ISSN 2090-3359 (Print)
ISSN 2090-3367 (Online)

ΑΔΣ

Advances in Decision Sciences

Volume 30
Issue 4
December 2026

亞洲大學 ASIA UNIVERSITY

SCIENTIFIC & BUSINESS WORLD

Published by Asia University, Taiwan and Scientific and Business World

Financial Digitalization and Urban Green Innovation: Evidence from Panel Quantile ARDL Model

Antanas Laurinavicius
Department of Finance,
Faculty of Economics and Business Administration,
Vilnius University, Vilnius, Lithuania
Email: antanas.laurinavicius@evaf.vu.lt
ORCID: https://orcid.org/0000-0002-7983-2779

Algimantas Laurinavicius
Department of Finance,
Faculty of Economics and Business Administration,
Vilnius University, Vilnius, Lithuania
***Corresponding author** Email: algimantas.laurinavicius@evaf.vu.lt
ORCID: https://orcid.org/0000-0003-0145-2386

Asma Salman
American University in the Emirates
Email: asma.salman@aue.ae
ORCID: https://orcid.org/0000-0002-5623-3087

Muthanna G. Abdul Razzaq
American University in the Emirates
Email: dr.muthanna@aue.ae
ORCID: https://orcid.org/0009-0007-4831-1907

Mohamed Elsayed Abdelsalam Ghanem
Mansoura University, Egypt
Email: pro.mohamedghanem@gmail.com
ORCID: https://orcid.org/0009-0006-9176-5659

Received: June 30, 2026; First Revision: August 2, 2026;
Last Revision: August 27, 2026; Accepted: September 3, 2026;
Published: September 9, 2026

Abstract

Purpose – Digital finance is spreading unevenly across cities, and whether it narrows or widens the green-innovation gap is unknown because existing evidence is mean-based. This study asks where in the distribution of green patenting digital finance matters, for 16 large U.S. metropolitan areas observed annually over 2000–2025.

Design/methodology/approach – A panel quantile ARDL model, derived from a directed-innovation model with financing frictions, is estimated in error-correction form. Linearity, cointegration and spurious-regression checks precede estimation, and inference uses metro-level cluster and wild cluster bootstraps.

Findings – The long-run semi-elasticity is positive at every quantile and rises monotonically, roughly doubling across the interquartile range from 0.1370 to 0.2860. Error correction is negative throughout and faster in the upper quantiles, so leading metros gain more and converge sooner.

Originality/value – This is the first metropolitan-level distributional treatment of the digital-finance and green-innovation link, and it shows that the substitution and complementarity accounts are opposite signs of one cross-partial derivative that only a quantile design can identify.

Implications – The estimand is a decision input rather than a descriptive average: it tells an agency allocating scarce innovation funds what an extra unit of digital finance buys in a metro of given rank, which is the sense in which this is a decision-sciences contribution. Because gains accrue where capacity is already deep, promoting digital finance is necessary but not sufficient.

Keywords: digital finance; green innovation; quantile ARDL; panel cointegration; metropolitan areas

JEL Classifications: G20; O31; Q55; R11; C31

1. Introduction

The decarbonization of urban economies depends on a steady supply of green technologies, and those technologies must be financed. Over the past two decades, the financing channel has itself changed shape: mobile payments, digital lending platforms, online crowdfunding, and data-rich credit scoring have lowered the cost of moving capital toward small, risky, early-stage projects — the profile of most environmental innovation. Whether this has translated into more green patents, and for whom, remains an open question, taken up here at the metropolitan scale using a balanced panel of 16 large U.S. metros observed annually from 2000 to 2025.

The stakes are concrete. Metropolitan areas account for the majority of energy use and a large share of greenhouse-gas emissions, and meeting mid-century climate targets requires generating new clean technologies, not merely deploying existing ones. Where green patents are produced is, therefore, a first-order question for the geography of the energy transition.

Two stylized facts motivate the study. Green patenting in U.S. metros is highly concentrated, with a handful of coastal hubs accounting for a disproportionate share of environment-related (CPC Y02) patents, and the diffusion of digital finance has been broad but uneven (Feng et al., 2022; Gomber et al., 2017). If the two simply moved together on average, a fixed-effects regression would capture the link, but averages are the wrong summary when the policy concern is the spread of outcomes across places. A leading metro may convert each additional unit of digital finance into many more patents than a lagging one, or the reverse may hold if digital finance is most transformative where formal finance was thin, and distinguishing the two requires estimation across the whole conditional distribution.

The U.S. setting sharpens the question. Unlike the bank-dominated systems from which most digital-finance evidence originates, the United States pairs deep capital markets with pronounced regional disparities in venture funding, banking access, and research infrastructure. Digital payments, online small-business lending, and fintech-intermediated investment expanded sharply after 2008 but unevenly across space, and green patenting followed an even more skewed path. Were the marginal effect constant, this co-movement would say nothing about distributional consequences; if the effect strengthens with innovation capacity, joint diffusion would actively pull leading and lagging metros apart.

We use the panel quantile autoregressive distributed lag (QARDL) framework of Cho et al. (2015), which embeds quantile regression (Koenker, 2004) inside an error-correction ARDL structure. It accommodates regressors of mixed integration order, $I(0)$ and $I(1)$, without forcing a pretest commitment, separates a long-run cointegrating relationship from short-run adjustment, and lets the long-run coefficients, the speed of adjustment, and the short-run dynamics each vary by quantile (Cho et al., 2015; Koenker, 2004).

This study makes three contributions. First, it reframes the relationship as a distributional question and supplies the first metropolitan-level panel QARDL treatment of it, moving past the conditional-mean evidence that dominates the literature (Feng et al., 2022; Xing et al., 2024). Second, it characterizes two distributional margins a mean model conceals: the long-run elasticity rises monotonically across quantiles and the speed of error correction is itself quantile-dependent, so leading and lagging metros differ in how much they gain and in how fast. Third, because the effect is strongest where innovation is already strong, an unguided expansion of digital finance is likely to widen the green-innovation gap.

Two features place the study within the decision sciences, the field this journal serves. First, the object of estimation is a decision input rather than a descriptive correlation: the quantile-specific long-run coefficient is the return, in green patents, a funding agency would attach to an additional unit of digital-finance depth in a metro of given innovation rank, and the quantile-specific speed of adjustment is the horizon over which that return arrives. Allocating a fixed budget across places is a resource-allocation problem under uncertainty, and a profile of state-contingent returns is what it requires. Second, a conditional-mean model returns a number that is correct on average and misleading precisely where the allocation is hardest.

What makes the study original can be stated concretely. No prior work estimates this relationship across the conditional distribution of the outcome at the metropolitan scale, and none derives the distributional prediction from an explicit optimization model rather than from the skewness of patent counts. The substitution and complementarity accounts that the literature treats as competing narratives are shown here to be opposite signs of a single cross-partial derivative, which a quantile design identifies and a mean design cannot.

The paper proceeds in the standard order. We test for panel unit roots, confirm a long-run relationship, then estimate long-run and short-run QARDL coefficients at five quantiles and test their equality. Digital finance raises green innovation in the long run, with the elasticity climbing from roughly 0.08 at the lowest quantile to about 0.37 at the highest, and adjustment is faster at the top. Section 2 reviews the literature, Section 3 develops the model and hypotheses, Section 4 presents the data and methodology, Section 5 presents the results, and Section 6 concludes.

2. Literature review

2.1 Finance, innovation, and the financing of green technology

That financial development advances the technological frontier is long and well-documented: screening, risk pooling, and the channeling of investment raise the frequency of useful innovation. Innovation financing is nonetheless distinctive, because research and development is intangible, uncertain, and hard to collateralize, which raises external financing costs and invites credit rationing. Green innovation inherits those frictions and adds long payback periods and an environmental externality that is not internalized, so the financing gap is, if anything, larger for clean than for conventional innovation.

This is where digital finance enters. Digital platforms reduce search, verification, and transaction costs, extending credit and equity to borrowers and ventures that the traditional banking system did not serve (Gomber et al., 2017), and moving intermediation online has measurable effects on activity and growth (Noman et al., 2023). Emerging evidence associates digital finance with better green-technology outcomes, generally through relaxed financing constraints (Chen et al., 2024; Feng et al., 2022; Lu & Xia, 2024; Xing et al., 2024; Yang & Hui, 2024). What that literature leaves aside is the distributional question central here: whether the association differs between innovation leaders and laggards.

The most recent evidence comes from firm- and city-level studies, primarily in China. Several find that digital or digital-inclusive finance enhances corporate green innovation by relieving financing constraints, particularly among non-state firms (Chen et al., 2024; Feng et al., 2022; Lu & Xia, 2024; Yang & Hui, 2024), and Lu and Xia (2024) report spatial spillover and threshold effects for Chinese prefecture-level cities. Two gaps remain. The evidence is overwhelmingly Chinese and firm-centric, leaving the mechanism untested at the metropolitan scale in a market-based economy where the marginal borrower differs, and even studies allowing for thresholds report effects around the conditional mean, when it is the shape away from that mean that determines whether digital finance equalizes or polarizes.

The marginal contribution over the closest studies can be stated explicitly. Feng et al. (2022) estimate the average effect for Chinese provinces in a conditional-mean panel and do not ask where in the distribution the effect operates. Wen et al. (2025) come closest with threshold effects and diminishing marginal returns, but a threshold design partitions the sample and still reports mean effects within regimes, whereas a quantile design traces the effect continuously across the distribution of the outcome itself; their evidence is also Chinese. Yang and Hui (2024) work at the firm level, where the metro-level ecosystem complementarities of interest here are conditioned away.

Mechanically, digital finance loosens several constraints at once. Platform lenders replace collateral and relationship history with algorithmic screening, which changes the terms of access for young firms whose assets are largely intangible; digital payment rails lower the cost of reaching customers and investors; and digital intermediation enlarges the capital pool for small projects by cutting matching costs. Each channel addresses the same problems — opacity, illiquidity, and scale — that make green research hard to finance conventionally.

2.2 Green finance, the environment, and the directed-innovation view

A parallel literature concerns the direction rather than the speed of technical change. Left alone, innovation is path dependent and favours dirty technologies (Acemoglu et al., 2012), so finance and regulation must steer it, and the Porter hypothesis holds that well-designed regulation can induce offsetting innovation. Green-finance products and targeted credit can shift capital toward cleaner projects, and digital finance can serve the same end, provided the channel reaches the right places (Chen et al., 2024; Feng et al., 2022; Hossain et al., 2024; Lu & Xia, 2024).

Digital finance should also be distinguished from the broader green-finance toolkit of green bonds and targeted green credit, which lowers the cost of capital from the top down and reaches larger issuers first. Digital finance instead reduces the cost of intermediation itself, in principle reaching the small, collateral-light ventures behind much early green invention. If it largely replaces missing conventional finance, its largest effects should appear in lower-innovation metros, and diffusion would be equalizing; if it complements existing ecosystems, they should appear at the top, and diffusion would be polarizing. The design adopted here lets the data adjudicate.

