



Munich Personal RePEc Archive

Minimum viable relationships (MVR): a relational readiness framework for African ventures

Mukiibi, Farouk Mark

Africa Market OS

15 September 2025

Online at <https://mpra.ub.uni-muenchen.de/126190/>
MPRA Paper No. 126190, posted 13 Oct 2025 23:32 UTC

Title: Minimum Viable Relationships (MVR): A Relational Readiness Framework for African Ventures

Author: Farouk Mark Mukiibi (ORCID: 0009-0009-8191-2098)

Affiliation

Independent Researcher; African Market OS (AMOS); Gatsby Marketing Agency UG Ltd.,
Kampala, Uganda

Correspondence

farouk.mark@gatsbymarketingug.com

Article type

Theoretical/Conceptual

License

CC BY 4.0

Version & Date

Version 2.0 — 15 September 2025

Keywords

relational readiness; embeddedness; social capital; diffusion of innovations; MVP; informal economy; Africa; market entry

Abstract

Startups that import MVP-first logics into high-context African markets frequently stall not because products lack utility, but because ventures lack permission to operate. This paper proposes Minimum Viable Relationships (MVR) as a relational-readiness gate that precedes MVP in such contexts. MVR formalizes the conditions under which counterparties (customers, distributors, institutions) grant access without triggering social or organizational sanction. The paper defines the construct, situates it against adjacent ideas (design thinking, diffusion of innovations), and argues that where social sanction outweighs functional risk, relationship viability must be validated before product experiments can be considered valid.

The paper contributes three artifacts: (1) a seven-dimensional MVR diagnostic with go/no-go thresholds that scores embeddedness, trust, guardian vouches, and channel permission; (2) an MVR Investment Memo template that enables funders to assess relational risk alongside

financial and operational risk; and (3) a practical design-lab of field tools for earning permission (pilot-site commitments, data-sharing agreements, compliance pathways, and referral loops). Comparative case vignettes (e.g., SafeBoda, KOKO Networks, Wave versus Dash, Sendy) illustrate how early permission artifacts are associated with durable traction.

The framework reframes early venture due diligence in Africa from “Can the product work?” to “Are we allowed to make it work here?” and offers testable propositions for founders, investors, and ecosystem builders. Limitations and avenues for empirical study are discussed.

1. Introduction

Lean-startup methods teach founders to learn quickly from the market by releasing a minimum viable product (MVP) and iterating (Ries, 2011; Blank, 2013). In many African contexts, however, seemingly “validated” pilots stall when they leave founder-managed environments. Teams report enthusiastic feedback during controlled trials, yet encounter blocked access, tepid referrals, and fragile traction once they attempt real distribution. These outcomes are often treated as product problems. I argue instead that they are frequently problems of **relational unreadiness** in **high-context** markets—settings in which coordination and sanction are enacted through dense social networks rather than through formal contracts alone (Granovetter, 1985; Fafchamps, 2004; Meagher, 2010; Hall, 1976).

I introduce **Minimum Viable Relationships (MVR)** to name and operationalize this missing antecedent to product testing. MVR is the minimally sufficient configuration of trust, embeddedness, and permission-to-operate that must be in place for early venture engagement to be legitimate and feasible. The emphasis on *permission* is deliberate: in high-context markets, gatekeeper endorsements, customary approvals, and reputational standing function as preconditions to participation. Without them, the same MVP that “works” in a founder’s pilot cannot circulate across real channels because it lacks social authorization. MVR therefore does not compete with MVP; it **precedes and scaffolds** it. The practical sequence is **MVR → MVP → PMF**.

This argument is grounded in well-established literatures. Work on embeddedness and social capital shows that economic action is situated within relationships, norms, and obligations that structure opportunity and risk (Granovetter, 1985; Coleman, 1988; Putnam, 2000; Uzzi, 1997). The relational view of the firm explains how interorganizational rents arise from relationship-specific assets, partner safeguards, and reputation (Dyer & Singh, 1998), while structural-holes theory clarifies how position in networks shapes reach and information advantages (Burt,

1992). Diffusion research emphasizes the role of opinion leaders and perceived legitimacy in accelerating adoption (Rogers, 2003). In African informal economies, these dynamics are amplified by limited formal enforcement and strong local sanctioning systems; community guardians and channel hosts can enable or block access at low cost (Fafchamps, 2004; Meagher, 2010). Taken together, these strands predict that ventures without minimal relational standing will generate ambiguous MVP signals: low uptake that reflects social veto rather than product unfitness.

The prevailing MVP-only approach produces a **measurement error**. When ventures test products without first earning sufficient endorsements and embedded access, negative or noisy results are easily misread as evidence that features, pricing, or messaging are at fault. Founders then “optimize” product attributes while the binding constraint remains relational. Investors and accelerators face a parallel misdiagnosis: diligence keyed to product-market fit overlooks **relational legitimacy**, a precursor to reliable acquisition, retention, and referral in trust-sensitive markets. The consequence is avoidable churn, inflated customer-acquisition costs, and discontinuous growth.

This paper makes four contributions. First, it **defines** MVR as a construct distinct from generic “relationships” or relationship marketing: MVR denotes a structural state of *minimal sufficiency* in trust, embeddedness, and permission within the specific networks a venture must traverse. Second, it **specifies** seven dimensions that make up MVR and links them to observable field indicators. Third, it **derives propositions** connecting MVR to early adoption, conversion, resilience, and growth, clarifying when and why an **MVR → MVP** sequence should outperform MVP-first approaches. Fourth, it **outlines measurement** via an MVR Index (0–100) and proposes mixed-methods designs—cross-sectional, experimental, and retrodictive—to validate the framework.

Scope and terminology warrant brief clarification. “African ventures” here includes both formal startups and informal enterprises operating in sectors where reputation, endorsements, and customary approvals are salient. “High context” refers to environments in which meaning and sanction are conveyed through shared histories, roles, and social cues, such that third-party vouching and fit with local scripts strongly condition cooperation (Hall, 1976). The intention is not to essentialize “Africa,” but to analyze contexts—common across many African markets—in which **social sanction can outweigh functional risk**. While examples and measures are drawn from these settings, the theory should travel to other high-context environments with appropriate calibration.

