the inverse equilibrium operator. Exact neutrality of proportional transfers implies $L_\theta p = 0$, although the nominal components of $D_a \mathbf{z}(0)[p]$ may be nonzero. □

Economically, (3) makes clear that the theorem does not hold portfolios or distributions fixed. The second term contains every induced response of the infinite continuation path. The Lagos–Wright calculation in Section 5 is a primitive evaluation of this reduced welfare derivative, not a partial-equilibrium transfer derivative.

The anticipated-permanent problem differs from a dated transfer because the policy changes the law of motion of nominal balances at every future date. It is therefore not enough to require a static homogeneity property of a single budget constraint. The complete equilibrium operator must be invariant to any date-specific rescaling that multiplies all nominal objects dated t—including balances, prices, transfers, payments, and any nominal ceiling—by the same strictly positive factor α_t. These arbitrary rescaling factors are distinct from the endogenous equilibrium scale factors $(\Gamma_t)_{t\geq 0}$.

Definition 1. A valued-money equilibrium path is *invariant to nominal rescaling* if any dated rescaling of its nominal objects remains an equilibrium with unchanged normalized real objects, provided the transfer rule and all nominal ceilings are rescaled by the same factors.

The definition is an equilibrium property, not a normalization convention. It requires household budget sets, bargaining payments, market clearing, the transfer policy, government issuance, and all continuation equations to transform coherently. In particular, a fixed nominal cap would violate scalability because a proportional expansion could move an agent against the cap even though all real balances were unchanged. A cap compatible with nominal rescaling must instead be expressed in normalized units or multiplied by the current nominal scale.

For the proportional permanent policy, the relevant scale sequence is endogenous. If the date-t nominal stock is multiplied by $1 + \varepsilon$, then every nominal balance and price at that date is multiplied by the same factor. Reapplying the rule at later dates produces a recursively determined scale sequence, while normalized portfolios, trades, and distributions remain fixed. This gives the following implication of Assumption 1.

Proposition 3. *Under Assumption 1, every sufficiently small εp has a unique representation with the zero-policy normalized equilibrium path and only a changed nominal scale. Hence $W(\varepsilon, p; \theta) = W(0; \theta)$ and $L_\theta p = 0$.*

Proof of Proposition 3. At date 0, a proportional transfer multiplies each agent's nominal post-policy balance and the aggregate money stock by the same factor, so normalized balances are unchanged. Invariance under date-specific nominal rescaling then maps the zero-policy continuation into a nominally rescaled continuation satisfying all household, market-clearing, policy, and government-budget equations. Repeating the argument date by date yields the same normalized allocation and distribution sequence. Uniqueness of the local branch in Assumption 3 identifies this rescaled path with the selected policy equilibrium. Since welfare depends only on normalized real allocations, it is unchanged. Differentiating at zero gives $L_\theta p = 0$. □

This proposition is the key distinction from one-time neutrality. Under a one-time unanticipated intervention, the inherited portfolio distribution is fixed and only the continuation after the transfer must be rescaled. Here the announcement precedes date-0 portfolio choice and the policy is repeated forever. Exact neutrality therefore requires households to choose the same normalized portfolios in anticipation of the rule and requires the entire future equilibrium path to remain on the same real branch.

3 The Implementation Theorem

The goal of this section is the Implementation Theorem: starting from any small policy change that would raise welfare to first order, construct an actual permanent monetary policy that agents report truthfully, that satisfies the feasibility restrictions, and that strictly raises welfare. The proof proceeds in two steps. First, Lemma 2 shows how to add a proportional component to a policy direction. A welfare-improving policy change may prescribe negative transfers or may decrease with reported balances, so it cannot yet be implemented. Adding a sufficiently large proportional component turns it into a nonnegative, strictly increasing policy. Second, Proposition 3 shows why this proportional augmentation does not undo the welfare gain: the proportional component only rescales nominal units and leaves real allocations and welfare un-

changed. The Implementation Theorem combines these two results. The remainder of the section then asks the converse question: when does no small policy change raise welfare to first order? Corollary 1 removes pure nominal rescalings from this search, and Proposition 4 identifies the exact first-order failure conditions that motivate the genericity analysis in Section 4.

Let $\mathcal{C} := \{e \in \mathcal{B} : e(0) \geq 0\}$. No monotonicity is imposed on e.

Lemma 2. *For each $e \in \mathcal{C}$, some $\lambda > 0$ makes $\widehat{e} = \lambda p + e \in \mathcal{K}$, with $\widehat{e}' > 0$ on X and $\widehat{e}(x) > 0$ for $x > 0$.*

Economically, the lemma separates the desired redistribution from the incentive constraint. The direction e contains the economically relevant variation across balance holders. The added term λp gives every additional reported unit of money the same sufficiently large extra transfer, so hiding balances becomes unprofitable. The lemma therefore answers how to make a favorable but possibly nonmonotone redistribution truthful; Proposition 3 answers why this incentive adjustment is welfare neutral.

Proof of Lemma 2. Choose $\lambda > \|e'\|_\infty$. Then $\widehat{e}'(x) = \lambda + e'(x) > 0$ and $\widehat{e}(0) = e(0) \geq 0$, so $\widehat{e}$ is nonnegative and strictly increasing. □

Theorem 1 (Implementation Theorem)**.** *Suppose Assumptions 1–4 hold and $L_\theta e > 0$ for some $e \in \mathcal{C}$. Then some nonproportional $\widehat{e} \in \mathcal{K}$ and $\bar{\varepsilon} > 0$ yield, for every $0 < \varepsilon < \bar{\varepsilon}$, a truthful, nearby valued-money permanent equilibrium with $W(\varepsilon, \widehat{e}; \theta) > W(0; \theta)$.*

If, in addition, the baseline pre-policy balance distribution has a positive mean and that mean varies continuously along the selected equilibrium branch, then the implementing policy is strictly money-financed for all sufficiently small $\varepsilon > 0$.

The theorem is local but fully dynamic. The favorable direction is evaluated after all anticipated portfolio and continuation effects have been incorporated into L_θ. The proportional component does not suppress those effects; it only changes nominal scale and restores feasibility.

A simple two-group example illustrates the construction. Suppose agents with low positive balances have a higher social marginal value of liquidity than agents with high balances. A signed direction that transfers liquidity from the high-balance group to the low-balance group then has a positive welfare derivative, but it may violate monotonicity because the high-balance group receives less. Adding λp raises transfers with reported balances while preserving the cross-group welfare gain. The Implementation Theorem formalizes exactly this logic for the full anticipated permanent equilibrium path.