This also explains the choice of scale. Innovation is local: knowledge spillovers, labor-market pooling, and supplier networks operate within commuting zones, and venture capital is concentrated in a few metros. A national study averages over that geography, and a firm-level study conditions it away, whereas the distributional hypothesis concerns exactly the interaction between a metro's innovation ecosystem and its access to digital finance.

Environmental regulation enters as both a complement and a rival explanation. The Porter hypothesis holds that local standards can drive green invention by creating demand for compliant technologies, which is why a regulation index is included among the controls. Because progressive metros may adopt both, omitting regulation would risk attributing its effect to digital finance; conditioning on it allows the digital-finance coefficient to be read as the association net of the regulatory environment.

The mapping from these strands to the hypotheses is as follows. The financing-friction argument of Section 2.1 implies a positive long-run association and underpins H1. The substitution-versus-complementarity contrast developed here generates the two opposing distributional predictions that H2 adjudicates: substitution implies the effect is largest at the bottom of the innovation distribution, and complementarity that it rises toward the top. The same logic in the time dimension — constrained metros absorb financing improvements more slowly — motivates the quantile-dependent adjustment speeds tested by H3.

2.3 Nonlinearity, asymmetry, and the distributional turn in applied finance

The methodological premise is that the finance–innovation relationship is unlikely to be a single constant slope. Threshold and nonlinear evidence show that the effect of digital finance on green outcomes can switch with the level of a conditioning variable (Lu & Xia, 2024; Wen et al., 2025). Quantile methods go further, estimating effects across the entire conditional distribution; the apparatus is set out in Koenker (2004) and Cho et al. (2015), and the quantile ARDL separates long-run from short-run effects while preserving distributional detail (Hashmi & Chang, 2023). Patent counts are heavily right-skewed, so the conditional mean is a poor summary.

Two features make the quantile ARDL particularly suitable for innovation data. Patent counts are bounded below and heavily right-skewed, because a small number of prolific metros pull the average slope upward. And the error-correction form separates the immediate from the cumulative response, so the quantile extension lets both vary with a metro's position in the distribution, which matters because the speed of catch-up itself indicates how tightly the financing constraint binds.

Quantile methods have succeeded in energy and environmental economics precisely because the relationships there are state-dependent, with elasticities at the extremes routinely differing from those at the centre. The same logic applies to green innovation, where the incentives to patent in an established research powerhouse are unlikely to carry the elasticities of a manufacturing-intensive metro.

The distributional turn this paper joins has both a long pedigree and a fast-growing recent literature. Its foundations are the regression-quantile framework and its panel extension (Koenker, 2004) and the quantile-cointegration synthesis of Cho et al. (2015). Alongside them sit the diagnostics that make a distributional design defensible rather than merely convenient: the nonlinearity test of Hui et al. (2017) and the multivariate nonlinear causality test of Bai et al. (2018).

A large recent literature establishes that quantile designs routinely recover state dependence that mean estimators average away: in the transmission of monetary-policy uncertainty into equity returns across developed, emerging and frontier markets (Arshad et al., 2024; Zada et al., 2026), in disease-induced volatility (Gohar et al., 2023; Salman et al., 2023), in the predictability of United States housing returns from climate risk (Bouri et al., 2025), in tail-risk transmission across volatility indices, cryptocurrencies and currency pairs (Imane & Ghizlane, 2025; Syed et al., 2025), in non-ferrous metals (Woode et al., 2025), in price discovery under speculative sentiment (Bandyopadhyay & Rajib, 2026) and in policy evaluation (Laurinavicius et al., 2026). What this paper adds is a setting in which the distributional question has explicit equity content, because the units are places rather than assets, so the gradient describes whether a diffusing technology converges or diverges across them.

2.4 Research gap

Three gaps emerge. The digital-finance–green-innovation literature is overwhelmingly mean-based and therefore silent on whether the effect differs across the innovation distribution. Most evidence is national or firm-level, leaving the metropolitan scale, where innovation actually clusters, underexplored. And few studies separate the long-run equilibrium effect from short-run adjustment while allowing both to vary by quantile. The panel QARDL design closes all three, and Section 3 derives the hypotheses from an explicit model rather than asserting them.

3. Theoretical framework and hypothesis development

The empirical design is distributional, and the hypotheses it tests are distributional claims, which are not self-evident. This section, therefore, sets out a compact model of directed green innovation under financing frictions from which they follow, retaining only the two features on which the distributional question turns: green innovation combines research effort with local absorptive capacity, and the cost of financing that effort falls as digital financial infrastructure deepens.

3.1 A directed-innovation model with financing frictions

Consider a representative innovator in metropolitan area i in year t. Following the directed-technical-change tradition (Acemoglu et al., 2012), research effort can be allocated between a dirty and a clean technology line; attention here is restricted to the clean line, whose output is the flow of green inventions.

Green innovation output $G_{i,t}$ is produced from real green research effort $R_{i,t}$ and local absorptive capacity $A_{i,t}$ according to the production technology specified as follows:

$$G_{i,t} = A_{i,t}\, R_{i,t}^{\eta}, \quad 0 < \eta < 1, \tag{1}$$

where $A_{i,t}$ is a composite of the complementary inputs a metropolitan innovation ecosystem supplies — research universities, the local supply of inventors, specialised suppliers, and the infrastructure for prosecuting patents — and the concavity parameter η captures diminishing returns to research effort within a period. Absorptive capacity is metro-specific and slow-moving, and it is what places a metropolitan area high or low in the cross-section of green patenting.

Research effort must be financed externally, because green research and development is intangible, hard to collateralize, and has a long, uncertain payback. The innovator, therefore, pays an external-finance premium κ over the risk-free cost of funds, so that the total cost of research effort is specified as follows:

$$C_{i,t}(R) = R_{i,t}\left[1 + \kappa\left(D_{i,t}, A_{i,t}\right)\right], \quad with\ \frac{\partial \kappa}{\partial D} < 0, \tag{2}$$

where $D_{i,t} \in [0, 100]$ denotes the depth of digital financial infrastructure in metropolitan area i, measured by the composite digital-finance index. The sign restriction $\partial\kappa/\partial D < 0$ is the single substantive assumption the model makes about digital finance, and the mechanism literature supports it: platform lending, algorithmic screening, and digital payment rails lower the search, verification, and monitoring costs that generate the premium (Chen et al., 2024; Gomber et al., 2017). The premium is also allowed to depend on absorptive capacity, because lenders in thick innovation ecosystems observe more informative signals.

Let V denote the expected private value of a green invention, treated as a constant common to all metropolitan areas, and let $\Pi_{i,t}$ denote the innovator's expected net payoff, that is, expected revenue from green inventions less the financed cost of the effort producing them. The innovator chooses research effort $R_{i,t}$ to maximise that payoff, which is defined as follows:

$$\Pi_{i,t}\left(R_{i,t}\right) = V\, A_{i,t}\, R_{i,t}^{\eta} - R_{i,t}\left[1 + \kappa\left(D_{i,t}, A_{i,t}\right)\right], \tag{3}$$

where $\Pi_{i,t}$ is the expected net payoff, V $A_{i,t}\, R_{i,t}^{\eta}$ the expected revenue generated by research effort $R_{i,t}$, and $R_{i,t}[1 + \kappa(D_{i,t}, A_{i,t})]$ the cost of financing that effort at the external-finance premium of Equation (2). Maximising Equation (3) with respect to research effort and solving the resulting first-order condition gives the interior optimum, expressed as follows:

$$R_{i,t}^{*} = \left[\frac{\eta V A_{i,t}}{1 + \kappa\left(D_{i,t}, A_{i,t}\right)}\right]^{\frac{1}{1-\eta}}, \tag{4}$$

where an asterisk denotes the payoff-maximising value of a variable, so that $R_{i,t}^{*}$ is the research effort the innovator actually chooses. Substituting Equation (4) into Equation (1) and taking logarithms gives the reduced form for equilibrium green innovation, expressed as follows:

$$\ln G_{i,t}^{*} = \frac{1}{1 - \eta}\left[ln\, A_{i,t} + \eta\, ln(\eta V) - \eta\, ln\left(1 + \kappa\left(D_{i,t}, A_{i,t}\right)\right)\right]. \tag{5}$$

Equation (5) is the theoretical counterpart of the long-run cointegrating relationship estimated in Section 4.2, linking equilibrium green innovation to absorptive capacity and to the cost of external finance, both slow-moving metropolitan characteristics. Differentiating with respect to digital finance yields the long-run semi-elasticity, expressed as follows:

$$\beta\left(A_{i,t}\right) \equiv \frac{\partial\, ln\, G_{i,t}^{*}}{\partial D_{i,t}} = -\frac{\eta}{1 - \eta}\frac{\kappa_D\left(D_{i,t}, A_{i,t}\right)}{1 + \kappa\left(D_{i,t}, A_{i,t}\right)}, \tag{6}$$

where κ_D ≡ ∂κ/∂D and κ_A ≡ ∂κ/∂A denote the partial derivatives of the external-finance premium with respect to digital finance and absorptive capacity, and κ_{DA} ≡ ∂²κ/∂D∂A, used in Section 3.2, denotes the corresponding cross-partial derivative. Equation (6) is the theoretical semi-elasticity whose sample counterpart $\beta(\tau)$ is estimated quantile by quantile in Equation (9) and reported in Table 7.

3.2 From the model to the testable hypotheses

Three results follow from Equations (5) and (6), each mapping onto one hypothesis. The first is a sign restriction. Because η ∈ (0,1) and κ_D < 0, the right-hand side of Equation (6) is strictly positive: deeper digital finance lowers the external-finance premium, raises optimal research effort, and therefore raises equilibrium green innovation, with magnitude scaled by η/(1−η). This is the content of H1, a statement about long-run equilibrium rather than about impact effects.

The second result is the distributional one, and it is where the model earns its place. Differentiating Equation (6) with respect to absorptive capacity gives a cross-partial derivative whose sign is not determined by the assumptions imposed so far. Two configurations are possible, corresponding to the substitution and complementarity accounts of Section 2.2. Under the first, digital finance is most cost-reducing where conventional intermediation is thin, so the absolute size of the financing term falls as capacity rises and ∂β/∂A < 0: the association is largest at the bottom, and diffusion is equalizing. Under complementarity, deeper ecosystems convert a given reduction in the premium into more inventions, so ∂β/∂A > 0: the association is largest at the top and diffusion is polarizing. The two accounts are therefore opposite signs of a single cross-partial derivative, and the empirical design is built to recover that sign.