The remainder proceeds as follows. Section 2 situates MVR in related literatures and articulates the specific gap the framework addresses. Section 3 develops the construct, detailing its seven

dimensions and distinguishing it from adjacent concepts. Section 4 translates the construct into field-measurable indicators and an index. Section 5 states propositions linking MVR to adoption, retention, and growth. Section 6 outlines measurement strategies and research designs for validation. Section 7 discusses implications for founders, investors, and accelerators, and Section 8 notes limitations and directions for future work. Section 9 concludes.

2. Background and Gap

A large body of social science treats economic action as situated in networks of ongoing relations rather than in atomized exchanges. In the **embeddedness** tradition, transactions are made possible and governable because actors rely on reputation, norms of reciprocity, and third-party monitoring carried by social ties (Granovetter, 1985). **Social capital** denotes the collectively produced resources—trust, obligations, information channels—that arise from these ties and lower the cost of coordination (Coleman, 1988; Putnam, 2000). For early ventures that lack track records, credible endorsements and repeated interactions can substitute for formal safeguards, enabling access that is otherwise prohibitively expensive.

Interorganizational research extends this logic to firm–partner relations. The **relational view** argues that durable advantage often comes from relationship-specific assets, complementary capabilities, and governance safeguards that reduce opportunism between partners (Dyer & Singh, 1998). Network structure also matters: actors positioned to **bridge structural holes** can reach diverse communities and accelerate information flow, but they still depend on localized legitimacy to convert reach into cooperation (Burt, 1992). Over-embeddedness, conversely, can produce lock-in and reduce adaptability when ties are too dense or insular (Uzzi, 1997). These findings collectively imply that the returns to a product or service are conditioned by where the venture sits in relevant networks and how it is perceived by influential others.

Adoption research underscores the same point from the demand side. The **diffusion of innovations** literature shows that opinion leaders, role models, and perceived legitimacy shape trial and spread: people learn what to adopt not only from instrumental payoffs but from who endorses an option and whether it fits prevailing norms (Rogers, 2003). In **high-context** environments—where meaning and sanction travel through shared histories, roles, and implicit cues—gatekeeper approval and culturally fluent presentation carry greater weight than in low-context settings (Hall, 1976). In many African markets, where **informal institutions** and community enforcement remain salient, customary permissions, group administrators, landlords, or trade associations may enable or veto access at low cost; dispute handling often follows recognized local fora and timetables (Fafchamps, 2004; Meagher, 2010). As a result,

ventures that are not yet recognized as legitimate participants can encounter hidden barriers even when their offers are functionally sound.

Despite this accumulated evidence, mainstream **MVP-first** practice treats relationships as a channel choice rather than as a **precondition**. The usual method asks teams to place a minimal product with early users, collect feedback, and iterate features and pricing. In high-context settings, however, such tests are easily confounded: pilots run inside founder-managed spaces or with friendly hosts generate encouraging signals, while broader trials falter because the venture lacks endorsements, embedded distribution, or dispute-resolution credibility. Teams then attribute the shortfall to product attributes and continue optimizing features, while the binding constraint remains relational. Investors and accelerators face a parallel diagnostic error when diligence keyed to product-market fit overlooks **relational legitimacy**—the permission and embedded access required for acquisition, retention, and referral to be reliable.

What is missing is a construct that names and measures the **minimal relational state** required for early market learning to be valid. Relationship marketing and partnership playbooks do not fully address this need: they focus on persuasion and campaigns rather than on whether a venture has crossed a community's threshold of acceptance. I label this threshold **Minimum Viable Relationships (MVR)** and propose it as a gating milestone that precedes MVP in trust-sensitive markets. By formalizing MVR—specifying its dimensions, mapping it to observable field indicators, and stating propositions about outcomes—we can separate relational veto from product failure and design studies that do not conflate the two. The next section develops the construct in detail.

3. Conceptualizing Minimum Viable Relationships (MVR)

3.1 Construct definition

I define **Minimum Viable Relationships (MVR)** as the *minimally sufficient configuration of trust, embeddedness, and permission-to-operate within the specific networks a venture must traverse* for early engagement to be both legitimate and feasible. MVR is a **state of the venture's position** in relevant social systems, not a marketing tactic. Its function is epistemic and practical: by crossing a threshold of relational legitimacy, a venture creates conditions under which MVP tests will yield interpretable signals rather than artifacts of social veto. In high-context markets, this threshold is achieved when credible third parties can vouch for the venture, when the venture is recognizably situated in existing channels and routines, and when dispute handling is expected to be swift and norm-concordant.

3.2 Boundary conditions and distinctions

MVR is adjacent to, but distinct from, several familiar ideas. It is **not** simply “relationship marketing,” which focuses on persuasion and post-acquisition retention campaigns; MVR concerns *ex ante* legitimacy and access. It is broader than “partnerships,” which can exist without conferring permission in the eyes of local guardians. It is narrower than “social capital” in general; MVR captures the *least* relational sufficiency needed for *this venture in this market* to operate without hidden blockages. Nor is MVR a synonym for “brand trust”: in many African contexts, localized endorsements, role fit, and credible dispute mechanisms outweigh brand familiarity. Finally, MVR is not a guarantee of success. Over-embeddedness can reduce adaptability, and high MVR may still fail if the underlying offer lacks value. The claim is conditional: where social sanction carries weight, **MVR typically precedes—rather than follows—MVP** in the order of operations.

3.3 The seven dimensions of MVR (substantive content in prose)

Gatekeeper inclusion. Ventures attain minimal legitimacy when identifiable gatekeepers—respected group administrators, association leaders, landlords, or channel hosts—can be named, reached, and are willing to endorse or host the venture. Inclusion is substantive, not symbolic: it entails a willingness by these actors to attach their reputational capital to the venture in bounded ways (e.g., hosting, co-signing access, introducing to their members).

Trust-signal density. People rely on cues to infer reliability when formal track records are thin. Trust-signal density denotes the concentration of credible third-party signals around the venture—public vouches, cross-member references, reciprocity histories—that jointly reduce perceived opportunism. A single testimonial rarely suffices; density matters because it reflects convergent evidence from multiple ties.

Reciprocity history. Before asking communities to transact, ventures that have given first—through pro-bono support, contributions to shared resources, or responsive service during an earlier phase—establish a visible record of reciprocity. Such histories create obligations and expectations that future exchanges will be fair. In practice this can include small, repeated acts (e.g., reliably topping up float, assisting a host group) that signal presence rather than extractiveness.