Proof of the Implementation Theorem. By Lemma 2, choose $\widehat{e} = \lambda p + e \in \mathcal{K}$. Equation (1) gives $L_\theta \widehat{e} = L_\theta e > 0$. If $\widehat{e}$ were proportional, then $e = \widehat{e} - \lambda p$ would be proportional and would have zero derivative, a contradiction. Assumption 4 implies strict welfare improvement for all sufficiently small positive ε, while Assumptions 1 and 3 give truthful implementation and a nearby permanent equilibrium. Strict positivity of $\widehat{e}$ at positive balances implies positive aggregate issuance whenever the policy-dependent distribution of pre-policy money holdings assigns positive probability to strictly positive balances. Since balances are nonnegative, a positive mean guarantees such positive probability. Continuity of the selected distribution branch keeps the mean positive near a positive-mean baseline, so issuance remains strictly positive for all sufficiently small policies. □

For redistributions among positive-balance agents, define $\mathcal{B}_0 := \{e \in \mathcal{B} : e(0) = 0\}$. Two directions in $\mathcal{B}_0$ that differ only by a multiple of the proportional policy p have the same welfare derivative, because $L_\theta p = 0$. We may therefore regard them as the same redistribution after removing the purely proportional change in nominal scale; the quotient space $\mathcal{B}_0/\operatorname{span}\{p\}$ records these equivalence classes. If one class has a nonzero welfare derivative, then either a representative direction or its negative raises welfare to first order, so the Implementation Theorem applies.

Failure at the first-order level means that no sufficiently small feasible policy raises welfare through its linear term. Proposition 4 shows that two separate conditions are

then necessary. First, every redistribution among agents with positive balances must have zero welfare derivative. Second, a policy that gives the same positive transfer to all positive-balance agents must have a weakly nonpositive welfare derivative. Theorem 2 in Section 4 studies the first condition: the exceptional case in which all positive-balance redistribution directions in a chosen finite-dimensional policy menu have zero first-order welfare effects.

Corollary 1. *On $\mathcal{B}_0 = \{e \in \mathcal{B} : e(0) = 0\}$, proportional neutrality makes L_θ a bounded functional on $\mathcal{B}_0/\operatorname{span}\{p\}$. If it is nonzero, the Implementation Theorem applies.*

The quotient formulation separates economically meaningful redistribution from a pure change of units. It also clarifies the role of a nondegenerate positive-balance distribution. At least two positive balance positions are needed for a nontrivial redistribution among positive-balance agents, although the Implementation Theorem itself may still apply through a uniform injection when the positive-balance face is degenerate.

Proof of Corollary 1. If $e_1 - e_2 = \lambda p$, Proposition 3 implies $L_\theta e_1 = L_\theta e_2$. Hence the derivative is well defined on equivalence classes modulo proportional scaling. If the quotient functional is nonzero, choose a representative $e \in \mathcal{B}_0$ with $L_\theta e \neq 0$ and reverse its sign if necessary. Both orientations satisfy $e(0) = 0$, so the favorable orientation belongs to $\mathcal{C}$ and the Implementation Theorem applies. □

Proposition 4. *Under proportional neutrality, $L_\theta e \leq 0$ for all $e \in \mathcal{C}$ if and only if $L_\theta e = 0$ for every $e \in \mathcal{B}_0$ and $L_\theta \mathbf{1} \leq 0$, where $\mathbf{1}$ is the constant-one policy. Equivalently, the first condition may be written $L_\theta e = 0$ for every $e \in \mathcal{B}_0$: every positive-balance redistribution direction has zero first-order welfare effect.*

The proposition identifies exactly when the welfare-improving policy construction in the Implementation Theorem can fail at first order. Failure of the theorem requires

every redistribution among positive-balance agents to have zero first-order welfare effect and a uniform positive injection to have a weakly nonpositive first-order welfare effect.

Proof of Proposition 4. Every $e \in \mathcal{B}$ has the unique decomposition $e = e(0)\mathbf{1} + s$ with $s \in \mathcal{B}_0$. If there exists some $s \in \mathcal{B}_0$ with $L_\theta s \neq 0$, then one of the two directions s and $-s$ raises welfare to first order and belongs to $\mathcal{C}$. If instead $L_\theta s = 0$ for every $s \in \mathcal{B}_0$ but $L_\theta \mathbf{1} > 0$, the constant policy is favorable. Conversely, if both displayed conditions hold, then for every $e \in \mathcal{C}$, $L_\theta e = e(0)L_\theta \mathbf{1} + L_\theta s \leq 0$, so no feasible comparison direction is favorable. □

The first requirement about the zero first-order welfare effect consists of infinitely many equalities, one for each positive-balance redistribution direction. The Genericity Theorem in Section 4 does not analyze this entire infinite-dimensional collection at once. Instead, it shows that all selected first-order welfare effects being zero on any fixed finite set of economically meaningful nonproportional policy directions is exceptional under a verifiable rank condition.

4 The Genericity Theorem

Proposition 4 in the previous section says that the Implementation Theorem can fail only if every redistribution among positive-balance agents has zero first-order welfare effect and a uniform positive injection does not help. Checking every smooth redistribution at once is an infinite-dimensional problem. This section therefore takes a finite and economically interpretable route. We choose a small menu of redistributions—for example, transfers from a high-balance group to a low-balance group—and ask whether all of their welfare derivatives can vanish simultaneously. The Genericity Theorem below shows that, when primitive parameters move these derivatives independently, all selected first-order welfare effects being zero occurs only on a lower-

dimensional exceptional set. The rank condition should therefore be interpreted and verified in economically meaningful primitive terms in each application.

4.1 Conditional Genericity

Fix an economically chosen ℓ-dimensional space of nonproportional policy directions $E \subset \mathcal{B}_0$ with $E \cap \operatorname{span}\{p\} = \{0\}$. Economically, this restriction excludes from E every nonzero policy that is merely proportional to current balances. Such a policy gives all positive-balance agents the same percentage increase, leaves their relative real balances unchanged, and only rescales the nominal unit; it is therefore not a genuine redistribution and is welfare neutral by proportional neutrality. Thus every nonzero direction in E must change the transfer schedule relative to proportional scaling.

The restriction also prevents the genericity argument from counting the identically neutral direction p as one of the ℓ economically distinct nonproportional policy directions: otherwise one component of the welfare-derivative vector would vanish identically and the full-rank condition below could not hold. Its basis directions $e^1, \ldots, e^\ell$ may represent, for example, transfers from high-balance groups to low-balance groups.

Define $G_E(\theta) := (L_\theta e^1, \ldots, L_\theta e^\ell) \in \mathbb{R}^\ell$ and $\mathcal{E}_E := G_E^{-1}(0)$. Thus $G_E(\theta)$ records the marginal welfare effects of the selected redistributions, and $\mathcal{E}_E$ is the set of economies in which all of them are neutral simultaneously. With one redistribution, an economy is in $\mathcal{E}_E$ only when that redistribution has exactly zero marginal welfare effect. With several redistributions, every corresponding marginal welfare effect must be zero at the same primitive vector.

Assumption 5. *The map $G_E : \Theta_0 \to \mathbb{R}^\ell$ is C^1 and $\operatorname{rank} D_\theta G_E(\theta) = \ell$ for all $\theta \in \mathcal{E}_E$.*

Assumption 5 says that these ℓ neutrality conditions are locally independent. Economically, primitive changes must be able to alter the selected welfare effects in distinct ways. For instance, if one nonproportional policy direction transfers liquidity

from a high-balance group to a low-balance group, a bargaining parameter may change the relative value of liquidity for those two groups. With two nonproportional policy directions, a bargaining parameter and a meeting-rate parameter may move the two marginal welfare effects separately. The assumption excludes parameterizations in which distinct redistributions are forced to have mechanically identical welfare effects.