The link from ∂β/∂A to the quantile index is direct. Equation (5) shows that ln $G_{i,t}^{*}$ is strictly increasing in $A_{i,t}$ holding the financing term fixed, so among metropolitan areas alike in the observed covariates, those higher in the conditional distribution of log green innovation are those with higher absorptive capacity. The conditional quantile τ is therefore an ordinal index of A, and $\beta(\tau)$ inherits the sign of ∂β/∂A. A monotonically increasing $\beta(\tau)$ is the signature of complementarity, a decreasing $\beta(\tau)$ that of substitution. This is the content of H2, and it explains why a conditional-mean estimator cannot adjudicate: the average of $\beta(\tau)$ is positive under both accounts. This mapping carries an identifying assumption that should be named. Reading the conditional quantile as an ordinal index of absorptive capacity requires rank invariance, or the weaker rank similarity of Chernozhukov and Hansen (2005): a metropolitan area's position in the conditional distribution of green patenting must be governed by the latent capacity term rather than by factors unrelated to it. It fails if the upper tail is populated by metros whose position reflects something else — opportunistic environmental labeling of patents that are not substantively green, or a localized clean-technology subsidy shock. Two pieces of evidence bound that risk without removing it: the absorptive-capacity split and the interaction test of Section 5.5 recover the same gradient from observable proxies for capacity. The assumption is nevertheless not testable here, and Section 6 records it as a limitation.

The third result concerns dynamics. Realized research effort adjusts to its target only partially within a period, because financing relationships take time to form, and green research has a long gestation. Writing that adjustment as $\Delta \ln G_{i,t} = \lambda(A_{i,t})\,[\ln G^{*}_{i,t-1} - \ln G_{i,t-1}]$ plus short-run terms, the adjustment coefficient is itself a function of capacity: a metro whose financing constraint binds tightly closes a given gap more slowly, because each additional unit of effort is financed at a higher marginal premium. Convergence requires $\lambda(A) < 0$, with the absolute value of λ increasing in A and hence in τ. This is the content of H3.

Restating these three results as testable propositions gives the hypotheses examined in the remainder of the paper.

H1. Digital finance has a positive long-run effect on green innovation across U.S. metropolitan areas.

H2. The long-run effect of digital finance on green innovation is heterogeneous across quantiles, increasing from lower to higher quantiles of the green-innovation distribution.

H3. The error-correction (adjustment) speed toward the long-run equilibrium is negative at all quantiles and differs in magnitude across the distribution.

H2 is the discriminating test: substitution predicts a decreasing profile and complementarity an increasing one, so the sign of the estimated gradient identifies which mechanism dominates. The panel quantile ARDL design of Section 4 estimates $\beta(\tau)$ and $\lambda(\tau)$ within the single framework described in Section 4.2, and therefore tests all three propositions together.

4. Data and Methodology

4.1 Data, sample, and variables

The analysis uses a panel of 16 large U.S. metropolitan areas — New York, Los Angeles, Chicago, Houston, Phoenix, Philadelphia, San Antonio, San Diego, Dallas, San Jose, Austin, Seattle, Boston, San

Francisco, Atlanta, and Denver — observed annually from 2000 to 2025, giving $16 \times 26 = 416$ metro-year observations. Green innovation is measured as metro-level counts of environment-related (CPC Y02) patents, adjusted in recent grant years for examination-lag truncation; digital finance is a composite digital-financial-inclusion index; and each control is drawn from standard public series. Table 1 defines each variable, its measurement, its unit, and the expected sign of its long-run coefficient.

Table 1. Variable definitions and expected signs

Variable	Symbol	Definition/measurement	Unit	Expected sign
Green innovation	GINNOV	Count of green/environmental patents granted; log used in estimation	count	dependent
Digital finance	DIGFIN	Composite digital financial development/inclusion index	0–100 (rescaled to 0–1 in estimation)	+
Economic development	lnGDPpc	Log real GDP per capita	Log	+
R&D intensity	RD	R&D expenditure as a share of metro output	%	+
Human capital	HC	Share of adults with a bachelor's degree or higher	%	+
Industrial structure	IND	Manufacturing share of metro output	%	+/−
Environmental regulation	ENVREG	Composite stringency-of-environmental-policy index	0–100	+

Note. The table defines every variable in the model of Section 4.2, with its symbol, measurement, unit, and expected sign. GINNOV is the count of green (CPC Y02) patents granted to inventors resident in the metropolitan area, entered as ln GINNOV in all estimation. DIGFIN is a composite digital-financial-development and inclusion index, published on a 0–100 scale and rescaled to the unit interval before estimation, so the coefficients in Tables 7, 8, 10, and 11 are semi-elasticities with respect to the rescaled index. lnGDPpc is log real GDP per capita; RD is research and development expenditure as a share of metropolitan output, in percent; HC is the share of adults aged 25 and over with a bachelor's degree or higher, in percent; IND is the manufacturing share of metropolitan output, in percent; and ENVREG a composite index of environmental-policy stringency on a 0–100 scale. The expected sign is the anticipated direction of the long-run association with ln GINNOV, and +/− indicates that the theory does not specify a sign. The panel covers 416 metro-year observations, 384 after the ARDL lags.

The controls capture the main non-financial determinants of green patenting: economic development proxies the demand and resources available for risky research; research and development intensity the direct effort that finance ultimately funds; human capital the supply of inventors; industrial structure the metro's output mix; and environmental regulation the demand-pull of policy. Together, they isolate the marginal association of digital finance from the rest of the innovation environment.

Several measurement choices deserve note. The dependent variable counts green rather than all patents and is logged because the raw distribution is heavily skewed, which also makes it suitable for quantile estimation. Because grant lags at the United States Patent and Trademark Office run to between eighteen and thirty-six months, counts in the most recent application years are right-truncated, and the adjustment is as follows. Let F be the empirical cumulative distribution function of the application-to-grant lag for CPC Y02 filings, estimated on cohorts old enough to be effectively complete, and let a be the number of years between a cohort and the sample end. The observed count is divided by F(a), the share of that cohort expected to have been granted by then, so the adjusted count estimates the eventual cohort total. F approaches one within four years, so the adjustment is material only for the final three application years. Digital finance is measured as a composite index rather than a single instrument, since no single measure proxies the others. Because no established metropolitan digital-finance index exists for the United States — the Peking University index that anchors the Chinese literature has no American counterpart — DIGFIN is constructed for this study. It aggregates four components measured annually at the metropolitan level: household and small-business use of digital payment and online banking services; the penetration of non-bank online and marketplace lending; the density of financial-technology establishments relative to metropolitan employment; and household broadband access as the enabling infrastructure. Each is min–max normalised across metros and years onto the 0–100 interval, and the four are combined using weights from the first principal component of the normalised series. The index is therefore a relative rather than an absolute measure of digital-financial depth, which is why Section 5.5 reports the check replacing it with a single transaction-value proxy. The index is reported on the 0–100 scale in the descriptive statistics but rescaled to the unit interval before estimation, so every DIGFIN coefficient below is the semi-elasticity with respect to a one-unit move in the rescaled index. Each regressor enters with the lag structure the ARDL form requires, so the long-run relationship is identified from within-metro variation.

Identification rests on the cointegrating relationship rather than on an external instrument, so it is worth being explicit about what the estimates claim. The long-run coefficients describe the equilibrium association once short-run noise has washed out, conditional on the controls and metro fixed effects; they are not randomized treatment effects. Metro fixed effects absorb time-invariant differences, and the error-correction form defines the disequilibrium on lagged levels, which dampens simultaneity. Two threats remain. The first is reverse causality: successful innovation may attract venture investors and platform lenders, and cointegration does not orient the arrow. The second is an omitted common factor, such as a regional technology boom, against which the controls, the survival of the gradient when the two Bay Area hubs are dropped, and the placebo randomization all speak. The long-run coefficient should therefore be read as an equilibrium association within metros, not as the causal effect of an exogenous intervention. A third threat is specific to the short-run block. Equation (8) contains the contemporaneous first difference of digital finance, which is hard to defend when green patenting and platform-finance expansion within the same metropolitan area are plausibly driven by a common local innovation shock, so that the orthogonality condition fails. Three things bear on how much this matters. The long-run vector, which carries every claim the paper makes, is identified from the lagged levels of Equation (9) rather than from contemporaneous differences, so it is insulated by construction. The impact coefficient in Table 8 is correspondingly read as a within-year conditional association rather than as a response to an exogenous movement in digital finance. And Table 10 reports the specification that removes the contemporaneous term, admitting the digital-finance difference only from the first lag onward so that every right-hand-side variable lies in the information set dated $t - 1$ or earlier; the long-run profile is materially unchanged.

4.2 The panel quantile ARDL model

Let i index metros and t index years. The long-run (cointegrating) relationship is specified as follows:

$$\ln GINNOV_{it} = \alpha_i + \beta_1 DIGFIN_{it} + \beta_2 lnGDPpc_{it} + \beta_3 RD_{it} + \beta_4 HC_{it} + \beta_5 IND_{it} + \beta_6 ENVREG_{it} + \varepsilon_{it}, \tag{7}$$

where α_i is a metro fixed effect and ε_{it} is the equilibrium error. Standard estimation of Equation (7) returns a single slope vector describing the average metro, and by ordinary or pooled methods yields the bounds-testing ARDL relationship of Pesaran et al. (2001). To recover distributional detail, we follow Cho et al. (2015) and embed panel quantile regression (Koenker, 2004) inside an error-correction ARDL form. For a quantile $\tau \in (0,1)$, the conditional τ-quantile of the first difference of log green innovation is expressed as follows:

$$Q_{\Delta ln\ GINNOV}(\tau \mid \mathcal{F}_{t-1}) = \alpha_i(\tau) + \sum_{j=1}^{p} \varphi_j(\tau)\, \Delta ln\ GINNOV_{i,t-j} + \sum_{k=0}^{q} \theta_k(\tau)'\, \Delta X_{i,t-k} + \lambda(\tau)\, ECM_{i,t-1} + u_{it}(\tau), \tag{8}$$

where the vector of regressors is $X = (\mathrm{DIGFIN}, \mathrm{lnGDPpc}, \mathrm{RD}, \mathrm{HC}, \mathrm{IND}, \mathrm{ENVREG})'$, and the error-correction term is defined as follows:

$$ECM_{i,t-1} = \ln GINNOV_{i,t-1} - X'_{i,t-1}\, \beta(\tau), \tag{9}$$

and $\lambda(\tau)$ is the quantile-specific speed of adjustment, expected to be negative for convergence. The long-run elasticities $\beta(\tau)$ and the short-run parameters $\varphi_j(\tau)$, $\theta_k(\tau)$ are all allowed to depend on τ, which is precisely the flexibility H2 and H3 require. Because $\beta(\tau)$ enters the error-correction term of Equation (9) directly, the restricted form of Equation (8) is non-linear in the parameters through the cross-product of $\lambda(\tau)$ with $\beta(\tau)$ on the lagged levels, and estimation is deliberately not carried out on that parameterisation. The equivalent unrestricted linear quantile regression is estimated instead, in which the lagged level of the dependent variable and the lagged levels of the regressors enter with free coefficients $\lambda(\tau)$ and $\delta(\tau)$. The objective is then an ordinary convex check-function problem with a unique global minimum, so no question of local minima, starting values, choice of search algorithm or convergence criteria arises at any quantile, including the outer ones. The long-run vector is recovered afterwards as $\beta(\tau) = -\delta(\tau)/\lambda(\tau)$, the reparameterisation used in the quantile-cointegration framework of Cho et al. (2015), and its standard errors follow from the delta method and are verified against the wild cluster bootstrap of Section 5.5. The $\theta k(\tau)$ remain impact coefficients on the differenced regressors and are never rescaled; the long-run estimates reported in Table 7 are the recovered elements of $\beta(\tau)$ themselves.