Cultural fluency. High-context environments encode meaning in etiquette, timing, and language choice. Cultural fluency is the venture’s demonstrated competence in these codes—how it addresses elders, who speaks in which setting, how consent is sought, and how sensitive

topics are framed. Fluency minimizes accidental disrespect, increases the intelligibility of the offer, and smooths conflict resolution when frictions arise.

Reputational capital. Separate from immediate vouchers, ventures accumulate a bank of verifiable references and success stories within the target network or its close neighbors. This capital is specific: what matters is evidence legible to the *same* or adjacent communities (e.g., a cooperative's secretary, a known aggregator), not distant accolades. Reputational capital makes endorsements portable across groups because it can be checked informally.

Informal-contract reliability. Many exchanges proceed on promises, not paper. Informal-contract reliability captures whether the venture keeps time-bound commitments and resolves disputes quickly in ways recognized as fair. Two parameters are central: *speed to closure* (aiming for hours or days, not weeks) and *recurrence* (whether issues resurface). Reliability on these parameters lowers the expected enforcement cost for partners and customers.

Community role fit. Entrants that map themselves onto familiar roles—agent, steward, host, facilitator—are easier to accept than those that arrive as category errors. Role fit does not mean mimicry; it means the venture's behavior, visuals, and touchpoints align with recognizable scripts so that others know how to interact with it and what obligations to expect. Role fit reduces ambiguity, which in turn reduces social risk for adopters and hosts.

3.4 Integrative logic

These dimensions cohere because they jointly address the three mechanisms that block participation in trust-sensitive systems: lack of authorization (gatekeeper inclusion and role fit), fear of opportunism (trust-signal density, reputational capital, informal-contract reliability), and misalignment with local scripts (cultural fluency and reciprocity history). When each mechanism is satisfied to a minimal degree, a venture crosses the MVR threshold. Below that threshold, market responses to an MVP are confounded: weak uptake may reflect social sanction rather than product value. Above it, MVP results more closely track functional fit and pricing because permission and access have ceased to be the binding constraints.

3.5 Relation to Diffusion, Design Thinking, and Co-creation

It is useful to situate MVR alongside adjacent approaches that might appear similar but ask different questions.

Diffusion of Innovations (Rogers, 2003). Diffusion theory presumes a candidate product exists and analyzes how quickly it spreads through a social system given attributes such as relative

advantage, observability, and the role of opinion leaders. By contrast, **MVR precedes the diffusion problem**: it asks whether the entrant has attained *minimal relational sufficiency*—authorization by gatekeepers, embedded access, dispute credibility—without which early trials are vetoed or confounded. Put differently, diffusion models *velocity conditional on admissibility*; MVR models the *admissibility* itself. The temporal relation is sequential: **MVR → MVP → (if value is present) Diffusion**. When MVR is below threshold, negative MVP signals are not diagnostic of poor value; they reflect social sanction rather than lack of functional fit.

Design Thinking (Brown, 2008; Liedtka, 2018). Design Thinking is a sense-making methodology that yields empathy, insight, and well-framed problems. It is orthogonal to MVR: one can conduct excellent human-centered research and still be denied meaningful access by guardians, hosts, or community norms. In trust-sensitive settings, **empathy without permission returns research, not access**. MVR therefore complements design: Design Thinking pinpoints what to build; MVR determines whether one is *allowed* to build and test in situ.

Co-creation (Prahalad & Ramaswamy, 2004). Co-creation describes value formation with stakeholders during design and delivery. Co-creation can *feed* MVR (for example, by producing reciprocity histories and trust signals), but it is not a substitute. One can co-create with small groups under founder protection and still lack **guardian inclusion, embedded distribution, or dispute credibility** at the community level. MVR is an entry governance: it gauges whether legitimacy translates into **repeatable motion** across counterparties rather than into one-off pilots.

Three distinctions make these boundaries precise. **Timing**: Diffusion and co-creation typically assume a candidate artifact; MVR is **pre-artifact governance** in high-context markets. **Unit of analysis**: Diffusion tracks adoption across populations; MVR operates at the level of **dyads and meso-institutions** (guardian–venture, host–venture, forum–venture). **Risk model**: Diffusion models **adoption velocity**; MVR models **sanction risk**—the permission, belonging, and obligation thresholds that determine whether trials are legitimate and interpretable. In short, diffusion explains speed, design explains fit, **MVR secures the visa**. Without the visa, an MVP is an **invalid experiment** because the social system has not admitted the entrant.

3.6 Discriminant Predictions and Falsifiability

To avoid tautology and sharpen adjudication, MVR yields **discriminant predictions** relative to adjacent frameworks.

1. **Admissibility vs. Velocity**. Conditional on comparable product attributes and spend, ventures with high MVR should exhibit **measurably higher rates of valid trial** (i.e., trials

occurring through authorized hosts with verified dispute coverage) than ventures with low MVR. By contrast, **conditional on admissibility**, diffusion variables (e.g., opinion-leader centrality) should better explain **spread** than **trial existence**.

2. **Orthogonality to Empathy Quality.** Holding constant independently rated Design-Thinking quality (e.g., depth of insight, prototype usability), **MVR indicators (GV, ED, DD) should predict admission to pilot sites and retention at launch.** If empathy alone were sufficient, MVR would add little explanatory power. Conversely, if MVR were merely “good design in disguise,” its effects would vanish when empathy is controlled. The prediction is that **both** matter, but on distinct axes.
3. **Beyond Co-creation Intensity.** After controlling for co-creation intensity (e.g., number of co-design sessions, stakeholder hours), **MVR should still predict embedded access and trusted referrals.** Co-creation can raise trust-signal density, but without **gatekeeper inclusion** and **host activation**, propagation stalls.
4. **Pre-registration of Fail Conditions.** The framework is falsifiable. It would be weakened if (a) MVP-first ventures matched to MVR→MVP ventures showed **no retention advantage** for the latter in high-context markets, or (b) MVR indicators failed to predict admissibility while diffusion or design measures did. It would be refuted if **high MVR consistently failed to raise the probability of valid trials** in settings where social sanction is demonstrably salient.