Theorem 2 (Genericity Theorem)**.** *Under Assumptions 1–5, $\mathcal{E}_E$ is empty or a relatively closed codimension-ℓ C^1 submanifold of Θ_0, hence nowhere dense and null. For $\theta \notin \mathcal{E}_E$, the Implementation Theorem applies.*

The Genericity Theorem supplies the favorable direction required by the Implementation Theorem. Outside $\mathcal{E}_E$, at least one selected redistribution has a nonzero welfare derivative. If a direction e lowers welfare, the reverse redistribution $-e$ raises welfare by the same first-order amount. The Implementation Theorem then augments that favorable direction with a proportional component into a nonnegative, strictly increasing, truthful, and welfare-improving permanent policy. The theorem therefore does not search the full infinite-dimensional policy space. It says that, within any fixed economically meaningful finite menu satisfying the rank condition, all selected first-order welfare effects being zero is confined to a lower-dimensional exceptional set.

Proof of the Genericity Theorem. Assumption 5 makes zero a regular value of G_E. The regular-value theorem, stated as Theorem A.2 in Appendix A, implies that $G_E^{-1}(0)$ is empty or a codimension-ℓ embedded submanifold; continuity makes it relatively closed, and positive codimension implies nowhere density and nullity. If $\theta \notin \mathcal{E}_E$, then $L_\theta e \neq 0$ for some $e \in E$. Since E is linear, $-e \in E$; hence one of e and $-e$ has a positive derivative. Because this orientation belongs to $\mathcal{C}$, the Implementation Theorem applies. □

The following adjacent-balance example illustrates the codimension conclusion of the Genericity Theorem. Suppose positive money balances take the values $x_1 < \cdots <$

x_J, and let μ_j be the mass of agents holding x_j. For each adjacent pair (x_k, x_{k+1}), let e^k represent a marginal redistribution of liquidity from agents at x_{k+1} to agents at x_k, calibrated so that aggregate liquidity remains unchanged and zero-balance agents are unaffected. Concretely, normalize the direction by setting the transfer per agent equal to μ_{k+1} for group k, $-\mu_k$ for group $k+1$, and zero for every other group. Then the aggregate change in liquidity is $\mu_k\mu_{k+1} + \mu_{k+1}(-\mu_k) = 0$. Thus this policy direction changes only the distribution of liquidity across the two groups, not its aggregate quantity. If $v_j(\theta)$ is the social marginal value of one additional unit of liquidity for group j, then $L_\theta e^k = \mu_k\mu_{k+1}\big[v_k(\theta) - v_{k+1}(\theta)\big]$. Hence the k-th redistribution is neutral exactly when the two adjacent groups have the same social marginal value of liquidity.

All $J-1$ adjacent redistributions are neutral only if $v_1(\theta) = v_2(\theta) = \cdots = v_J(\theta)$. These are $J-1$ equalities. If bargaining, matching, or transaction primitives can move the $J-1$ adjacent marginal-value differences independently, Assumption 5 holds with $\ell = J-1$, and the Genericity Theorem implies that all selected first-order welfare effects being zero is locally a codimension-$(J-1)$ event. Economically, this is a knife-edge outcome: groups with different bargaining positions, meeting opportunities, or transaction technologies will not ordinarily have exactly the same social marginal value of liquidity.

The conclusion is invariant to a smooth reparameterization of primitives. A non-singular change of coordinates preserves the rank of DG_E at its zeros and therefore preserves the codimension of $\mathcal{E}_E$. Parameter restrictions must nevertheless be built into the primitive manifold. For example, with n type shares summing to one, only $n-1$ shares are independent and should be used as coordinates.

Corollary 2. *Let $E_1 \subset E_2 \subset \cdots \subset \mathcal{B}_0$ be finite-dimensional with $E_n \cap \text{span}\{p\} = \{0\}$, and set $\mathcal{E}_n = \{\theta : L_\theta e = 0$ for every $e \in E_n\}$. Then*

$$\mathcal{N} := \{\theta : L_\theta e \leq 0 \text{ for every } e \in \mathcal{C}\} \subseteq \bigcap_{n \geq 1} \mathcal{E}_n.$$

If $\bigcup_n E_n$ *is dense in* $\mathcal{B}_0$*, the intersection equals* $\{\theta : L_\theta e = 0$ *for every* $e \in \mathcal{B}_0\}$*.*

Corollary 2 explains what happens when the menu of nonproportional policy directions is enlarged. With one direction, failure requires one redistribution to be neutral; after adding a second, both must be neutral; and each further economically distinct direction adds another test. The possible failure set can therefore only shrink. The corollary should not be read as a genericity result on the entire infinite-dimensional policy space. Rather, it organizes a sequence of finite-dimensional tests: each chosen space of nonproportional policy directions yields one such test, and enlarging the space makes the test progressively more demanding.

Proof of Corollary 2. If $\theta \in \mathcal{N}$, no direction in the linear space E_n can have a nonzero derivative, because one of its two orientations would then be favorable. Thus $\theta \in \mathcal{E}_n$ for every n. If the union is dense and L_θ vanishes on every E_n, continuity extends the equality to all of $\mathcal{B}_0$. The converse is immediate. □

For applications, the objective is not to make E as large as possible. Instead, one should select a few economically interpretable nonproportional policy directions and ask whether the model's primitive parameters can vary their welfare effects independently. This is the substantive content of the rank condition. The next subsection gives a primitive-level sufficient condition for Assumption 5.

4.2 Economic Interpretation of the Rank Condition

Assumption 5 requires the selected welfare effects to respond independently to changes in economic primitives. Economically, this means that matching rates, bargaining shares, preferences, or other structural parameters do not move the liquidity values of all balance groups in exactly the same way. If distinct primitive changes alter distinct balance-group comparisons, then the equalities defining all selected first-order welfare effects being zero impose genuinely independent restrictions rather than repeating the same condition in different notation.

This interpretation is useful for applications, but it does not require a separate genericity theorem. The appropriate check is model specific. One may verify directly that the Jacobian of the selected welfare derivatives has full rank at every relevant neutrality point. Alternatively, an application may establish the stronger conclusion that at least one economically meaningful nonproportional policy direction has a strictly nonzero welfare derivative throughout the parameter region of interest. In that case the exceptional set is empty and no rank calculation is needed.

The Lagos–Wright application below takes the second route. It constructs a redistribution from a high-balance group to a low-balance group and shows that its welfare derivative is strictly positive in a neighborhood of the benchmark. Thus the application verifies Assumption 5 vacuously, because there is no parameter value in that neighborhood at which all selected nonproportional policy directions are neutral.

5 A Full Primitive Application

This section gives a complete application of the Implementation Theorem and the Genericity Theorem in a three-type Lagos–Wright economy. We proceed in one line: construct a regular stationary monetary equilibrium, extend it to the anticipated permanent-policy branch, verify truthful reporting and proportional neutrality, compute the welfare derivative, and exhibit a favorable nonproportional policy direction. Every substantive assumption used by the Implementation Theorem and the Genericity Theorem is verified inside the model.

5.1 Environment and Stationary Equilibrium

There is a unit mass of infinitely lived agents divided into permanent types $j \in \{0, 1, 2\}$ with shares $\mu_j > 0$ and $\sum_j \mu_j = 1$. Each agent receives her permanent type $j \in \{0, 1, 2\}$. The population shares and the type-dependent primitives are common knowledge.