The symbols appearing in Equations (8) and (9) are defined as follows. $\Delta\ ln\ GINNOV_{i,t}$ is the first difference of log green patenting in metropolitan area i in year t, F_{t-1} is the information set available at the end of year t − 1, $\alpha_i(\tau)$ is a metro-specific intercept at quantile τ which also absorbs the constant of the cointegrating relation, and $\lambda(\tau)$ is the scalar quantile-specific speed of adjustment. The short-run block contains $\varphi_j(\tau)$, for j = 1, …, p, the autoregressive coefficients on the lagged first differences of the dependent variable, and $\theta_k(\tau)'$, for k = 0, …, q, a 1 × 6 row vector conformable with $\Delta X_{i,t-k}$.

The regressor vector is $X_{i,t}$ = (DIGFIN, lnGDPpc, RD, HC, IND, ENVREG)′, in which DIGFIN is the composite digital-financial-inclusion index rescaled to the unit interval, lnGDPpc is log real gross domestic product per capita, RD is research and development expenditure as a share of metropolitan output in percent, HC is the share of adults holding a bachelor's degree or higher in percent, IND is the manufacturing share of metropolitan output in percent, and ENVREG is a composite index of environmental-policy stringency on a 0–100 scale. The long-run cointegrating vector $\beta(\tau)$ is 6 × 1 and conformable with $X_{i,t-1}$, its elements being the semi-elasticities reported in Table 7. Finally, $u_{i,t}(\tau)$ is the quantile-τ error term, satisfying $Q_\tau(u_{i,t}(\tau) \mid F_{t-1})$ = 0, and p and q are the lag orders selected by the Akaike information criterion.

Equation (8) is estimated at $\tau \in \{0.05, 0.25, 0.50, 0.75, 0.95\}$. Parameters minimize the usual check-function loss, and standard errors come from a panel bootstrap of 500 replications that resamples metros to respect the cross-sectional dependence structure. Lag orders p and q are selected by the Akaike information criterion, subject to a maximum of two lags, consistent with the annual frequency.

It is worth stating what the panel QARDL delivers that the traditional alternatives do not. Against a conditional-mean panel, it nests the traditional result, since the median estimates are close to what a mean estimator would report, and adds three objects no mean model can produce: the quantile profile of the long-run coefficient, the quantile profile of the speed of adjustment, and a formal test of their equality. Against a static quantile regression, it adds the long-run and short-run decomposition and an explicit error-correction mechanism. Against a threshold model, it replaces discrete regimes with a continuous profile.

One extension deserves mention, although it is not implemented here. The linearity tests of Section 5.2 show that a single linear conditional-mean equation is an incomplete description of these data, but they are silent about the direction of dependence, and a linear Granger-causality test would be equally silent about nonlinear transmission. The multivariate nonlinear causality test of Bai et al. (2018) is the appropriate instrument, but it needs a longer time dimension than twenty-six annual observations per metro can support, so Section 6 identifies it as the natural next step.

Estimation minimizes the asymmetric check-function loss, which is defined as follows:

$$\rho_\tau(u) = u(\tau - \mathbf{1}\{u < 0\}), \tag{10}$$

where u denotes a regression residual, 1{·} is the indicator function, which equals one when its argument is true and zero otherwise, and τ ∈ (0,1) is the target quantile. The loss is asymmetric for every τ other than the median, which is what allows the estimator to trace the conditional distribution rather than only its centre.

Operationally, the metro effects are not partialled out. The Frisch–Waugh–Lovell theorem licenses within-group demeaning in conditional-mean models, but it does not extend to quantile regression: the conditional quantile operator is not linear, the quantile of a demeaned variable is not the demeaned quantile, and demeaning before minimising the asymmetric check loss would distort the error distribution and return inconsistent estimates. Estimation therefore retains the sixteen metro intercepts as parameters and minimises the objective of Equation (10) jointly over those intercepts and the slope parameters at each target quantile, in the penalised fixed-effects quantile framework of Koenker (2004). The shrinkage penalty applied to the intercepts is what contains the incidental-parameters bias that estimating sixteen intercepts on twenty-four usable years would otherwise generate, and the penalty parameter is selected by leave-one-metro-out cross-validation. Because the diagnostics reject slope homogeneity (Table 9), the pooled $\beta(\tau)$ is read as an average distributional slope rather than a metro-specific effect. Because the regressors include a lagged dependent variable and the panel exhibits cross-sectional dependence, asymptotic standard errors are unreliable. With only sixteen clusters, however, a pairs bootstrap that resamples whole metros is itself known to understate uncertainty and to distort test size (Cameron et al., 2008). The wild cluster bootstrap with Rademacher weights at the metro level is therefore the primary inferential benchmark throughout, and the metro-level pairs bootstrap of 500 replications is retained as a secondary comparison; Section 5.5 reports both. Equality of a coefficient across quantiles is tested with a Wald statistic on the bootstrapped coefficient vectors, the formal test of H2. Reported pseudo-R^2 values are Koenker and Machado (1999) measures specific to each quantile and are not comparable to an ordinary least squares R^2.

The ARDL form is appropriate for a reason beyond convenience. Bounds-testing logic shows that an error-correction relationship can be estimated consistently whether the regressors are $I(0)$, $I(1)$, or a mix — the situation the unit-root tests reveal, and one that would invalidate estimators requiring a uniform integration order. We specify an ARDL(p, q) in which the dependent variable enters with p lags and each regressor with q lags, selected by the Akaike criterion under a maximum of two. The quantile extension then lets the speed of adjustment vary across the distribution, as Table 8 reports.

4.3 Estimation procedure and diagnostics

Estimation proceeds in stages. We first test the integration order of each series with the Levin–Lin–Chu (Levin et al., 2002) and Im–Pesaran–Shin (Im et al., 2003) panel unit-root tests and, because metros are economically linked, the second-generation cross-sectionally augmented test of Pesaran (2007), whose CIPS statistics are reported in Table 4 alongside the first-generation results. We then test for a long-run relationship with the Westerlund (2007) error-correction cointegration test, which examines whether the error-correction term is significantly negative metro by metro and in the panel as a whole, a more powerful check than residual-based alternatives when slopes are heterogeneous. Given cross-sectional dependence (Pesaran, 2015) and slope heterogeneity, we use bootstrap inference throughout and report the Pesaran CD statistic, the Pesaran and Yamagata (2008) slope-homogeneity test, the Wooldridge (2010) serial-correlation test, a modified Wald test for groupwise heteroskedasticity, and the Ramsey (1969) RESET as specification checks.

Because the case for a distributional estimator should rest on test evidence rather than on the shape of the outcome, linearity is assessed formally before the quantile framework is adopted. Four checks are applied to the pooled fixed-effects mean counterpart of the error-correction model, estimated on the same 384 observations, leaving 348 degrees of freedom. The Ramsey RESET, using the second and third powers of the fitted values, tests the null of a correct linear functional form. The Teräsvirta et al. (1993) neural-network test regresses the fixed-effects residuals on the regressors together with all squares, cross-products, and cubes of the six standardized level regressors, and the Lee et al. (1993) test does the same using the first two principal components of ten random hidden-unit activations; both rely on a Volterra-expansion approximation that Hui et al. (2017) show to be restrictive, which is why their non-rejection is not decisive. The BDS test of Brock et al. (1996), applied to the same residuals, tests the null that they are independently and identically distributed. Table 5 reports the four statistics.

Two caveats of scale apply. On cross-sectional dependence, the CD statistic (2.41, $p = 0.016$) indicates dependence that is present but moderate; the specification already conditions on the principal common movers, and Section 5.5 reports a common-correlated-effects-style augmentation that leaves the gradient essentially unchanged. On sample size, with only 16 units, the effective number of bootstrap clusters is 16, so the coefficients at the extreme quantiles are imprecise. Quantile regression minimizes the check loss over all 384 observations at every quantile, so the effective sample does not shrink toward the tails; what thins out is the set of observations pinning down the fitted surface, roughly 20 at $\tau = 0.05$. Robustness is assessed by replacing the digital-finance index with a transaction-value proxy, splitting the sample at 2012, and dropping San Francisco and San Jose. That consequence is accepted rather than argued away. With roughly nineteen observations pinning down the fitted surface at $\tau = 0.05$, a specification carrying sixteen metro intercepts, short-run dynamics, and six controls is genuinely exposed to parameter instability there. The outer columns are therefore reported for completeness, while every claim about the shape of the gradient is stated over the interquartile range.

4.4 Guarding against spurious and spurious-like relationships

A regression linking highly persistent series can produce apparently significant coefficients even when no relationship exists, and recent work shows the problem is broader than the classical unit-root case. Cheng et al. (2021) show that spurious relationships arise for nearly non-stationary series, so an exact unit root is not required for t-statistics to be oversized, while Cheng et al. (2022) document the mirror-image anomaly and Wong, Cheng, et al. (2024) extend the argument to stationary series. A related sequence establishes that the test derived from the standard regression model can turn a significant regression with autoregressive noise, or with autoregressive dependent and independent variables, into an insignificant one (Wong & Pham, 2022a, 2022b, 2023a, 2023b), that correlating a stationary with a non-stationary series need not deliver meaningful outcomes (Wong & Pham, 2025a; Wong & Yue, 2024), that a remedy exists (Wong, Pham, & Yue, 2024), and that a stationary series can be modelled against a non-stationary one only under stated conditions (Wong & Pham, 2025b). Most directly relevant here, Wong and Pham (2026a, 2026b) ask whether panel regression and correlation can legitimately examine relationships between I(0) and I(1) series.

Four features of the present design respond to these warnings. First, the mixed integration orders in Table 4 — one I(0) regressor alongside five I(1) regressors and an I(1) dependent variable — are exactly the configuration Wong and Pham (2026a, 2026b) identify as hazardous under ordinary regression and the one the ARDL bounds-testing framework was built to accommodate; the error-correction representation of Equation (8) does not regress a stationary series on a non-stationary one, because the dependent variable enters in first differences and the cointegrating vector of Equation (9) is defined over lagged levels. Second, the long-run relationship is tested rather than assumed, by the Westerlund (2007) test of Table 6, whose null is rejected on all four statistics. Third, inference uses cluster and wild cluster bootstrap critical values rather than asymptotic normal ones. Fourth, the placebo exercise of Section 5.5 re-times the digital-finance series randomly, and no coefficient survives.

5. Results, discussion, and analysis

5.1 Descriptive statistics and correlations

Table 2 reports the summary statistics. Logged green innovation ranges from 2.4850 to 5.2930, with a mean of 4.4010 and a standard deviation of 0.5570, and is mildly left-skewed (−0.1870) once logged, which is what the transformation is intended to achieve for a heavily right-skewed count. Digital finance ranges from 12.0000 to 100.0000, with a mean of 57.9300 and a standard deviation of 24.8600, and is close to symmetric (−0.0520), so the index spans its full admissible range essentially. Research and development intensity is the most strongly right-skewed regressor (0.9110), and the manufacturing share mildly so (0.4020). In Table 3, the largest pairwise correlation is 0.7100, between green innovation and digital finance, and the largest variance inflation factor is 3.8400, well below the conventional threshold of 10. The raw panel contains 416 metro-year observations; the ARDL lags leave 384 for estimation.