Methodologically, discriminant validity can be tested in a multi-model design that includes **MVR indicators, Design-Thinking quality scores, and diffusion variables** (e.g., opinion-leader degree centrality, observability indices). Explanatory power should partition as follows: **MVR → admissibility and early retention, diffusion → spread conditional on admissibility, design → fit and satisfaction.** Pre-registered hierarchical models, with mediation and moderation specified ex ante, make such partitions adjudicable.

MVR Index construction (concise specification)

Let the five indicators be GV, ED, RM, ME, and DD. For each indicator, compute a robust standardized score within a sector–city reference class using median and MAD (median absolute deviation). Combine the five standardized values with the baseline weights (GV 0.25, ED 0.25, RM 0.20, ME 0.15, DD 0.15) to form a composite, then linearly rescale to 0–100 so that 50 equals the reference median. Always report the composite and the five components with IQRs, and run sensitivity checks with (a) equal weights and (b) data-driven loadings (PCA or confirmatory factors). Treat structurally missing components neutrally and flag measurement-process missingness.

3.7 Clarifying terms: legitimacy, permission, and repeatability

In high-context systems, three notions that often travel together must be pried apart to avoid equivocation. **Legitimacy** is an audience judgment that the entrant is appropriate for the role it claims—“this actor belongs here.” Legitimacy is necessary but not sufficient for operation.

Permission is an actionable authorization from actors with the power to enable or veto access (guardians, hosts, administrators, compliance officers). Permission is *situated* and *revocable*; it is demonstrated by artifacts such as hosted posts, letters of introduction, slotting on a shelf, or the opening of a group workflow. **Repeatability** is the ability for early actions to be reproduced across counterparties without collapsing into one-off heroics; it is revealed by durable referral momentum, routine entrenchment, and dispute closure that does not require founder escalation. **MVR** concerns the *minimal joint presence* of these three properties in the relevant network. A venture can enjoy diffuse legitimacy yet lack permission; it can enjoy one-time permission yet lack repeatability. By requiring observable vouchers, embedded access, routinized ties, and enforceable closure, MVR operationalizes the threshold at which legitimacy becomes motion.

4. Observable Indicators and the MVR Index

Translating the seven conceptual dimensions into field measurement requires indicators that are (a) observable under real venture conditions, (b) comparable across teams and sites, and (c) resistant to trivial gaming. I operationalize MVR with five indicators that jointly capture authorization, embedded access, social propagation, routine integration, and enforcement credibility. Each indicator is defined, tied to concrete data sources, and accompanied by reliability checks.

Guardian vouchers (GV). Guardian vouchers capture whether identifiable gatekeepers endorse or host the venture in bounded, verifiable ways. In practice, guardians may be cooperative officials, market or building managers, group administrators, landlords, parish or ward chairs, or trade-association officers. Measurement proceeds by enumerating unique guardians who have provided explicit endorsement (e.g., signed statements, on-record audio, hosted posts in member groups, or co-branded notices) and weighting each by local credibility and reach. To reduce inflation, endorsements are validated by spot-calling, link verification, or attendance records for hosted sessions. Because legitimacy is localized, vouchers must be legible within the *target* network rather than being generic testimonials from distant actors.

Embedded distribution (ED). Embedded distribution reflects the proportion of customers acquired or fulfilled through existing, trusted channels rather than cold, firm-owned funnels. Typical examples include co-selling inside SACCOs and producer cooperatives; placing an offer with anchor tenants or established merchants; piggy-backing on agent networks; or being hosted in WhatsApp/Telegram groups with active admins. ED is measured as the share of new customers or transactions attributable to embedded hosts over a defined window (e.g., 14–30 days), using referral codes, admin-post logs, or POS tags. Reliability checks include reconciling self-reports against host-side records and auditing a random subset of transactions for provenance.

Referral momentum (RM). Referral momentum is the pace at which the venture propagates through peer channels absent heavy spend. The core statistic is the median number of days per net-new customer acquired through verified referrals (“whisper velocity”), with secondary measures such as referral-to-paid conversion. RM is computed from instrumented referral links or codes and from admin push logs in hosted groups. Because early bursts can be noisy, momentum is reported with robust statistics (median and interquartile range) and with spend-normalized ratios (referred vs. paid-acquisition shares). Where feasible, outlier spikes driven by one-off events are flagged rather than folded into trend estimates.

Micro-entrenchment (ME). Micro-entrenchment assesses whether the venture is woven into everyday routines in ways that are inexpensive to maintain yet costly to abandon. Typical mechanisms include float accounts or pre-positioned inventory, shared ledgers or standing orders, pinned workflow groups with active admins, and light integrations with a host’s existing processes. ME is recorded as the count of distinct routine-binding mechanisms in live use, multiplied by their observed usage frequency over the measurement window. Audits verify that mechanisms are operational (not merely set up) and that use is attributable to the venture’s activity rather than generic background traffic.

Dispute dynamics (DD). Dispute dynamics quantify the venture’s promise-keeping and enforcement credibility. The key variables are median time-to-closure for disputes and the 30-day recurrence rate, coded alongside whether resolution followed locally recognized fora and scripts. Evidence sources include ticketing systems, WhatsApp thread exports with admin notes, and simple closure logs co-signed by counterparties. Because speed can be gamed by closing disputes prematurely, recurrence is paired with time-to-closure in a joint summary; unusually low times alongside high recurrence trigger review.

Constructing a single index from these indicators is useful for screening, sequencing decisions, and empirical work. I define the **MVR Index** as a standardized, weighted composite that maps a venture’s indicator profile onto a 0–100 scale. Construction proceeds in four steps. First, each

raw indicator is normalized within a relevant reference class—typically sector × city or region—using robust standardization (median–MAD or winsorized z-scores) to minimize leverage from extremes. Second, missingness is handled conservatively: if an indicator is missing for structural reasons (e.g., no disputes occurred), a neutral value is imputed and flagged; if it is missing because measurement was not attempted, the observation is retained but down-weighted in sensitivity analyses. Third, standardized indicators are combined by a vector of theory-informed weights that reflect their distinct roles: in the baseline specification, **GV** and **ED** receive higher weights because they most directly encode authorization and access, while **RM**, **ME**, and **DD** capture propagation, stickiness, and enforcement (illustratively, 0.25, 0.25, 0.20, 0.15, 0.15, respectively). Fourth, the weighted sum is linearly rescaled to 0–100 for interpretability, with 50 corresponding to the median of the reference class.