Each date has a centralized market (CM), a policy stage, and a decentralized bilateral-trade market (DM) in this order. Agents' types differ only in their DM role probability σ_j and DM bargaining share η_j; preferences, technologies, the discount factor, and all other primitives are common across types. Within a bilateral DM trade, each participant observes the counterparty's type before bargaining; this bilateral observation is not public and does not reveal type to the monetary authority. Accordingly, the transfer policy may condition on the agent's reported money balance, but not directly on type.

If types were publicly observable to the monetary authority, the unique type-specific portfolio choices characterized below would reveal each type's equilibrium money balance. The authority could then condition transfers directly on the observable type, or equivalently on its implied equilibrium balance, so truthful balance reporting would not be a separate implementation problem on the local equilibrium branch. We retain types that are unobservable to the monetary authority, precisely to make the reporting constraint substantive.

At the policy stage, an agent presents a balance, the monetary authority delivers the transfer prescribed by the permanent policy in newly issued money, and aggregate issuance updates the nominal scale before the resulting post-policy balance is carried into the DM. This stage implements the transfer policy and the associated nominal accounting.

Money is fiat: it provides no direct utility, enters neither production nor consumption, and carries no redemption value or fiscal claim. DM trade is subject to no monitoring: an individual's DM actions are not publicly observable to the economy as a whole, so neither the monetary authority nor other agents can observe individual transactions or infer resulting balances from a public history. Unsecured credit and public-history enforcement are therefore unavailable, and fiat money is essential for DM trade. The transfer policy can depend only on a balance report, not on monitored trade or credit records.

In the DM, type j is a buyer with probability $\sigma_j \in (0, 1/2)$, a seller with the same probability σ_j, and does not find a partner with probability $1-2\sigma_j$. The utility from having no partner in the DM is normalized to 0. Conditional on being a buyer and meeting a seller, type j receives the proportional bargaining share $\eta_j \in (0, 1)$ of total match surplus, following Kalai (1977). The bargaining share depends only on the buyer's permanent type and not on the seller's type. Thus, when a type-j buyer meets a type-k seller, the buyer receives share η_j and the seller receives share $1-\eta_j$, for every k. Matching is feasible when $\sum_j \mu_j \sigma_j < 1/2$.

In a DM meeting, buyer utility and seller cost are $u(q) = Aq - Bq^2/2$ and $c(q) = q$ with $A > 1$ and $B > 0$. CM preferences are quasi-linear. Agents discount future utility or cost with discount factor $\beta \in (0, 1)$.

Total surplus is $S(q) = (A-1)q - Bq^2/2$, maximized at $q^* = (A-1)/B$. Under the Kalai proportional solution, the real-balance payment required for quantity q is $g(q, \eta) = u(q) - \eta S(q) = \eta q + (1-\eta)u(q)$. For balances below the efficient-liquidity threshold, let $Q(x, \eta)$ solve $g(Q, \eta) = x$ and define realized surplus $F(x, \eta) = S(Q(x, \eta))$. Then

$$F_x(x, \eta) = \frac{S'(Q)}{g_q(Q, \eta)} > 0 \quad \text{and} \quad F_{xx}(x, \eta) = -\frac{B}{g_q(Q, \eta)^3} < 0. \tag{4}$$

Thus liquidity raises DM surplus at a decreasing rate.

Quasi-linearity reduces type j's stationary real-balance choice to

$$\max_{x \geq 0} \left\{ -(1-\beta)x + \beta \sigma_j \eta_j F(x, \eta_j) \right\}. \tag{5}$$

Choose target balances $0 = x_0 < x_1 < x_2$, with both x_1 and x_2 strictly below their efficient-liquidity thresholds. For active types $j = 1, 2$, set

$$\sigma_j = \frac{1-\beta}{\beta \eta_j F_x(x_j, \eta_j)} \tag{6}$$

and choose type-0 primitives so that the right derivative of (5) at zero is strictly negative. By (4), each objective is strictly concave on the binding branch, so these primitives generate the unique stationary support $0 = x_0 < x_1 < x_2$. Strict probability and matching inequalities define a nonempty open primitive region.

Let $\bar{x} = \mu_1 x_1 + \mu_2 x_2$ and let M be the stationary nominal stock. The nominal scale is $P = M/\bar{x}$, giving support points $p_j = P x_j$. Individual choices are continuous, while the stationary distribution is the nondegenerate atomic distribution $\sum_j \mu_j \delta_{p_j}$, where δ_{p_j} denotes the Dirac measure that assigns unit mass to the point p_j.

Lemma 3. *On the binding-liquidity region, Q and F are smooth, with $F_x > 0$ and $F_{xx} < 0$. Extending F constantly beyond efficient liquidity gives an increasing, concave, C^1 realized-surplus function.*

Lemma 3 translates the bilateral bargaining problem into a well-behaved reduced-form liquidity technology. Economically, an additional unit of real balances relaxes the buyer's payment constraint and raises the quantity traded, but the associated surplus gain becomes smaller as the efficient quantity is approached. The flat extension beyond efficient liquidity records that extra money is then redundant for DM trade. Mathematically, smoothness and strict concavity make household portfolio choices regular and locally unique and provide the curvature used later in the welfare derivative.

Proof of Lemma 3. For $q \leq q^*$,

$$g_q(q, \eta) = \eta + (1 - \eta) u'(q) \geq 1,$$

so the inverse-function theorem gives a smooth Q. Differentiating $F(x, \eta) = S(Q(x, \eta))$ yields the formulas in (4). At the efficient quantity $S'(q^*) = 0$, so the derivative joins continuously to the flat branch. □

Strict concavity makes each prescribed active balance the unique global optimum. Here, there is no exogenous physical storage nor handling cost. There is however

an endogenous intertemporal opportunity cost represented by the term $-(1-\beta)x$ in (5): acquiring an additional unit of real balances in the CM requires the agent to give up one unit of current CM consumption (or to supply the equivalent amount of labor), while an unspent unit carried back to the next CM is worth only its discounted continuation value β. The resulting net cost is therefore $x - \beta x = (1-\beta)x$. Balances above the efficient-liquidity threshold are dominated because they raise this endogenous opportunity cost without increasing DM surplus.

Proposition 5. *For any $0 = x_0 < x_1 < x_2$ below the relevant efficient-liquidity thresholds, a nonempty open set of primitives makes* (6) *interior, matching feasibility strict, $x_0 = 0$ a strict type-0 corner, and x_1 and x_2 the unique stationary portfolios of types 1 and 2.*

Proposition 5 shows that the desired three-point monetary distribution is not produced by knife-edge parameter choices. Type 0 optimally remains at the reporting boundary with zero balances, whereas types 1 and 2 hold distinct positive balances because their expected DM gains differ. The strict inequalities imply robustness: small changes in preferences, bargaining shares, or meeting probabilities preserve the same qualitative equilibrium. This open-set result supplies the primitive baseline required for a genuine complete application rather than a single calibrated example.

Proof of Proposition 5. Choose β sufficiently close to one and bargaining shares in a compact subset of $(0,1)$. Since $F_x(x_j, \eta_j) > 0$, formula (6) then gives probabilities strictly between zero and one. Type 0 can be assigned a sufficiently small buyer probability so that the right derivative of its portfolio objective at zero is negative. Lemma 3 makes the active objectives strictly concave and their first-order conditions hold at the prescribed targets. The probability, corner, branch, and matching restrictions are strict and therefore persist on an open neighborhood. □

Three types are minimal for the full application: the zero-balance type creates the feasibility and reporting boundary, while the two positive-balance types permit a

nonproportional redistributive policy direction.