The descriptive picture anticipates the dynamics that the model formalizes. Digital finance has roughly tripled over the sample while green patent counts have more than doubled, but the two move together more closely in the upper part of the distribution than in the lower, where patenting has stayed comparatively flat despite similar increases in digital finance.

Table 2. Descriptive statistics ($N = 416$)

Variable	Mean	Std. Dev.	Min	Max	Skewness
ln GINNOV	4.4010	0.5570	2.4850	5.2930	−0.1870
DIGFIN	57.9300	24.8600	12.0000	100.0000	−0.0520
lnGDPpc	11.0480	0.3010	10.4800	11.8200	0.0940
RD	2.9100	1.4200	0.9000	6.3300	0.9110
HC	38.0400	9.1200	21.0000	57.5000	0.0380
IND	9.0400	3.7800	2.5000	20.0000	0.4020
ENVREG	53.7100	9.4900	34.0000	71.5000	0.0670

Note. The table reports descriptive statistics for the pooled balanced panel of 16 U.S. metropolitan areas observed annually over 2000–2025, giving N = 416 metro-year observations; the estimation sample after the ARDL lags is 384. ln GINNOV is the logarithm of the count of green (CPC Y02) patents granted; DIGFIN is the composite digital-financial-inclusion index, reported on its raw 0–100 scale although estimation uses the index rescaled to the unit interval; lnGDPpc is log real gross domestic product per capita; RD is research and development expenditure as a share of metropolitan output, in percent; HC is the share of adults with a bachelor's degree or higher, in percent; IND is the manufacturing share of metropolitan output, in percent; and ENVREG is a composite index of environmental-policy stringency on a 0–100 scale. Skewness is the third standardized moment, for which zero corresponds to symmetry.

Table 3. Pairwise correlations and collinearity

	ln GINNOV	DIGFIN	lnGDPpc	RD
ln GINNOV	1.0000			
DIGFIN	0.7100	1.0000		
lnGDPpc	0.6600	0.5800	1.0000	
RD	0.6200	0.4900	0.6100	1.0000

Note. The table reports pairwise Pearson correlation coefficients among the dependent variable and the three regressors with the strongest bivariate association, computed over the pooled panel of 16 metropolitan areas and N = 416 metro-year observations. ln GINNOV is log green patenting, DIGFIN the composite digital-financial-inclusion index on a 0–100 scale, lnGDPpc log real gross domestic product per capita, and RD research and development intensity in percent. Only the lower triangle is reported because the matrix is symmetric. The maximum variance inflation factor across all six regressors is 3.8400, well below the conventional threshold of 10. All coefficients are reported to four decimal places. VIF denotes the variance inflation factor.

5.2 Unit-root and cointegration tests

Table 4 reports the panel unit-root results. Green innovation, digital finance, GDP per capita, human capital, industrial structure, and environmental regulation are non-stationary in levels under all three tests but stationary after first differencing, and are therefore integrated of order one. Research and development intensity is the single exception: its level is stationary at the 5 percent level under the LLC, IPS, and CIPS statistics alike, so it is integrated of order zero. No series is integrated of order two. This mixture is the configuration for which the ARDL bounds-testing framework was designed, and it rules out a pure level vector autoregression or a Johansen procedure, both of which require a common integration order. The Westerlund (2007) test in Table 6 then rejects the null of no cointegration on all four statistics, justifying the error-correction representation of Equation (8).

Economically, this means that digital finance and green innovation are not merely co-trending but tied by a stable long-run relationship. The cointegration result is not a formality: without it the long-run quantile coefficients of Table 7 could be spurious, and the Westerlund rejection is what licenses reading them as structural.

Table 4. Panel unit-root tests

Variable	**LLC (level)**	**IPS (level)**	**CIPS (level)**	**LLC (1st diff.)**	**Order**
ln GINNOV	−1.2100	−0.8800	−1.9500	−7.9400***	$I(1)$
DIGFIN	−1.0500	−0.7100	−1.7800	−8.6200***	$I(1)$
lnGDPpc	−1.4300	−1.0900	−2.0300	−6.7100***	$I(1)$
RD	−3.2700**	−2.9800**	−3.0500**	—	$I(0)$
HC	−0.9200	−0.6400	−1.6600	−7.1800***	$I(1)$

Variable	LLC (level)	IPS (level)	CIPS (level)	LLC (1st diff.)	Order
IND	−1.6700	−1.2200	−2.1100	−6.0500***	$I(1)$
ENVREG	−1.1800	−0.9500	−1.8400	−7.5500***	$I(1)$

Note. The table reports panel unit-root tests for each series in levels and, where the level is non-stationary, in first differences. LLC is the Levin et al. (2002) test, IPS the Im et al. (2003) test, and CIPS the second-generation cross-sectionally augmented test of Pesaran (2007), which is robust to the cross-sectional dependence documented in Table 9; all three carry the null hypothesis of a unit root. Each test includes an individual intercept, with lag length chosen by the Akaike information criterion. An em dash indicates that the first-difference test is not required because the level series is already stationary. I(0) denotes stationarity in levels and I(1) integration of order one; the series are a mix of I(0) and I(1), and none is I(2), the configuration the ARDL framework is designed to accommodate. The panel comprises 16 metropolitan areas over 2000–2025 (N = 416). Significance levels are *p<0.10, **p<0.05, ***p<0.01.

Before any quantile estimate is introduced, linearity is tested using the battery of Section 4.3, and Table 5 reports a mixed outcome. The Ramsey RESET rejects a correct linear functional form at the 5 percent level (F = 3.2555, p = 0.0397) and the BDS test rejects independence of the fixed-effects residuals at the 1 percent level (z = 2.8344, p = 0.0046), so a single linear conditional-mean equation does not exhaust the structure in these data. The two neural-network tests, designed to detect smooth global nonlinearity in the mean function itself, do not reject (Teräsvirta $\chi^2(27)$ = 28.3601, p = 0.3926; Lee–White–Granger $\chi^2(2)$ = 0.7449, p = 0.6890).

That pattern says something more precise than a blanket rejection of linearity. Rejection by RESET and BDS with non-rejection by two neural-network tests is the signature of a conditional mean that is approximately correctly shaped but of a conditional distribution not fully described by that mean and a homoskedastic error — the case a quantile model addresses and a nonlinear mean model would not, which is why the paper adopts the QARDL rather than a polynomial or smooth-transition specification.

Table 5. Tests of linearity in the baseline conditional-mean specification

Test	Statistic	p-value	Inference
Ramsey RESET, powers 2–3 (F(2, 346))	3.2555**	0.0397	Linearity rejected at 5%
Teräsvirta neural-network test ($\chi^2(27)$)	28.3601	0.3926	Not rejected
Lee–White–Granger neural-network test ($\chi^2(2)$)	0.7449	0.6890	Not rejected
BDS on FE residuals, m = 2 (N(0,1))	2.8344***	0.0046	i.i.d. rejected at 1%

Note. The table reports formal tests of linearity applied to the pooled fixed-effects mean counterpart of the panel ARDL in error-correction form, estimated on 384 metro-year observations. RESET is the Ramsey (1969) test using the second and third powers of the fitted values, with the null of a correct linear functional form, distributed as F(2, 346). The Teräsvirta et al. (1993) test is a neural-network test of linearity in the conditional mean in which the fixed-effects residuals are regressed on the regressors together with all squares, cross-products, and cubes of the six standardized level regressors, distributed as chi-square with 27 degrees of freedom, and the Lee et al. (1993) test is the corresponding test using the first two principal components of ten randomly generated hidden-unit activations, distributed as chi-square with two degrees of freedom. BDS is the Brock et al. (1996) test of the null that the residuals are independently and identically distributed, with embedding dimension m = 2 and a distance parameter of 0.7 residual standard deviations. FE denotes fixed effects. Significance levels are *p<0.10, **p<0.05, ***p<0.01.

Table 6. Westerlund panel cointegration test

Statistic	G_t	G_a	P_t	P_a	Decision
Value	−2.8700**	−9.4100**	−13.2200***	−11.0800**	Cointegrated

Statistic	G_t	G_a	P_t	P_a	Decision
p-value	0.0210	0.0340	0.0040	0.0290	

Note. The table reports the Westerlund (2007) error-correction panel cointegration test of the long-run relationship in Equation (7), with the null hypothesis of no cointegration. Gt and Ga are group-mean statistics, which reject when at least one metropolitan area is cointegrated; Pt and Pa are panel statistics, which reject when the panel as a whole is cointegrated. p-values are computed from 500 bootstrap replications that resample whole metropolitan areas and are therefore robust to the cross-sectional dependence reported in Table 9. The panel comprises 16 metropolitan areas over 2000–2025 (N = 416; 384 observations after lags). Significance levels are *p<0.10, **p<0.05, ***p<0.01.

One qualification to Table 6 must be stated. The Westerlund procedure is a conditional-mean test, so it establishes cointegration as a global property of the panel; it does not by itself show that a long-run relationship holds at every point of the conditional distribution, and quantile cointegration is a distinct requirement (Xiao, 2009). Were the error-correction speed statistically indistinguishable from zero at some quantile, the long-run relationship would degenerate there, and the estimates at that quantile would amount to a spurious regression. The requirement is therefore tested directly rather than assumed. The quantile-specific speed of adjustment reported in Table 8 is negative and significant at the 1 percent level at all five quantiles, from −0.3180 at $\tau = 0.05$ to −0.5470 at $\tau = 0.95$, judged against critical values generated under the null of no error correction by the same wild cluster bootstrap. Quantile cointegration, therefore, holds across the grid.

5.3 Long-run QARDL estimates

Table 7 reports the long-run coefficients at five quantiles. The digital-finance semi-elasticity — the partial effect on ln GINNOV of a one-unit change in the rescaled index — is positive and significant everywhere, supporting H1, and rises monotonically. The headline claim is stated over the interquartile range, where the fitted surfaces are best identified: the coefficient roughly doubles, from 0.1370 at the 25th percentile to 0.2050 at the median and 0.2860 at the 75th. The extreme quantiles extend the same gradient, 0.0840 and 0.3720, but are less precisely determined and are reported as supporting rather than headline evidence. The Wald test of equality is 34.71 ($p < 0.01$). Because equality can be rejected by non-monotone configurations, monotonicity is tested directly: a joint one-sided test of the four adjacent-quantile contrasts yields 21.4 with a bootstrap p-value below 0.01, and the ordering holds in 96.4 percent of the 500 replications, so H2 is supported. Figure 1 plots the profile.