Two implementation details enhance credibility. **Auditability** is preserved by retaining source artifacts (endorsement links, host records, referral exports, workflow screenshots, dispute logs) so that third parties can verify claims without relying on self-report alone. **Robustness** is assessed by re-estimating the index with alternative weightings (equal weights; data-driven weights from principal components or confirmatory factor models) and by checking that substantive conclusions—particularly go/no-go decisions and treatment effects—do not hinge on a single indicator. Because the relational structure of markets varies, **context calibration** is explicit: reference classes and, where necessary, weights are re-estimated when a venture shifts sector or city, ensuring that an MVR score of, say, 65 in Kampala retail is comparable to 65 in Nairobi mobility only after appropriate normalization.

While the index is convenient for dashboards and models, practitioners should resist reducing MVR to a single number. In sequencing decisions, **profile shape** matters: a venture with strong vouchers and embedded access but weak dispute dynamics faces different risks than one with high momentum but no recognized guardians. For that reason, the index is always reported alongside the five underlying indicators and a short narrative interpreting the pattern in light of local gatekeepers, hosts, and norms.

5. Propositions

The framework yields a set of testable propositions that locate MVR’s effects in well-specified outcomes and mechanisms. Throughout, outcomes can be defined at the venture–period level (e.g., weekly or monthly windows), with standard controls for price, product quality proxies, marketing spend, and seasonality. Where noted, mechanisms such as perceived legitimacy, perceived risk, and perceived norms can be measured with short validated scales or behavioral

proxies; market “high-contextness” can be captured with a composite (share of informal transactions, reliance on group administrators, salience of customary dispute fora).

Proposition 1 (Legitimacy). Ventures with higher guardian vouchers (GV) will exhibit higher early adoption—measured as sign-ups, trials, or first purchases—conditional on price and product features. The mechanism is perceived relational legitimacy: credible third-party vouching lowers the social cost of trying a new entrant. A mediation test can use a brief legitimacy scale administered to prospects or observational proxies such as the rate at which invited members act on admin posts relative to non-admin posts.

Proposition 2 (Access). A higher share of embedded distribution (ED) will predict a greater conversion from pilot to paid usage, net of marketing spend. The mechanism is frictionless access through trusted channels, which lowers search and enforcement costs. Identification can leverage within-venture variation when a new host channel is activated, comparing pre/post conversion while holding offer and price constant.

Proposition 3 (Propagation). Faster referral momentum (RM)—shorter median days per net-new referred customer—will predict a higher proportion of organic growth and a lower cost of customer acquisition (CAC). The mechanism is descriptive norm formation (“people like me are using this”), which increases the credibility of referral invitations. A mediation design can pair RM with a short perceived-norms scale and test indirect effects on growth and CAC.

Proposition 4 (Resilience). Greater micro-entrenchment (ME)—live, routine-binding mechanisms such as float accounts, standing orders, or shared ledgers—will predict lower churn and faster usage recovery following operational shocks (e.g., stockouts, price changes). The mechanism is switching and coordination cost: embedded routines make abandonment costlier and resumption easier. Event-study designs around shocks can estimate ME’s buffering effect.

Proposition 5 (Repair). More favorable dispute dynamics (DD)—shorter median time-to-closure with low 30-day recurrence and norm-concordant resolution—will predict higher customer lifetime value (LTV) and referral rates among affected customers. The mechanism is trust repair through visible, fair enforcement. Analyses can compare cohorts with versus without disputes, matched on baseline value, and estimate the differential in subsequent LTV and referrals.

Proposition 6 (Sequencing). Ventures that sequence **MVR→MVP** will be more likely to meet predefined retention thresholds at launch than otherwise similar ventures that run MVP first. The mechanism is elimination of relational veto prior to product testing, yielding interpretable MVP signals. A field experiment can randomize ventures (or local rollouts within a venture) to

run a two-week MVR Sprint before MVP versus MVP-first, with retention at 30/60 days as the primary endpoint.

Proposition 7 (Context moderation). The effects in Propositions 1–6 will be stronger as market high-contextness increases. Operationally, the interaction between each indicator (GV, ED, RM, ME, DD) and a high-contextness index should be positive when predicting adoption, conversion, resilience, and growth. Cross-city or cross-sector samples (e.g., Kampala vs. Nairobi; mobility vs. e-commerce) allow hierarchical models to estimate varying slopes by context.

Proposition 8 (Investor screening). A higher pre-MVP MVR Index will predict faster time-to-first-meaningful-revenue and lower early CAC relative to matched peers. The mechanism is that ventures with relational sufficiency face fewer hidden access constraints, allowing spend to translate efficiently into paying users. A matched cohort design (propensity scores on sector, team experience, and runway) can estimate the association; robustness checks include alternative index weightings and equal-weight composites.

Two identification concerns deserve note. First, reverse causality (e.g., growth attracting endorsements) can be mitigated by lagging MVR measures and exploiting discrete channel activations or endorsement events. Second, unobserved venture quality can be addressed with fixed effects in panel designs, matched samples, or—where feasible—instrumental variation such as exogenous gatekeeper turnover or randomized assignment of host groups. Across propositions, pre-registration of endpoints and analysis choices on OSF reduces researcher degrees of freedom and clarifies confirmatory versus exploratory components.

6. Measurement and Research Design

The empirical program proceeds on two parallel tracks: field-observable indicators for practice and modeling (guardian vouchers, embedded distribution, referral momentum, micro-entrenchment, dispute dynamics) and a psychometrically validated scale that captures the seven latent dimensions of MVR at the venture–market dyad level. Outcomes include early adoption (sign-ups, trials, first purchases), conversion from pilot to paid use, retention and churn, referral share, customer acquisition cost, and time-to-first-meaningful-revenue. Analyses use venture–period panels, with consistent windows (e.g., weekly in early stages, monthly thereafter), and include standard controls for price, offer breadth, spend, calendar effects, and sector–city fixed effects.