5.2 Truthful Reporting and Permanent Continuation

A permanent normalized policy a is announced before the date-0 CM portfolio choice and applied after the portfolio choice at every date. At a date with nominal scale P_t, a truthful type-j portfolio $x_j(\varepsilon, a)$ corresponds to nominal balances $P_t x_j(\varepsilon, a)$ and receives transfer $P_t \varepsilon a(x_j(\varepsilon, a))$. Aggregate issuance changes the next nominal scale by

$$\Gamma(\varepsilon, a) = 1 + \frac{\varepsilon \sum_j \mu_j a(x_j(\varepsilon, a))}{\sum_j \mu_j x_j(\varepsilon, a)}. \tag{7}$$

The normalized post-policy balance is

$$y_j(\varepsilon, a) = \frac{x_j(\varepsilon, a) + \varepsilon a(x_j(\varepsilon, a))}{\Gamma(\varepsilon, a)}. \tag{8}$$

These formulas imply $\sum_j \mu_j y_j(\varepsilon, a) = \sum_j \mu_j x_j(\varepsilon, a)$: the policy redistributes normalized liquidity but creates no real resources.

The permanent portfolio problem is the stationary recursive problem obtained by inserting (8) into current DM surplus and the next CM continuation. Quasi-linearity makes inherited nominal wealth enter the CM value function affinely. Consequently, after the affine term is removed, the portfolio first-order conditions and market clearing form a finite system in (x_0, x_1, x_2, P), while the remaining value constants and wealth distributions are uniquely reconstructed. At the baseline the active-type Hessians are strictly negative, type 0 has a strict corner, and all liquidity and ceiling constraints are slack. The finite Jacobian is therefore nonsingular on a sufficiently small open primitive set.

The policy-stage formulas can be derived directly from nominal accounting. Let M_t be the nominal stock before the policy stage and let P_t be the nominal price of

one unit of normalized balances. CM money-market clearing requires

$$M_t = P_t \sum_{j=0}^{2} \mu_j x_{j,t}. \tag{9}$$

The permanent policy issues $P_t \varepsilon a(x_{j,t})$ to type j. Hence the post-issuance stock is

$$M_{t+1} = P_t \left[\sum_j \mu_j x_{j,t} + \varepsilon \sum_j \mu_j a(x_{j,t}) \right].$$

Dividing it by (9) gives (7). On the nominal-scale-invariant branch $P_{t+1} = \Gamma_t P_t$, and dividing each post-transfer nominal balance by P_{t+1} gives (8). Aggregating across types yields

$$\sum_j \mu_j y_{j,t} = \sum_j \mu_j x_{j,t}. \tag{10}$$

Thus money creation changes the nominal stock and redistributes a fixed quantity of aggregate real liquidity through the policy stage. It does not create real resources.

Government accounting is correspondingly simple. Fiat money is not redeemed and the authority purchases no goods, so the government budget condition is the law of motion for the nominal stock itself. Starting from $M_0 > 0$, the sequence (Γ_t) uniquely determines (M_t), and (9) then determines prices. No taxes, bonds, or additional fiscal instrument are required.

Quasi-linearity closes the remaining infinite-horizon objects. Let w be inherited real monetary wealth on entering the CM and let $\mathbf{s}_t$ denote the normalized public state. The CM value function has the affine form

$$V_{j,t}(w, \mathbf{s}_t) = w + v_{j,t}(\mathbf{s}_t). \tag{11}$$

Indeed, substituting the linear CM budget constraint into quasi-linear utility makes inherited wealth enter additively with coefficient one. The next portfolio is chosen by maximizing the remaining term and is therefore independent of w.

Let $v_{j,t}(\mathbf{s}_t)$ denote the wealth-independent component of the type-j CM value function in (11). Given the normalized allocation at date t, let $b_{j,t}(\mathbf{s}_t)$ denote the corresponding one-period expected payoff, including CM utility, expected DM surplus, portfolio-acquisition costs, and policy-stage payoffs, but excluding the discounted continuation constant. Suppressing the public-state argument when no confusion can arise, these additive constants solve the recursions

$$v_{j,t} = b_{j,t} + \beta v_{j,t+1}, \tag{12}$$

On an exponentially weighted sequence space with weight strictly larger than β, the forward operator has norm below one. Hence (12) has the unique bounded solution

$$v_{j,t} = \sum_{s=t}^{\infty} \beta^{s-t} b_{j,s}.$$

The inherited-wealth distributions satisfy analogous affine forward equations once portfolios, meetings, payments, and the policy stage are fixed. Their coefficients are probabilities and remain uniformly bounded on the local branch, so the predetermined date-0 distribution generates a unique bounded distribution sequence.

Proposition 6. *Every strict baseline admits a unique C^1 anticipated permanent-policy equilibrium branch with bounded portfolios, unchanged active branches and type-0 corner. The finite portfolio–policy–clearing block is locally one-to-one with the complete normalized infinite-horizon equilibrium: omitted values satisfy*

$$v_{j,t} = \sum_{n \geq 0} \beta^n b_{j,t+n},$$

the predetermined date-0 nominal distribution and stock uniquely reconstruct the bounded wealth-distribution, price, money-stock, transfer, and payment paths, and any ceiling indexed to the nominal scale can be chosen uniformly slack. For $a = p$, $\Gamma = 1 + \varepsilon$ and $y_j = x_j$, so the complete normalized path is unchanged.

Proposition 6 says that, once households' portfolio choices and the policy-stage accounting are determined, the rest of the economy is not free to drift: continuation values, wealth distributions, nominal money stocks, prices, transfers, and payments are pinned down by the predetermined initial state and the permanent policy. The C^1 branch means that small policy changes generate small and differentiable equilibrium responses, which is exactly what the local welfare analysis requires. The proportional-policy clause identifies the nominal rescaling direction: it expands all nominal balances and the price level proportionally while leaving every normalized real allocation unchanged, thereby verifying the neutrality hypothesis used in the Implementation Theorem.

Proof of Proposition 6. Strict concavity for active types, the strict type-0 corner, and the positive baseline value gaps imply that, after shrinking the primitive-policy neighborhood, each household's local solution is its unique global portfolio optimum and the active bargaining branches remain fixed. Order the finite unknowns as (x_1, x_2, P, Γ). At the baseline the Jacobian has block form

$$DR = \begin{pmatrix} H & b \\ c' & D \end{pmatrix},$$

where the portfolio block H is diagonal with entries $\beta\sigma_j\eta_j F_{xx}(x_j, \eta_j) < 0$, while the accounting block has nonzero diagonal terms $\partial R_M/\partial P = \bar{x} > 0$ and $\partial R_S/\partial\Gamma = 1$. Atomlessness removes direct individual effects on aggregates, and the remaining aggregate feedback is continuous. Hence the Jacobian remains nonsingular on a sufficiently small open neighborhood, and Theorem A.1 gives a unique C^1 finite branch.