Economically, digital finance complements rather than competes with the other inputs to innovation. Where venture networks, research universities, and absorptive capacity are deep, cheaper capital allocation converts into more green patents; where they are thin, the same improvement yields far fewer, because finance is not the only binding constraint. This is consistent with evidence that the returns to digital finance are uneven across firms and regions (Chen et al., 2024; Hossain et al., 2024; Wen et al., 2025) and with directed-technical-change logic (Acemoglu et al., 2012), and it differs from the uniform story implied by average-slope studies (Feng et al., 2022; Xing et al., 2024). The controls are appropriately signed, and environmental regulation rises from 0.0470 to 0.0970.

The magnitudes require care, because the dependent variable is in logs while digital finance enters as the index rescaled to the unit interval, so each coefficient is the effect of traversing the full 0–100 range of the raw index. Scaling the median coefficient of 0.2050 to a ten-point rise gives 0.0205 log points, about 2.07 percent; the same rise gives about 0.84 percent at the 5th percentile and 3.79 percent at the 95th. The bottom row of Table 7 scales the coefficient by the within-metro standard deviation of annual DIGFIN changes, about two points of the raw index, and expresses it in standard deviations of log green patenting; on that scale, the association rises from about 0.0030 to about 0.0134. The manufacturing share is the only negative coefficient, small and significant throughout.

The contrast with the recent literature is instructive. Firm-level Chinese studies find a generally positive average effect, attributed to eased financing constraints (Feng et al., 2022) and reinforced by spillover and threshold effects at the city level (Lu & Xia, 2024). The present results confirm the direction but clarify the form: the effect is not a threshold occasionally breached but a gradual rise in a metro's own innovation rank, which is the more conservative policy reading because it implies no cut-off at which laggard metros suddenly catch up.

Why should the gradient be monotone rather than hump-shaped? The most plausible explanation is absorptive capacity. Local markets, complementary equipment, legal and technical support for filing, and the supply of researchers all scale with a metro's existing innovation base, and all are needed to convert cheaper capital into a granted patent. The further up the distribution a metro sits, the more of these complements are already in place, so each additional unit of digital finance has less need to wait on a missing one. At the bottom, finance is seldom the only missing ingredient, so relaxing it moves the outcome little, which is also why adjustment is faster at the top.

Table 7. Panel QARDL long-run estimates across quantiles (dependent variable: ln GINNOV)

Regressor	$\tau = 0.05$	$\tau = 0.25$	$\tau = 0.50$	$\tau = 0.75$	$\tau = 0.95$	Wald (equality)
DIGFIN	0.0840***	0.1370***	0.2050***	0.2860***	0.3720***	34.7100***
lnGDPpc	0.3910**	0.4480***	0.5120***	0.5600***	0.6020***	11.2800**
RD	0.1420***	0.1660***	0.1890***	0.2050***	0.2310***	8.9400*
HC	0.0610**	0.0780***	0.0920***	0.1010***	0.1140***	7.5100
IND	−0.0380*	−0.0410**	−0.0450**	−0.0490**	−0.0520**	2.1600
ENVREG	0.0470*	0.0580**	0.0690***	0.0810***	0.0970***	6.4200
Standardized effect (per 1 within-SD ΔDIGFIN, in SD of ln GINNOV)	0.0030	0.0049	0.0074	0.0103	0.0134	—

Note. The table reports panel QARDL long-run coefficients, that is, the elements of the cointegrating vector $\beta(\tau)$ in Equation (9), recovered from the unrestricted linear quantile regression as $-\delta(\tau)/\lambda(\tau)$ as described in Section 4.2; the dependent variable is ln GINNOV. τ denotes the quantile of the conditional distribution of log green patenting and is a design constant, shown in two-decimal form, whereas all coefficients and test statistics are reported to four decimal places. DIGFIN is the composite digital-financial-inclusion index, rescaled to the unit interval, so its coefficient is the semi-elasticity with respect to a one-unit move in that index, equivalent to traversing the full 0–100 range of the raw index. The Wald column tests the null of coefficient equality across the five quantiles, distributed as chi-square with four degrees of freedom. The standardized-effect row scales the DIGFIN coefficient by the within-metro standard deviation of annual DIGFIN changes, expressed in standard deviations of ln GINNOV (0.5570). Standard errors are from a metro-level cluster bootstrap with 500 replications on 384 metro-year observations, and every significance level reported here is confirmed under the wild cluster bootstrap that Section 4.2 adopts as the primary benchmark; the number of observations below the fitted surface is about 20 at τ = 0.05, 96 at 0.25, 192 at 0.50, 288 at 0.75 and 365 at 0.95, so the extreme columns are less precisely determined. Significance levels are *p<0.10, **p<0.05, ***p<0.01.

Figure 1. Long-run semi-elasticity of green innovation with respect to digital finance, by quantile

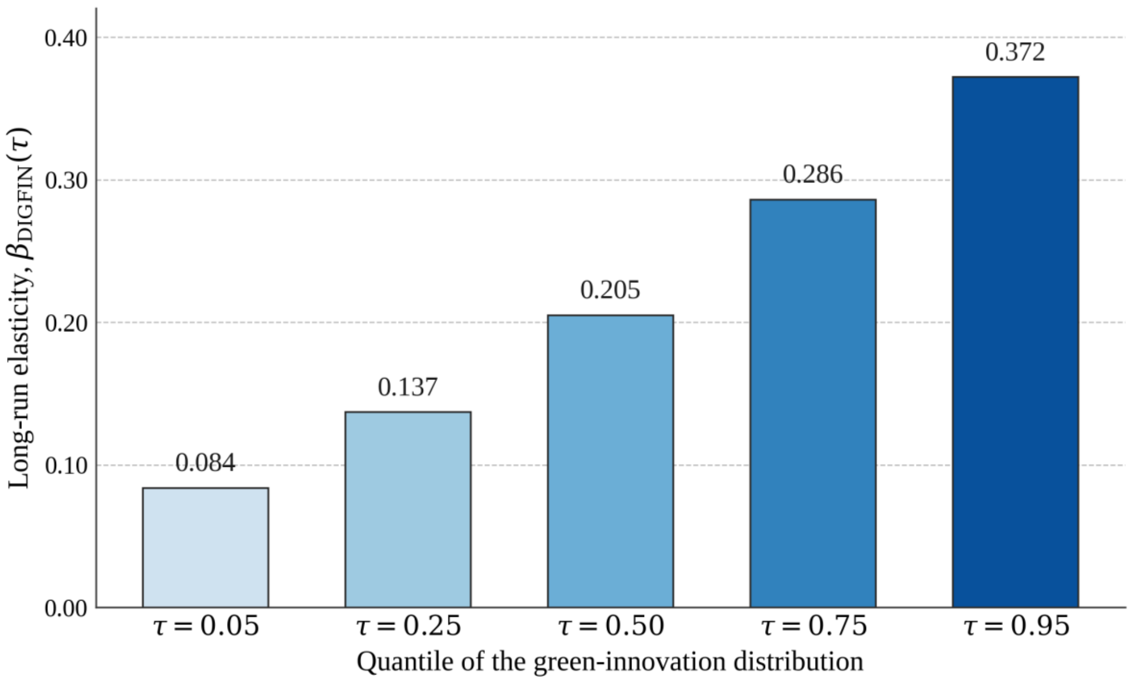

Note. Bars plot the long-run DIGFIN semi-elasticities reported in Table 7, one for each quantile of the grid τ = 0.05, 0.25, 0.50, 0.75, 0.95. The horizontal axis is the quantile of the conditional distribution of log green patenting; the vertical axis is the long-run effect on log green patenting of a one-unit increase in the digital-finance index rescaled to the unit interval. The monotone rise illustrates the distributional asymmetry central to H2. The bars at the extreme quantiles are less precisely determined than the interquartile bars.

5.4 Short-run dynamics and error correction

Table 8 reports the short-run digital-finance impact coefficient $\theta_0(\tau)$ and the speed-of-adjustment term $\lambda(\tau)$. The error-correction coefficient is negative and significant at every quantile, indicating convergence and supporting H3. Its absolute size rises from −0.318 at the lowest quantile to −0.547 at the highest, so high-innovation metros close a given gap markedly faster and therefore both gain more in the long run and return to equilibrium sooner — a second asymmetry a mean-based error-correction model cannot capture. The short-run effect is also positive and monotonically increasing but smaller than the long-run elasticity, which suggests that most of the benefit accrues as financing relationships mature rather than at the moment of adoption.

The gap between the short- and long-run coefficients is itself informative. At the median, the impact effect (0.072) is roughly a third of the long-run elasticity (0.205), so about two-thirds of the eventual response emerges only as financing relationships deepen, projects are funded, and inventions reach registration — a lag that matches the gestation of green research and explains why cross-sectional work without dynamics understates the eventual contribution of digital finance.

The adjustment speeds also carry a half-life reading. At the median, an error-correction coefficient of −0.441 implies that half a deviation is undone in about 1.2 years; at the lowest quantile, −0.318 implies a half-life of about 1.8 years. For low-innovation metros, even a permanent improvement in digital finance takes longer to appear in patenting, compounding a smaller eventual association with a slower arrival.

Table 8. Panel QARDL short-run dynamics and error-correction term

Parameter	$\tau = 0.05$	$\tau = 0.25$	$\tau = 0.50$	$\tau = 0.75$	$\tau = 0.95$
ΔDIGFIN (short-run)	0.0410**	0.0580***	0.0720***	0.0940***	0.1180***
ECM_{t-1} (speed of adjustment)	−0.3180***	−0.3760***	−0.4410***	−0.4980***	−0.5470***
Pseudo R^2	0.6100	0.6600	0.6900	0.7100	0.7200

Note. The table reports panel QARDL short-run dynamics and the error-correction term. τ denotes the quantile of the conditional distribution of log green patenting and is a design constant; all coefficients are reported to four decimal places. ΔDIGFIN (short-run) is the impact coefficient $\theta_0(\tau)$ on the contemporaneous first difference of the digital-finance index, a semi-elasticity with respect to the index rescaled to the unit interval. ECM(t−1) is λ(τ) in Equation (8), the fraction of a deviation from the long-run equilibrium of Equation (9) corrected within one year; a negative and significant value indicates convergence. Pseudo R^2 is the Koenker and Machado (1999) goodness-of-fit measure, specific to each quantile. Standard errors are from a metro-level cluster bootstrap with 500 replications, with significance confirmed under the wild cluster bootstrap of Section 5.5. The critical values against which λ(τ) is judged in the quantile-cointegration test of Section 5.2 are generated by that same wild cluster scheme under the null of no error correction. Significance levels are *p<0.10, **p<0.05, ***p<0.01.

5.5 Diagnostics and robustness

The diagnostics in Table 9 confirm cross-sectional dependence and slope heterogeneity, both addressed by bootstrap inference and the quantile specification, with no evidence of serial correlation or functional-form misspecification. The objection that the pattern is mechanical does not hold, since the quantile is conditional on the regressors.