Indicator measurement privileges auditability over self-report. Guardian vouchers are enumerated from verifiable artifacts such as signed notices, on-record audio, hosted admin posts, or co-branded announcements, each tethered to specific gatekeepers whose legitimacy

is legible in the target network. Embedded distribution is computed as the proportion of new customers or fulfilled transactions attributable to identifiable hosts, established via referral codes, admin logs, or POS tags, and reconciled with host-side records on a random audit schedule. Referral momentum is derived from instrumented links or codes and group admin pushes, with the primary statistic defined as the median number of days per net-new referred customer; because early bursts can be volatile, the interquartile range is reported and marketing spend is used to normalize organic shares. Micro-entrenchment is recorded as the number of distinct routine-binding mechanisms in live use—float accounts, standing orders, shared ledgers, pinned workflow groups—multiplied by their observed usage frequency over the measurement window; verification requires evidence that mechanisms are active rather than merely configured. Dispute dynamics are coded from ticketing systems or message exports with admin notes, summarizing median time-to-closure, norm concordance of the forum used, and 30-day recurrence; unusually rapid closures paired with high recurrence are flagged for review to deter superficial “fast” resolutions.

To complement these behavioral indicators, a short relational scale will be developed to measure the seven MVR dimensions directly. Item generation draws on interviews with market administrators, cooperative leaders, group admins, and founders across Kampala, Nairobi, and Lagos to ensure content validity and local legibility. Items are phrased at the venture–market level (for example, “Recognized gatekeepers in this community would be willing to host or introduce us” for gatekeeper inclusion; “We have a visible history of giving before asking in this group” for reciprocity history; “If a dispute arose, there is a predictable, legitimate forum that would resolve it quickly” for informal-contract reliability) and use 5- or 7-point Likert response formats anchored in concrete behaviors. A pilot sample of approximately 120 ventures will be used for exploratory factor analysis and item reduction, targeting two to three items per dimension. A confirmatory sample of at least 300 venture–market observations will test a seven-factor model; reliability will be assessed via coefficient omega, convergent validity via average variance extracted, and discriminant validity via the Fornell–Larcker criterion and HTMT ratios. Measurement invariance across cities and sectors will be examined sequentially (configural, metric, scalar), and partial invariance will be documented where full invariance is not attainable. The behavioral indicators (GV, ED, RM, ME, DD) will be correlated with the latent scale to establish criterion validity and to test whether the composite MVR Index tracks the latent construct.

Three designs map onto the propositions. Design A is a mixed-methods field study. Ventures are recruited through accelerators, incubators, and trade associations in the three focal cities, with eligibility defined as active commercial activity and an intention to launch or extend into a defined community within the study horizon. After consent, teams provide baseline indicator

artifacts, complete the relational scale, and share outcomes via lightweight data-sharing agreements or structured exports (for example, anonymized referral logs). Outcomes are tracked for sixty and one hundred and twenty days, and qualitative interviews with gatekeepers and hosts are conducted to elaborate mechanisms such as how vouchers are earned or how disputes are actually handled. This design supports contemporaneous tests of associations and mediation (for perceived legitimacy, perceived risk, and perceived norms), and its qualitative arm provides construct clarity and helps interpret heterogeneous effects.

Design B is a field experiment to test sequencing. Within ventures planning multi-neighborhood or multi-group rollouts, local sites (or time blocks) are randomized to either run a two-week MVR sprint before MVP or proceed with MVP-first. The sprint is standardized in a replicable protocol: secure at least three gatekeeper vouchers with consented references, activate one embedded host channel, instrument referrals and achieve a targeted whisper velocity, entrench at least one routine-binding mechanism, and dry-run dispute resolution with a consenting counterparty. Primary endpoints are retention thresholds at 30 and 60 days; secondary endpoints include referral share, CAC, and time-to-first-meaningful-revenue. Randomization is blocked by sector and city to improve balance, and pre-specified covariate adjustment increases precision. Spillovers are minimized by separating treated and control sites geographically or by time. A pre-analysis plan registered on OSF specifies primary outcomes, estimands, and decision rules to limit researcher degrees of freedom.

Design C is a retrodictive multiple-case analysis. Past ventures are purposively sampled for maximum variation on outcomes (failures after promising pilots; resilient growers) within comparable sectors and cities. Archival materials, news coverage, and retrospective interviews are used to reconstruct indicator trajectories and to score the seven relational dimensions just prior to key launch moments. Two independent coders rate each dimension using explicit rubrics; intercoder agreement is assessed with weighted kappa, and disagreements are resolved by adjudication. The analysis tests whether pre-MVP MVR profiles differentiate subsequent success from failure when controlling for product sophistication, runway, and market timing. Triangulation across documentary, interview, and artifact evidence reduces recall bias.

Across designs, the analytic strategy follows a common structure. Indicator and outcome series are aligned into venture–period panels and summarized with robust statistics that limit leverage from extreme observations. Baseline models estimate associations between each indicator (and the composite index) and outcomes using fixed effects to absorb time-invariant venture characteristics and city-by-sector interactions to absorb stable context. Mediation is assessed with latent-variable path models or with two-stage procedures using short scales for perceived legitimacy, risk, and norms. Moderation by high-contextness is estimated with

interactions between indicators and a context index constructed from external data (for example, reliance on group administrators, prevalence of informal dispute fora, and the share of transactions occurring outside formal contracts). For Design B, intent-to-treat effects and, where compliance permits, complier average causal effects are reported, with randomization inference as a robustness check. Multiple-comparison adjustments are applied within families of related outcomes, and sensitivity analyses include equal-weight composites, alternative normalizations, leave-one-indicator-out indices, and re-weighting by principal components.

Data quality and ethics are addressed explicitly. Participation is voluntary and can proceed under minimal-risk review; identifying information is stored separately from analysis files, and any message exports from hosted groups are collected with admin consent and scrubbed of personal identifiers. Where disputes are analyzed, summaries are coded at the case level without names; any quotations used for qualitative illustration are paraphrased and member-checked. The study will register a public OSF page containing the pre-analysis plan, survey instruments, coding rubrics, and code for index construction and statistical analysis; de-identified replication datasets will be shared to the extent allowed by host agreements. A short data-management plan specifies retention, access control, and destruction timelines consistent with funder and institutional norms.

Power calculations are design-specific but use conservative assumptions. In Design A, with approximately 300 venture–market observations and within-venture panels over at least eight periods, detectable standardized effects around 0.20–0.25 on primary outcomes are feasible at 80% power with two-sided $\alpha = .05$. In Design B, cluster-randomized rollouts require accounting for intra-cluster correlation; blocking by city and sector and pre-specifying retention thresholds improves efficiency. Simulation-based power using historical variance from Design A will guide site counts per arm before fielding. In Design C, pattern matching focuses on logical triangulation rather than hypothesis testing; nonetheless, coder reliability and cross-case consistency are reported.