It remains to show that this finite branch is a complete equilibrium rather than a truncated system. Equation (11) removes inherited wealth from portfolio optimality. On the weighted sequence space, the value recursion is $(I - \beta S)v_j = b_j$, where S is the forward shift and the maintained weight gives $\beta\|S\| < 1$. The Neumann series therefore gives the displayed unique bounded C^1 value sequence. Conditional on

portfolios, meetings, payments, and the policy stage, each type's inherited-wealth distribution obeys an affine forward equation through the mapping $\mathcal{Q}_{j,t}$. Its coefficients are probabilities, remain uniformly bounded, and vary C^1 on the branch; iteration from the predetermined date-0 distribution therefore gives a unique bounded path in $\mathcal{D}_\omega$.

Let $\widetilde{\nu}_{j,0}$ denote the predetermined type-j nominal-wealth distribution and let

$$\mathsf{M}_0^- := \sum_j \mu_j \int m \, \widetilde{\nu}_{j,0}(dm) > 0.$$

For a normalized date-0 portfolio with $\bar{x}_0 > 0$, money-market clearing uniquely determines

$$P_0 = \frac{\mathsf{M}_0^-}{\bar{x}_0}.$$

Pushing the initial nominal distributions forward under $m \mapsto m/P_0$ initializes the real distributional recursion. Starting from the predetermined nominal stock, equations (7) and (9) then reconstruct the entire nominal stock and price paths, and scalability reconstructs transfers and payments. Conversely, normalizing any nearby nominal equilibrium recovers the same finite system. Thus the finite and complete formulations are locally one-to-one.

If a ceiling indexed to the nominal scale is desired, compactness of the local portfolio-policy set provides a common upper bound on all $x_{j,t}$ and $y_{j,t}$; choosing the normalized ceiling strictly above that bound makes it uniformly slack. Finally, for $a = p$, equation (7) gives $\Gamma = 1 + \varepsilon$ and equation (8) gives $y_j = x_j$, proving exact neutrality of proportional transfers. □

Economically, the proposition shows that the permanent-policy experiment is a genuine equilibrium perturbation of the original economy rather than an independent reset of its continuation variables.

Proposition 7. *On the local branch, every $a \in \mathcal{K}$ is truthful against all feasible*

downward reports; truth telling is strict if a is strictly increasing on reachable balances.

Because a nondecreasing policy never rewards a lower report, concealment cannot improve either current liquidity or continuation opportunities. Hence the proposition verifies the incentive-compatibility requirement of the abstract theorems; strict monotonicity removes indifference and makes truthful reporting unique.

Proof of Proposition 7. An agent with true normalized balance x who reports $\widehat{x} \leq x$ retains the hidden balance and receives $\varepsilon a(\widehat{x})$. Monotonicity gives $a(\widehat{x}) \leq a(x)$. Atomlessness makes the aggregate scale factor and public continuation state independent of one agent's report. Because additional post-policy money weakly expands current DM and future CM opportunities, truthful reporting weakly dominates concealment. Strict monotonicity makes the resource inequality strict for every lower report that changes the transfer. □

5.3 Welfare Derivative and Proportional Neutrality

The anticipated permanent experiment differs fundamentally from a one-time transfer. A household chooses its CM portfolio knowing that the same policy will be applied after every future portfolio choice. A perturbation of the transfer level is therefore partly offset by endogenous portfolio adjustment. The first-order real effect is governed by the slope of the transfer schedule relative to aggregate proportional scaling.

Let $x_j(\varepsilon, a)$ be type j's pre-policy portfolio, let $y_j(\varepsilon, a)$ be its post-policy real balance, and define

$$\kappa(a) := \frac{\sum_{j=1}^{2} \mu_j a(x_j)}{\bar{x}}.$$

For a fixed policy direction a, let $\Gamma(\varepsilon, a)$ denote the equilibrium scale factor along the selected permanent-policy branch. Holding the primitives and the direction a fixed while portfolios adjust endogenously along the equilibrium branch, differentiating (7)

with respect to the policy intensity at the zero-policy baseline gives

$$\left.\frac{\partial \Gamma(\varepsilon, a)}{\partial \varepsilon}\right|_{\varepsilon=0} = \kappa(a).$$

Because CM utility is quasi-linear, aggregate affine wealth transfers cancel in utilitarian welfare, while the policy stage uses no real resources. Hence, up to a policy-independent constant,

$$\mathcal{W}(\varepsilon, a) = \frac{\beta}{1-\beta} \sum_{j=0}^{2} \mu_j \sigma_j F(y_j(\varepsilon, a), \eta_j). \tag{13}$$

The private portfolio cost remains in household optimization but is not a net social resource cost. This private–social distinction is essential for the derivative.

Proposition 8. *For $a \in \mathcal{B}$ with $a(0) = 0$,*

$$\left.\frac{\partial y_j(\varepsilon, a)}{\partial \varepsilon}\right|_{\varepsilon=0} = -\frac{a'(x_j) - \kappa(a)}{\beta \sigma_j \eta_j F_{xx}(x_j, \eta_j)} \quad \text{for } j = 1, 2, \tag{14}$$

$$La = -\sum_{j=1}^{2} \frac{\mu_j \left[a'(x_j) - \kappa(a)\right]}{\beta \sigma_j \eta_j^2 F_{xx}(x_j, \eta_j)}. \tag{15}$$

Equation (15) *gives the first-order welfare effect of the policy direction a and thereby defines the linear welfare functional $L \in \mathcal{B}^*$. Moreover, $Lp = 0$.*

Proposition 8 converts the complete equilibrium response into the linear welfare functional required by the general theory. The term $a'(x_j)$ measures how the policy changes the marginal reward to holding liquidity at type j's equilibrium balance, while $\kappa(a)$ removes the component that is merely a common proportional expansion of nominal balances. Curvature $F_{xx} < 0$ determines how strongly portfolios adjust, and the factor $1/\eta_j$ reflects the wedge between the private share of match surplus and the full surplus valued by the planner. The equality $Lp = 0$ is therefore both an algebraic and an economic neutrality statement: pure proportional rescaling has no first-order welfare effect.

Proof of Proposition 8. Differentiate the active types' permanent portfolio first-order conditions. At the baseline, the derivative of the marginal return generated at the policy stage is $a'(x_j) - \kappa(a)$, while the derivative through realized DM surplus is

$$\beta\sigma_j\eta_j F_{xx}(x_j, \eta_j) \left.\frac{\partial y_j(\varepsilon, a)}{\partial \varepsilon}\right|_{\varepsilon=0}.$$

This gives (14). Differentiating (13), substituting (14), and using the stationary portfolio condition $\beta\sigma_j\eta_j F_x(x_j, \eta_j) = 1 - \beta$ yields (15). For $p(x) = x$, $p'(x_j) = 1$ and $\kappa(p) = 1$, so every term vanishes. Exact neutrality of proportional transfers also follows directly from Proposition 6. □

Equation (14) is the central distinction between anticipated permanent and one-time policy. A pure level perturbation is crowded out through the CM portfolio response. What survives is the change in the marginal transfer, $a'(x_j)$, relative to aggregate proportional scaling, $\kappa(a)$. The welfare coefficient reflects that the buyer internalizes only the Kalai share η_j of total DM surplus, while social welfare values the entire surplus.