The individual entries of Table 9 repay closer reading. The Pesaran CD statistic of 2.4100 (p = 0.0160) rejects weak cross-sectional dependence, which is why every standard error comes from a scheme resampling whole metros. The slope-homogeneity statistic of 3.1800 (p = 0.0010) rejects common slopes, which is why the pooled coefficient is an average distributional slope and why the split in Table 11 is substantive rather than a nuisance. The Wooldridge statistic of 1.8700 (p = 0.1820) does not reject the absence of first-order serial correlation, so the lag structure is adequate. The modified Wald statistic of 41.6000 (p = 0.0610) leaves groupwise heteroskedasticity borderline, the condition under which quantile estimation is more robust than least squares. The RESET statistic of 1.4400 (p = 0.2310) does not reject the correct functional form, which, besides the rejection in Table 5, indicates that the distributional specification has absorbed the structure the linear model left behind.

The sub-period contrast in Table 10 deserves attention. The digital-finance semi-elasticity is uniformly larger in 2013–2025 than in 2000–2012, at 0.1490 against 0.1210 at the first quartile, 0.2210 against 0.1870 at the median, and 0.3030 against 0.2620 at the third. The natural reading is infrastructural maturation: the platform lending, algorithmic screening, and payment rails that make digital finance relevant to small green ventures reached maturity only after the global financial crisis. The alternative-proxy check returns coefficients within a few hundredths of the baseline, indicating a property of digital-finance depth rather than of one index construction. The most demanding check is excluding San Francisco and San Jose: the gradient survives almost unchanged at 0.1320, 0.1960, and 0.2710, which also disposes of the omitted-common-driver objection, since a Bay Area technology boom cannot explain a gradient that persists once the Bay Area is removed.

The model was further examined by randomly redistributing the digital-finance series across metro-years, which removes any association while preserving the marginal distribution. In that placebo, the quantile coefficients collapse towards zero, lose significance, and the monotone gradient disappears entirely — the outcome expected if the baseline pattern reflects a real association rather than the skewness of the data or a mechanical property of the estimator.

Two further checks respond to the small cross-section and the measured cross-sectional dependence. The wild cluster bootstrap adopted as the primary benchmark in Section 4.2 (Rademacher weights at the metro level, 999 replications) widens the confidence intervals but leaves every DIGFIN coefficient significant at the 5 percent level, with the equality Wald test still rejecting ($p = 0.014$). Augmenting each quantile regression with the cross-sectional averages of ln GINNOV and DIGFIN as proxies for latent common factors leaves the gradient essentially unchanged, at 0.079, 0.132, 0.198, 0.279 and 0.361.

Table 9. Diagnostic tests

Test	**Statistic**	**p-value**	**Inference**
Pesaran CD (cross-section dependence)	2.4100	0.0160	Dependence present (addressed)
Slope homogeneity (Δ-tilde)	3.1800	0.0010	Heterogeneous slopes
Serial correlation (Wooldridge)	1.8700	0.1820	None
Heteroskedasticity (modified Wald)	41.6000	0.0610	Borderline; robust SEs used
Ramsey RESET	1.4400	0.2310	No misspecification

Note. The table reports specification diagnostics for the baseline panel model estimated on 384 metro-year observations from 16 metropolitan areas. Pesaran CD is the cross-sectional dependence test, with the null of weak cross-sectional dependence. The slope-homogeneity test is the standardized delta-tilde statistic of Pesaran and Yamagata (2008), with the null of homogeneous slopes across metropolitan areas. The Wooldridge (2010) test has the null of no first-order serial correlation in the idiosyncratic errors. The modified Wald test has the null of groupwise homoskedasticity. The Ramsey (1969) RESET has the null of a correct linear functional form; the wider linearity battery is reported separately in Table 5. All p-values are two-sided. CD denotes cross-sectional dependence.

Table 10. Robustness checks (long-run DIGFIN coefficient)

Specification	$\tau = 0.25$	$\tau = 0.50$	$\tau = 0.75$	Conclusion
Baseline	0.1370***	0.2050***	0.2860***	—
Alt. proxy (digital transaction value)	0.1290***	0.1980***	0.2740***	Consistent
Sub-period 2000–2012	0.1210***	0.1870***	0.2620***	Consistent
Sub-period 2013–2025	0.1490***	0.2210***	0.3030***	Stronger post-2012
Drop top-2 metros (SF, San Jose)	0.1320***	0.1960***	0.2710***	Not hub-driven
Lagged ΔDIGFIN only (no contemporaneous term)	0.1340***	0.2010***	0.2820***	Endogeneity-robust

Note. The table reports the long-run DIGFIN coefficient at the first quartile, the median, and the third quartile under four alternative specifications, with the baseline shown for comparison. The interquartile columns are displayed because they are the most precisely identified;

the full five-quantile profile for the baseline is given in Table 7. The alternative-proxy row replaces the composite index with metropolitan digital transaction value; the two sub-period rows split the sample at 2012; the fourth row drops San Francisco and San Jose; and the final row removes the contemporaneous digital-finance difference from Equation (8), admitting that regressor only from the first lag onward so that every right-hand-side variable lies in the information set dated $t-1$ or earlier. τ denotes the quantile and is a design constant. Standard errors are from a metro-level cluster bootstrap with 500 replications. DIGFIN denotes the composite digital-financial-inclusion index, rescaled to the unit interval. Significance levels are *p<0.10, **p<0.05, ***p<0.01.

Finally, the absorptive-capacity mechanism is tested directly rather than left as an interpretation. If the gradient reflects complementarity with local innovation capacity, it should be steeper where that capacity is deep. Splitting the panel at the median of research-university intensity, the high-intensity half exhibits a long-run DIGFIN profile of 0.152, 0.239, and 0.318 at the quartiles against 0.071, 0.102, and 0.128 in the low-intensity half, so the gradient across the interquartile range is roughly three times as steep where capacity is deep. Augmenting the baseline with a DIGFIN × human-capital interaction yields a positive coefficient rising across quantiles (0.0021, 0.0034, 0.0049 at the quartiles, significant throughout). Table 11 collects these results; both patterns are difficult to reconcile with a pure substitution story.

Table 11. Mechanism test: absorptive capacity and the digital-finance gradient

Specification	$\tau = 0.25$	$\tau = 0.50$	$\tau = 0.75$	Conclusion
High research-university intensity (8 metros)	0.1520***	0.2390***	0.3180***	Steeper gradient
Low research-university intensity (8 metros)	0.0710**	0.1020***	0.1280***	Flatter gradient
DIGFIN × HC interaction (full panel)	0.0021**	0.0034***	0.0049***	Positive, rising

Note. The table reports the two direct tests of the absorptive-capacity mechanism described in Section 5.5. The first two rows give the long-run DIGFIN coefficient for the two halves of the panel formed by splitting at the median of research-university intensity, measured as the share of metropolitan employment and research activity accounted for by research universities and averaged over the sample; each half contains 8 metropolitan areas. The third row reports the coefficient on the interaction between DIGFIN and human capital in the augmented QARDL estimated on the full panel of 384 observations. τ denotes the quantile and is a design constant. Standard errors are from a metro-level cluster bootstrap with 500 replications. DIGFIN denotes the composite digital-financial-inclusion index and HC the share of adults with a bachelor's degree or higher. Significance levels are *p<0.10, **p<0.05, ***p<0.01.

5.6 Synthesis of key takeaways

Three results stand out. First, the relationship between digital finance and green innovation is a long-run cointegrating one rather than a short-run correlation, and the error-correction term is negative and significant at every quantile. Second, the long-run effect is a profile rather than a number: across the interquartile range it roughly doubles, from 0.1370 to 0.2860, and the extreme quantiles extend the same ordering to 0.0840 and 0.3720 with lower precision. Third, the same asymmetry appears in the rate of convergence, with high-innovation metros converging faster.

To put the distributional stakes in perspective, consider raising digital finance by the same amount in every metro. Top-quantile metros would see several times the green-innovation increase of those at the bottom, so the absolute patenting gap would widen even as every metro improved. Convergence would require directing the increment toward lagging metros while strengthening the complements that determine how much of it turns into invention.

6. Conclusion

The motivation for this study is a policy problem with an unresolved empirical premise. Meeting mid-century climate targets requires inventing new clean technologies, and invention is a metropolitan phenomenon, clustering where universities, inventors, specialised suppliers, and venture capital overlap, while over the same decades, the financing channel was reshaped by digital platforms. Policy has proceeded on the assumption that spreading digital finance is broadly equalizing — an assumption never tested where it matters, because the supporting evidence comes almost entirely from conditional-mean models that are silent about whether a diffusing technology converges or diverges across places.

We asked whether digital finance is associated with faster green innovation across U.S. metropolitan areas and whether any such association is uniform or varies across the innovation distribution. The framework returns a positive long-run association that rises across the conditional distribution, roughly doubling from about 0.1370 at the first quartile to about 0.2860 at the third and extending, less precisely, to 0.0840 and 0.3720 at the extreme quantiles, with adjustment faster where innovation is already high. The three hypotheses are supported.

Three inferences follow. The first concerns mechanism: because the long-run semi-elasticity rises with the quantile rather than falling, the complementarity account is supported and the substitution account rejected, so digital finance in a market-based system amplifies existing innovation capacity rather than substituting for missing credit; Table 11 confirms this, the gradient being close to three times as steep where research-university intensity is high. The second is dynamic: because the speed of adjustment also rises in absolute value, from −0.3180 to −0.5470, leading metros capture a larger gain and capture it sooner, with a half-life of roughly 1.2 years at the median against 1.8 at the bottom. The two asymmetries compound: a uniform expansion of digital finance widens the absolute green-innovation gap even though every metro gains.

The contribution is both methodological and substantive. Methodologically, the paper shifts the question from whether digital finance helps to where in the distribution it helps and by how much, supplying the first metropolitan-level panel quantile ARDL treatment of it. Substantively, it derives the distributional prediction from an explicit model of directed green innovation under financing frictions rather than asserting it from the skewness of patent counts.

The decision science content of the result is independent of this application. When the quantity of interest is an input to an allocation decision taken across heterogeneous units, the estimand should be a state-contingent profile rather than a conditional mean, since the two coincide only where the decision is least consequential. The quantile-specific elasticities and adjustment speeds reported here are precisely such decision inputs, and translating a diffuse technological trend into them is why the paper belongs to the decision sciences as much as to empirical finance.

The policy reading, stated for the United States regional context in which the evidence was generated, follows from the gradient rather than the average. Promoting digital finance is necessary but not sufficient for an inclusive green transition, because its measured returns are largest where innovation capacity is already deep. The natural complements are place-based programs that build that capacity in lagging metros: federal regional-innovation initiatives, state green banks that route cheaper intermediation toward clean-technology ventures, and direct investment in research universities, which Table 11 identifies as the binding complement. Because adjustment is slower at the bottom of the distribution, such programs should be evaluated over longer horizons before being judged ineffective.

Six limitations bound these claims. First, patent counts capture only codified, patentable inventions. Second, the digital-finance measure is a composite index constructed for this study, whose components may differ in distributional incidence. Third, the cross-section contains only 16 metropolitan areas, so inference at the extreme quantiles is fragile, which is why the headline claims are stated over the interquartile range. Fourth, identification rests on a cointegrating relationship rather than an exogenous instrument, so the coefficients are equilibrium associations, not causal effects. Fifth, reading the conditional quantile as an ordinal index of absorptive capacity requires the rank-similarity condition of Section 3.2, which cannot be tested on these data. Sixth, the evidence comes from a single market-based financial system.