Taken together, this measurement and design strategy is intended to satisfy three criteria that typically concern skeptical readers: construct validity (clear mapping from theory to observable and psychometric measures), identification (within-venture designs, randomization where feasible, and documented robustness), and transparency (pre-registration, artifact audits, and open materials). With these in place, the propositions in Section 5 become empirically adjudicable rather than rhetorical.

7. Implications

The framework implies a reordering of early-stage practice in high-context markets. Rather than treating relationships as a marketing function that follows product trials, ventures should treat **relational readiness as a gating milestone** that determines whether product signals will be interpretable. For founders, this means budgeting and calendar time for an **MVR phase** prior to MVP, with explicit go/no-go criteria based on guardian vouchers, embedded access, referral momentum, routine entrenchment, and dispute credibility. Teams commonly try to buy their way around relational veto with promotions and paid traffic; the present framework predicts that these tactics generate expensive, non-durable spikes when authorization and legitimacy are not yet in place. Sequencing **MVR → MVP** instead reframes launch from a one-shot event to a two-stage process in which permission, access, and propagation capacity are established before feature optimization begins. Practically, this shifts early hiring toward operators who can map and engage gatekeepers, design referral scripts that administrators will endorse, implement routine-binding mechanisms with hosts, and codify dispute pathways that match local norms.

For investors, the implications flow through screening, coaching, and portfolio risk. Diligence templates geared to product-market fit and financial runway often miss the **relational constraint** that determines whether spend converts to revenue. The MVR Index and the five underlying indicators provide **auditable leading indicators** of near-term traction: a team with strong vouchers and embedded access but weak dispute dynamics faces a different risk than one with fast referral momentum but no recognized guardians. As a result, investment committees can (i) identify **false negatives**—ventures with modest product maturity but strong MVR that are likely to translate spend efficiently—and (ii) flag **false positives**—ventures with polished products but weak MVR that will burn cash on acquisition without retention. Post-investment, coaching shifts toward helping teams earn permission and entrench routines in targeted communities before scaling paid acquisition. Because the indicators are evidence-based and portable across sectors once calibrated, they also enable **comparable reporting** across a portfolio, reducing reliance on narrative updates.

Accelerators and ecosystem builders can incorporate MVR directly into **readiness screens and curricula**. Application forms can require evidence of guardian vouchers and at least one embedded host channel; bootcamps can include a two-week MVR sprint with artifact audits; and demo days can report an MVR profile alongside product demos. Doing so improves program efficiency by reducing the number of teams that enter growth phases without the relational foundations required for durable adoption. Policymakers and development organizations that support entrepreneurship can take a related step by formalizing **light-touch interfaces between formal and informal governance**—for example, creating registries of recognized group administrators, simplifying letters of introduction, or supporting mediation

fora that shorten time-to-closure for routine disputes. Such measures reduce the transaction cost of earning permission and improve the reliability of informal enforcement, thereby widening the set of entrants who can operate legitimately.

For researchers, the framework offers a tractable way to **separate social sanction from product value** in empirical work. By measuring MVR explicitly and, where feasible, randomizing the sequencing of MVR and MVP, scholars can clarify when negative product signals reflect relational veto rather than low instrumental value. This distinction is crucial for cumulative knowledge, because it prevents the misclassification of contexts as “non-responsive to innovation” when the binding constraint is authorization. Moreover, explicit measurement enables comparative analysis across cities and sectors: with appropriate calibration of reference classes, one can study how much relational sufficiency is “enough” in different settings and how quickly MVR decays or compounds over time.

Finally, the framework carries **ethical implications**. Because MVR involves endorsements and community roles, the pursuit of permission must avoid extractive practices that offload risk onto hosts and guardians. Transparent communication about benefits and obligations, genuine reciprocity (including non-commercial contributions), and fast, fair dispute resolution are not only effective tactics but also **normatively desirable** in systems where social sanction substitutes for formal enforcement. In short, the same behaviors that raise an MVR score also constitute good citizenship in high-context markets.

8. Limitations and Future Work

This framework is intentionally narrow in scope: it theorizes conditions under which **relational readiness precedes product readiness** in high-context markets. Several limitations follow. First, **construct validity** remains a live concern. The seven dimensions of MVR are substantively motivated, but their empirical distinctness must be demonstrated psychometrically; overlapping variance between, for example, trust-signal density and reputational capital could collapse into fewer factors in some settings. The proposed program explicitly addresses this through item development, factor analysis, and tests of convergent and discriminant validity, yet replication across sectors and cities will be required for cumulative confidence.

Second, the **observational indicators** (guardian vouchers, embedded distribution, referral momentum, micro-entrenchment, dispute dynamics) are designed for auditability but are not immune to gaming or measurement error. A venture could, for instance, collect nominal endorsements from peripheral actors or close disputes quickly but superficially. To mitigate

these risks, the measurement plan pairs speed with recurrence (for disputes), weights vouchers by *local* credibility rather than raw counts, reconciles embedded distribution with host-side records, and emphasizes robust statistics (medians, IQRs) over means. Even so, indicator integrity will vary with the quality of local record-keeping and the willingness of hosts to share artifacts; sensitivity analyses and independent audits should therefore be routine.

Third, **endogeneity and reverse causality** are plausible: traction can attract endorsements, and successful teams may gain access to embedded channels ex post. The empirical designs respond with lag structures, within-venture variation (before/after activation of hosts), fixed-effects panel models, and a field experiment that randomizes sequencing. Nonetheless, instruments for relational variables are hard to defend convincingly; where credible quasi-experimental shocks exist (e.g., exogenous gatekeeper turnover or sudden closure of a host channel), they should be exploited and reported transparently alongside non-experimental estimates.

Fourth, **boundary conditions** matter. In low-context settings with thick formal enforcement or where product externalities dominate social sanction, MVP-first may be sufficient and efficient. Even in high-context environments, **over-embeddedness** can create path dependence and slow adaptation if ventures become beholden to a narrow set of guardians or routines. The framework therefore recommends *minimum* sufficiency rather than maximal embeddedness. Future work should map the response surface: how much MVR is “enough,” where diminishing returns begin, and when over-embeddedness harms adaptability.