5.4 A Favorable Policy Direction

Fix a primitive vector in the maintained open region. Choose a smooth cutoff equal to one on a neighborhood of $[0, x_2]$ and define a bounded C^2 policy direction e satisfying

$$e(0) = e(x_1) = e(x_2) = 0, \quad e'(x_1) = 1, \quad e'(x_2) = 0. \tag{16}$$

To see explicitly that such a direction exists, consider on a neighborhood of $[0, x_2]$ the polynomial

$$\frac{x(x - x_1)(x - x_2)^2}{x_1(x_1 - x_2)^2}.$$

The factors x, $x - x_1$, and $(x - x_2)^2$ ensure, respectively, that the polynomial is zero at 0 and x_1, and that both its value and derivative are zero at x_2. The denominator

normalizes its derivative at x_1 to one. Multiplying this polynomial by a smooth cutoff function that equals one on this neighborhood and vanishes outside a slightly larger compact set gives a bounded C^2 policy direction e satisfying (16). Since e vanishes at all support points, $\kappa(e) = 0$.

Proposition 9. *For the policy direction e in* (16)*,*

$$Le = -\frac{\mu_1}{\beta\sigma_1\eta_1^2 F_{xx}(x_1, \eta_1)} > 0. \tag{17}$$

Hence a sufficiently steep proportional component is nonnegative, strictly increasing, truthful, and welfare improving for all small positive intensities.

Proposition 9 supplies the economically favorable redistribution direction that the Implementation Theorem needs. The policy direction defined by (16) leaves transfers at all actually held balances unchanged but raises the marginal transfer around type 1's balance. It therefore alters the incentive to acquire liquidity without creating a mechanical first-order redistribution at the stationary support.

The reason this marginal incentive correction raises welfare is that type 1 chooses liquidity to maximize the private objective in (5). In a DM match, an additional unit of liquidity raises total match surplus by $F_x(x_1, \eta_1)$, but the buyer receives only the Kalai share η_1 of that gain. Hence the private first-order condition values the marginal DM return at $\eta_1 F_x(x_1, \eta_1)$, whereas utilitarian welfare in (13) values the full increase $F_x(x_1, \eta_1)$.

Since $\eta_1 < 1$, the buyer holds too little liquidity relative to the welfare criterion on the binding branch. Raising the marginal transfer at x_1 induces a first-order increase in type 1's post-policy real balances; because $F_{xx}(x_1, \eta_1) < 0$, this response is finite and has the sign shown in (14). The resulting welfare gain is precisely the positive derivative in (17). Adding a neutral proportional component then restores feasibility and monotonicity without destroying this positive welfare derivative.

Proof of Proposition 9. Substitute $\kappa(e) = 0$, $e'(x_1) = 1$, and $e'(x_2) = 0$ into (15).

Strict positivity follows from $F_{xx}(x_1, \eta_1) < 0$. Lemma 2 and the Implementation Theorem then supply the resulting feasible policy and the strict welfare gain. □

Choosing the proportional component $\widehat{e} = \lambda p + e$ with $\lambda > \|e'\|_\infty$ makes the resulting feasible policy strictly increasing. Proposition 7 then gives unique truthful reporting, while Lemma 2 and the Implementation Theorem give feasibility and a strict welfare gain for all sufficiently small positive intensities.

5.5 Complete Application

Fix a benchmark primitive vector θ_0 and the associated policy direction e_{θ_0} defined by (16), and set $E = \text{span}\{e_{\theta_0}\}$. Proposition 6 verifies invariance under nominal rescaling, complete C^1 equilibrium reconstruction, global portfolio optimality, and exact proportional neutrality. Proposition 7 verifies dynamic truthful reporting. Proposition 8 derives the welfare derivative, and Proposition 9 gives $L_{\theta_0} e_{\theta_0} > 0$.

By continuity of $\theta \mapsto L_\theta e_{\theta_0}$, after shrinking to a neighborhood U of θ_0, $L_\theta e_{\theta_0} > 0$ for every $\theta \in U$. Thus the exceptional set $\{\theta \in U : L_\theta e = 0 \text{ for every } e \in E\}$ is empty. Assumption 5 is satisfied vacuously, and the Genericity Theorem applies with no local exceptional primitives. The policy direction in (16) is a bounded global C^2 function, so its restriction belongs to the abstract policy space $\mathcal{B} = C^2(X)$.

Theorem 3 (Complete Application)**.** *A nonempty open set of three-type Lagos–Wright primitives satisfies every hypothesis of the Implementation Theorem and the Genericity Theorem. It has a complete, regular anticipated permanent equilibrium with uniquely reconstructed values, distributions, nominal scales, and prices. For the resulting feasible policy, reporting is uniquely truthful; proportional policy is exactly neutral; and welfare satisfies* (15)*. At every maintained primitive, the policy direction in* (16) *satisfies* (17)*; locally, a fixed one-dimensional space of nonproportional policy directions has no exceptional primitive.*

Theorem 3 provides a complete application of the preceding abstract results to

the three-type Lagos–Wright economy. It shows that, although money balances are private and households respond optimally to policy by changing their portfolios, the monetary authority can introduce a small permanent transfer schedule contingent on reported balances that raises welfare. The conclusion accounts for the induced adjustment of money holdings and for the resulting changes in all future prices, nominal variables, and balance distributions. Moreover, the result does not depend on a single calibrated parameter vector: it holds for a nonempty open set of primitives, and the same policy direction continues to have a positive first-order welfare effect throughout a neighborhood. Hence the hypotheses and conclusions of the Implementation Theorem and the Genericity Theorem are derived from the primitives of the monetary economy rather than assumed in reduced form.

Proof of Theorem 3. Proposition 5 provides the nonempty open primitive region. Proposition 6 gives the complete regular permanent-policy branch, unique reconstruction of omitted values and distributions, global portfolio optimality, nominal-scale-indexed-ceiling compatibility, and exact proportional neutrality. Proposition 7 gives dynamic truthful reporting, which is unique for the strictly increasing resulting feasible policy. Proposition 8 proves the welfare derivative, and Proposition 9 gives the favorable nonproportional policy direction. Fixing $E = \operatorname{span}\{e_{\theta_0}\}$ and using continuity makes the exceptional set empty on a neighborhood of every maintained primitive, so the Implementation Theorem and the Genericity Theorem apply completely. □

Thus the three-type Lagos–Wright economy is a complete application: all equilibrium, incentive, neutrality, welfare-derivative, and rank requirements are proved from primitives, and the resulting feasible policy is welfare improving on a nonempty open region.

6 Conclusion

This paper establishes a local Wallace theorem for anticipated permanent money creation in pure-currency economies with private balances. The argument has two parts. The Implementation Theorem starts from any policy direction with a positive first-order welfare effect. Although that direction may prescribe negative transfers or fail monotonicity, adding a sufficiently large proportional component leaves its welfare derivative unchanged while making the resulting policy nonnegative and strictly increasing. The resulting permanent rule is therefore feasible and truthfully implemented even when households can conceal part of their money holdings.

The Genericity Theorem explains why a favorable direction normally exists within a finite-dimensional family of economically chosen policy directions. Under a full-rank condition with respect to the primitives, all selected welfare derivatives vanish simultaneously only on a lower-dimensional exceptional set. Outside that set, at least one direction is favorable after choosing its sign, and the Implementation Theorem converts it into a feasible welfare-improving policy.