Fifth, there are **ethical and equity considerations**. Because permission often flows through local elites or administrators, the pursuit of vouchers can entrench gatekeeping or replicate exclusionary practices. The proposed practices—genuine reciprocity, transparent obligations, and norm-concordant dispute resolution—are intended to counter extractive dynamics, but empirical work should still examine distributional consequences: who gains or loses from formalizing MVR, and whether marginalized entrants face systematically higher relational thresholds.

Sixth, **external validity** is conditional. Although the theory likely transports to other high-context arenas (e.g., tightly networked diasporas, informal labor markets outside Africa), calibration is essential. Reference classes (for standardization), indicator weights, and even the salience of particular dimensions may shift with sector (health vs. mobility), channel architecture (agent networks vs. digital platforms), or city. Cross-site studies with explicit calibration protocols will help delineate what travels and what must be localized.

Finally, **research logistics** pose familiar constraints. Power for cluster-randomized sequencing experiments depends on intra-cluster correlation and feasible site counts; attrition and data-sharing frictions can erode effective sample sizes. The program anticipates this with blocking, covariate adjustment, pre-registration, and open materials, but credible nulls should be reported when they arise. Future work might also integrate **computational traces** (e.g., admin post engagement, message-level dispute timelines) to reduce recall bias and enrich mechanism tests. A complementary direction is to study **decay and compounding**: does MVR erode without maintenance, and under what conditions do vouchers and embedded routines produce flywheel effects that lower acquisition costs over time?

9. Conclusion

In high-context markets where informal sanction and reputational logics structure exchange, **relational readiness precedes product readiness**. The construct of **Minimum Viable Relationships (MVR)** formalizes the minimally sufficient configuration of trust, embedded access, and permission-to-operate that allows MVP tests to yield interpretable signals. By specifying seven dimensions of MVR, mapping them to field-measurable indicators, and proposing an auditable **MVR Index**, the paper separates social veto from product failure and offers a tractable agenda for empirical validation. The practical implication is a re-sequencing of early venture activity—**MVR → MVP → PMF**—which, if supported by evidence, should reduce avoidable false negatives, lower customer-acquisition costs, and increase the resilience of traction. For researchers, explicit measurement and (where feasible) randomized sequencing open identification strategies that move beyond rhetoric to adjudicable claims. For founders, investors, and ecosystem builders, treating permission, embeddedness, and dispute credibility as first-order constraints reframes launch from feature optimization to **earning the right to operate**.

Author note: Funding: None. Conflicts of interest: None. Ethics: Conceptual/theoretical work; no human participants or sensitive data.

References (APA)

- Blank, S. (2013). *The four steps to the epiphany* (2nd ed.). K&S Ranch.
- Burt, R. S. (1992). *Structural holes: The social structure of competition*. Harvard University Press.
- Coleman, J. S. (1988). Social capital in the creation of human capital. *American Journal of*

Sociology, 94(S1), S95–S120.

Dyer, J. H., & Singh, H. (1998). The relational view: Cooperative strategy and sources of interorganizational competitive advantage. *Academy of Management Review*, 23(4), 660–679.

Fafchamps, M. (2004). *Market institutions in Sub-Saharan Africa: Theory and evidence*. MIT Press.

Fornell, C., & Larcker, D. F. (1981). Evaluating structural equation models with unobservable variables and measurement error. *Journal of Marketing Research*, 18(1), 39–50.

Granovetter, M. (1985). Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, 91(3), 481–510.

Hall, E. T. (1976). *Beyond culture*. Anchor.

Henseler, J., Ringle, C. M., & Sarstedt, M. (2015). A new criterion for assessing discriminant validity in variance-based SEM. *Journal of the Academy of Marketing Science*, 43(1), 115–135.

McDonald, R. P. (1999). *Test theory: A unified treatment*. Lawrence Erlbaum.

Meagher, K. (2010). *Identity economics: Social networks and the informal economy in Nigeria*. James Currey.

Putnam, R. D. (2000). *Bowling alone: The collapse and revival of American community*. Simon & Schuster.

Ries, E. (2011). *The lean startup*. Crown Business.

Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.

Uzzi, B. (1997). Social structure and competition in interfirm networks: The paradox of embeddedness. *Administrative Science Quarterly*, 42(1), 35–67.

Brown, T. (2008). Design thinking. *Harvard Business Review*, 86(6), 84–92.

Liedtka, J. (2018). Why design thinking works. *Harvard Business Review*, 96(5), 72–79.

Prahalad, C. K., & Ramaswamy, V. (2004). Co-creating unique value with customers. *Strategy & Leadership*, 32(3), 4–9.

Appendix A. Two-Week MVR Sprint (standardized protocol for Design B)

This appendix details the pre-MVP sprint referenced in Section 6. The sprint’s objective is to raise a venture above the **MVR threshold** so that subsequent MVP results are interpretable.

Days 1–2: Gatekeeper inclusion. Map the relevant guardians (e.g., association leaders, group admins, building managers). Secure at least three explicit vouchers—hosted posts, signed notes, or on-record audio—and record contactability for audit.

Days 3–5: Embedded channel activation. Establish one live host channel (co-sell inside a cooperative/SACCO, shelf/slot with an anchor merchant, or a hosted WhatsApp group with active admin). Instrument source tagging so acquisition from the host can be attributed.

Days 6–8: Referral momentum. Deploy simple scripts and instrumented links/codes in the host environment. Target a **median ≤ 5 days per net-new referred customer** over the window; log admin pushes and peer shares.

Days 9–11: Micro-entrenchment. Put at least two routine-binding mechanisms into live use—e.g., float accounts, standing orders, shared ledgers, pinned workflow groups. Verify usage frequency with artifacts (screens, receipts, logs).

Days 12–14: Dispute protocol. Agree a three-step, locally recognized path for routine issues; dry-run with a consenting counterparty. Track median time-to-closure and 30-day recurrence.

Go/No-Go thresholds. Before commencing MVP, aim for: **GV ≥ 4** , **ED $\geq 30\%$** of new customers via hosts, **RM ≥ 3** net-new referred customers per 14 days (median interval ≤ 5 days), **ME ≥ 2** mechanisms in active use, and **DD median closure < 48 hours** with low recurrence. Where thresholds are missed, recycle the sprint before moving to MVP.