The three-type Lagos–Wright economy verifies this logic from primitives. It establishes the local anticipated equilibrium branch, exact neutrality of proportional transfers, the evolution of values and balance distributions, truthful reporting, and the welfare derivative that incorporates endogenous portfolio adjustment. The policy direction in (16) has a strictly positive welfare derivative throughout a nonempty open set of primitives. Hence the application is not merely illustrative: it provides a complete monetary economy in which the assumptions and conclusions of both the Implementation Theorem and the Genericity Theorem can be checked directly.

The results are local. They apply to sufficiently small permanent policies on the selected valued-money equilibrium branch and do not rank large interventions or distant equilibria. Locality, however, does not make the policy experiment static. Households choose date-0 portfolios in anticipation of a rule that remains in force permanently, and the welfare derivative incorporates the induced changes in all future

allocations, values, prices, nominal scales, and balance distributions.

A Mathematical Results Used

This appendix records the two standard mathematical results invoked in the main text, in the forms needed here.

Theorem A.1 (Banach-space implicit function theorem)**.** *Let X, Y, and Z be Banach spaces, let $U \subset X \times Y$ be open, and let $F : U \to Z$ be C^1. Suppose that $F(x_0, y_0) = 0$ and that $D_y F(x_0, y_0) : Y \to Z$ is a bounded linear isomorphism. Then there exist neighborhoods U_X of x_0 and U_Y of y_0, and a unique C^1 map $\varphi : U_X \to U_Y$, such that $F(x, \varphi(x)) = 0$ for every $x \in U_X$, and every zero of F in $U_X \times U_Y$ is of this form. Moreover, $D\varphi(x) = -\big[D_y F(x, \varphi(x))\big]^{-1} D_x F(x, \varphi(x))$.*

A standard Banach-space formulation is given by Deimling (1985, Chapter 4). In Section 2.3, Theorem A.1 is applied with the policy and primitive variables as x, the equilibrium object as y, and the equilibrium residual map $\mathcal{F}$ as F. It also justifies the finite-dimensional continuation argument in Proposition 6.

Theorem A.2 (Regular-value theorem)**.** *Let M be a C^r manifold of dimension n, let $f : M \to \mathbb{R}^q$ be C^r, where $r \geq 1$, and suppose that 0 is a regular value of f; that is, $Df(x) : T_x M \to \mathbb{R}^q$ is surjective at every $x \in f^{-1}(0)$. Then $f^{-1}(0)$ is empty or a C^r embedded submanifold of M of codimension q. In particular, if it is nonempty, then $q \leq n$.*

This is the preimage, or regular-value, theorem; see Hirsch (1976, Chapter 1). The proof of the Genericity Theorem applies Theorem A.2 to $f = G_E$, $M = \Theta_0$, and $q = \ell$. A positive-codimension embedded submanifold of the Euclidean open set Θ_0 has empty relative interior and Lebesgue measure zero, which yields the nowhere-density and nullity conclusions used there.

References

Amendola, Nicola, Leo Ferraris, and Fabrizio Mattesini. 2021. "Liquidity Externalities and the Wallace Conjecture." *Macroeconomic Dynamics* 25 (6): 1578–1588. https://doi.org/10.1017/S1365100519000798.

Bajaj, Ayushi, Tai-Wei Hu, Guillaume Rocheteau, and Mario Rafael Silva. 2017. "Decentralizing Constrained-Efficient Allocations in the Lagos–Wright Pure Currency Economy." *Journal of Economic Theory* 167:1–13. https://doi.org/10.1016/j.jet.2016.09.011.

Berentsen, Aleksander, Gabriele Camera, and Christopher J. Waller. 2005. "The Distribution of Money Balances and the Nonneutrality of Money." *International Economic Review* 46 (2): 465–487. https://doi.org/10.1111/j.1468-2354.2005.00329.x.

Deimling, Klaus. 1985. *Nonlinear Functional Analysis.* Berlin: Springer-Verlag. https://doi.org/10.1007/978-3-662-00547-7.

Deviatov, Alexei. 2006. "Money Creation in a Random Matching Model." *The B.E. Journal of Macroeconomics* 6 (3): 1–20. https://doi.org/10.2202/1534-5998.1499.

Deviatov, Alexei, and Neil Wallace. 2001. "Another Example in Which Lump-Sum Money Creation Is Beneficial." *Advances in Macroeconomics* 1 (1): 1. https://doi.org/10.2202/1534-6013.1001.

Fukai, Hiroki. 2026. "The Wallace Conjecture: Welfare Gains from Monetary Transfers." Unpublished manuscript, July 29.

Hirsch, Morris W. 1976. *Differential Topology.* Graduate Texts in Mathematics 33. New York: Springer-Verlag. https://doi.org/10.1007/978-1-4684-9449-5.

Kalai, Ehud. 1977. "Proportional Solutions to Bargaining Situations: Interpersonal Utility Comparisons." *Econometrica* 45 (7): 1623–1630. https://doi.org/10.2307/1913954.

Kiyotaki, Nobuhiro, and Randall Wright. 1989. "On Money as a Medium of Exchange." *Journal of Political Economy* 97 (4): 927–954. https://doi.org/10.1086/261634.

———. 1993. "A Search-Theoretic Approach to Monetary Economics." *American Economic Review* 83 (1): 63–77.

Lagos, Ricardo, and Randall Wright. 2005. "A Unified Framework for Monetary Theory and Policy Analysis." *Journal of Political Economy* 113 (3): 463–484. https://doi.org/10.1086/429804.

Molico, Miguel. 2006. "The Distribution of Money and Prices in Search Equilibrium." *International Economic Review* 47 (3): 701–722. https://doi.org/10.1111/j.1468-2354.2006.00393.x.

Nozawa, Wataru, and Hoonsik Yang. 2021. "Optimal Intervention in a Random-Matching Model of Money." *Bulletin of Economic Research* 73 (4): 721–735. https://doi.org/10.1111/boer.12270.

Rocheteau, Guillaume, Pierre-Olivier Weill, and Tsz-Nga Wong. 2018. "A Tractable Model of Monetary Exchange with Ex Post Heterogeneity." *Theoretical Economics* 13 (3): 1369–1423. https://doi.org/10.3982/TE2821.

Rocheteau, Guillaume, and Randall Wright. 2005. "Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium." *Econometrica* 73 (1): 175–202. https://doi.org/10.1111/j.1468-0262.2005.00568.x.

Trejos, Alberto, and Randall Wright. 1995. "Search, Bargaining, Money, and Prices." *Journal of Political Economy* 103 (1): 118–141. https://doi.org/10.1086/261978.

Wallace, Neil. 1981. "A Modigliani–Miller Theorem for Open-Market Operations." *American Economic Review* 71 (3): 267–274.

———. 2014. "Optimal Money Creation in "Pure Currency" Economies: A Conjecture." *The Quarterly Journal of Economics* 129 (1): 259–274. https://doi.org/10.1093/qje/qjt030.

Zhou, Ruilin. 1999. "Individual and Aggregate Real Balances in a Random-Matching Model." *International Economic Review* 40 (4): 1009–1038. https://doi.org/10.1111/1468-2354.00051.