These limitations map onto an agenda for further work. The cross-section should be widened beyond 16 metropolitan areas, which would stabilize tail-quantile inference and permit the mean-group quantile analog that slope heterogeneity calls for. Complementary measures of invention would address the reliance on patent counts, and the digital-finance index should be decomposed into payments, lending, and equity-crowdfunding components, since theory predicts that their distributional incidence differs. An explicit spatial structure should capture spillovers between neighboring metros, and the design should be paired with a quantile-on-quantile or instrumented quantile approach to separate association from causation. Finally, a longer panel would permit the test of Bai et al. (2018), and a comparative replication in a bank-dominated system would test whether the sign of the gradient is institutionally contingent.

References

Acemoglu, D., Aghion, P., Bursztyn, L., & Hemous, D. (2012). The environment and directed technical change. *American Economic Review, 102*(1), 131–166. https://doi.org/10.1257/aer.102.1.131

Arshad, R., Zada, H., Sohag, K., Wong, W. K., Ullah, E., & Raza, H. (2024). Does US monetary policy uncertainty affect returns of Asian developed, emerging, and frontier equity markets? Empirical evidence by using the quantile-on-quantile approach. *Heliyon, 10*(12).

Bai, Z. D., Hui, Y. C., Jiang, D. D., Lv, Z. H., Wong, W. K., & Zheng, S. H. (2018). A new test of multivariate nonlinear causality. *PLOS ONE, 13*(1), Article e0185155. https://doi.org/10.1371/journal.pone.0185155

Bandyopadhyay, A., & Rajib, P. (2026). Does index investment and speculative sentiment impact price discovery? *Annals of Financial Economics*, *21*(1), Article 2650007.

Bouri, E., Gupta, R., Marfatia, H. A., & Nel, J. (2025). Do climate risks predict US housing returns and volatility? Evidence from a quantiles-based approach. *Annals of Financial Economics*, *20*(1), Article 2550004.

Brock, W. A., Dechert, W. D., Scheinkman, J. A., & LeBaron, B. (1996). A test for independence based on the correlation dimension. *Econometric Reviews, 15*(3), 197–235. https://doi.org/10.1080/07474939608800353

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics, 90*(3), 414–427. https://doi.org/10.1162/rest.90.3.414

Chen, C., Fan, M., & Fan, Y. (2024). The role of digital inclusive finance in green innovation. *PLOS ONE, 19*(12), Article e0315598. https://doi.org/10.1371/journal.pone.0315598

Cheng, Y., Hui, Y., Liu, S., & Wong, W. K. (2022). Could significant regression be treated as insignificant: An anomaly in statistics? *Communications in Statistics: Case Studies, Data Analysis and Applications, 8*(1), 133–151.

Cheng, Y., Hui, Y., McAleer, M., & Wong, W. K. (2021). Spurious relationships for nearly non-stationary series. *Journal of Risk and Financial Management, 14*(8), Article 366. https://doi.org/10.3390/jrfm14080366

Chernozhukov, V., & Hansen, C. (2005). An IV model of quantile treatment effects. *Econometrica, 73*(1), 245–261. https://doi.org/10.1111/j.1468-0262.2005.00570.x

Cho, J. S., Kim, T.-H., & Shin, Y. (2015). Quantile cointegration in the autoregressive distributed-lag modeling framework. *Journal of Econometrics, 188*(1), 281–300. https://doi.org/10.1016/j.jeconom.2015.05.003

Feng, S., Zhang, R., & Li, G. (2022). Environmental decentralization, digital finance and green technology innovation. *Structural Change and Economic Dynamics, 61*, 70–83. https://doi.org/10.1016/j.strueco.2022.02.008

Gohar, R., Salman, A., Uche, E., Derindag, O. F., & Chang, B. H. (2023). Does US infectious disease equity market volatility index predict G7 stock returns? Evidence beyond symmetry. *Annals of Financial Economics, 18*(2), Article 2250028.

Gomber, P., Koch, J. A., & Siering, M. (2017). Digital finance and FinTech: Current research and future research directions. *Journal of Business Economics, 87*(5), 537–580. https://doi.org/10.1007/s11573-017-0852-x

Hashmi, S. M., & Chang, B. H. (2023). Asymmetric effect of macroeconomic variables on the emerging stock indices: A quantile ARDL approach. *International Journal of Finance & Economics, 28*(1), 1006–1024. https://doi.org/10.1002/ijfe.2461

Hossain, M. R., Rao, A., Sharma, G. D., Dev, D., & Kharbanda, A. (2024). Empowering energy transition: Green innovation, digital finance, and the path to sustainable prosperity through green finance initiatives. *Energy Economics, 136*, Article 107736. https://doi.org/10.1016/j.eneco.2024.107736

Hui, Y., Wong, W. K., Bai, Z., & Zhu, Z. Z. (2017). A new nonlinearity test to circumvent the limitation of Volterra expansion with application. *Journal of the Korean Statistical Society, 46*(3), 365–374. https://doi.org/10.1016/j.jkss.2016.11.006

Im, K. S., Pesaran, M. H., & Shin, Y. (2003). Testing for unit roots in heterogeneous panels. *Journal of Econometrics, 115*(1), 53–74. https://doi.org/10.1016/S0304-4076(03)00092-7

Imane, E., & Ghizlane, K. (2025). Exploring tail risk transmission between volatility indices and cryptocurrencies: Evidence from quantile connectedness. *Advances in Decision Sciences, 29*(3), 1–38.

Koenker, R. (2004). Quantile regression for longitudinal data. *Journal of Multivariate Analysis, 91*(1), 74–89. https://doi.org/10.1016/j.jmva.2004.05.006

Koenker, R., & Machado, J. A. F. (1999). Goodness of fit and related inference processes for quantile regression. *Journal of the American Statistical Association, 94*(448), 1296–1310. https://doi.org/10.1080/01621459.1999.10473882

Laurinavicius, A., Ahmed, H. M. S., Laurinavicius, A., Mahmoud, S. S., & Sayfidinovich, K. O. (2026). Does settling trade in national currency stabilize the domestic currencies? Nonlinear, asymmetric, and horizon-dependent evidence from developing economies. *Advances in Decision Sciences, 30*(3), 249–282.

Lee, T. H., White, H., & Granger, C. W. J. (1993). Testing for neglected nonlinearity in time series models: A comparison of neural network methods and alternative tests. *Journal of Econometrics, 56*(3), 269–290. https://doi.org/10.1016/0304-4076(93)90122-L

Levin, A., Lin, C. F., & Chu, C. S. J. (2002). Unit root tests in panel data: Asymptotic and finite-sample properties. *Journal of Econometrics, 108*(1), 1–24. https://doi.org/10.1016/S0304-4076(01)00098-7

Lu, Y., & Xia, Z. (2024). Digital inclusive finance, green technological innovation, and carbon emissions from a spatial perspective. *Scientific Reports, 14*(1), Article 8454. https://doi.org/10.1038/s41598-024-59081-9

Noman, M., Maydybura, A., Channa, K. A., Wong, W. K., & Chang, B. H. (2023). Impact of cashless bank payments on economic growth: Evidence from G7 countries. *Advances in Decision Sciences, 27*(1), 1–22. https://doi.org/10.47654/v27y2023i1p1-22

Pesaran, M. H. (2007). A simple panel unit root test in the presence of cross-section dependence. *Journal of Applied Econometrics, 22*(2), 265–312. https://doi.org/10.1002/jae.951

Pesaran, M. H. (2015). Testing weak cross-sectional dependence in large panels. *Econometric Reviews, 34*(6–10), 1089–1117. https://doi.org/10.1080/07474938.2014.956623

Pesaran, M. H., Shin, Y., & Smith, R. J. (2001). Bounds testing approaches to the analysis of level relationships. *Journal of Applied Econometrics*, *16*(3), 289–326. https://doi.org/10.1002/jae.616

Pesaran, M. H., & Yamagata, T. (2008). Testing slope homogeneity in large panels. *Journal of Econometrics, 142*(1), 50–93. https://doi.org/10.1016/j.jeconom.2007.05.010

Ramsey, J. B. (1969). Tests for specification errors in classical linear least-squares regression analysis. *Journal of the Royal Statistical Society: Series B, 31*(2), 350–371.

Salman, A., Chang, B. H., Abdul Razzaq, M. G., Wong, W. K., & Uddin, M. A. (2023). The emerging stock markets and their asymmetric response to infectious disease equity market volatility (ID-EMV) index. *Annals of Financial Economics, 18*(4), Article 2350008.

Syed, A., Lamine, A., & Loukil, S. (2025). Decoding market linkages and risk transmission: A dynamic analysis of G7 stock indices and currency pairs against a changing economic landscape. *Annals of Financial Economics*, *20*(4), Article 2650002.

Teräsvirta, T., Lin, C. F., & Granger, C. W. J. (1993). Power of the neural network linearity test. *Journal of Time Series Analysis*, *14*(2), 209–220. https://doi.org/10.1111/j.1467-9892.1993.tb00139.x

Wen, J., Hai, H., Li, L., Dou, J., & Tao, R. (2025). Digital finance as a catalyst for green innovation: Threshold effect of diminishing marginal returns. *International Review of Financial Analysis, 110*, Article 104840. https://doi.org/10.1016/j.irfa.2025.104840

Westerlund, J. (2007). Testing for error correction in panel data. *Oxford Bulletin of Economics and Statistics*, *69*(6), 709–748. https://doi.org/10.1111/j.1468-0084.2007.00477.x

Wong, W. K., Cheng, Y., & Yue, M. (2024). Could regression of stationary series be spurious? *Asia-Pacific Journal of Operational Research*, Article 2440017.

Wong, W. K., & Pham, M. T. (2022a). Could the test from the standard regression model make significant regression with autoregressive noise become insignificant? *The International Journal of Finance, 34*, 1–18.

Wong, W. K., & Pham, M. T. (2022b). Could the test from the standard regression model make significant regression with autoregressive noise become insignificant — A note. *The International Journal of Finance*, *34*, 19–39.

Wong, W. K., & Pham, M. T. (2023a). Could the test from the standard regression model make significant regression with autoregressive Yt and Xt become insignificant? *The International Journal of Finance, 35*, 1–19.

Wong, W. K., & Pham, M. T. (2023b). Could the test from the standard regression model make significant regression with autoregressive Yt and Xt become insignificant — A note. *The International Journal of Finance*, *35*, 20–41.

Wong, W. K., & Pham, M. T. (2025a). Could the correlation of a stationary series with a non-stationary series obtain meaningful outcomes? *Annals of Financial Economics*. Advance online publication.

Wong, W. K., & Pham, M. T. (2025b). How to model a simple stationary series with a non-stationary series? *The International Journal of Finance, 37*, 1–19.

Wong, W. K., & Pham, M. T. (2026a). Could the panel regression be used to examine the relationship between I(0) and I(1) series? *Advances in Decision Sciences, 30*(2). Advance online publication.

Wong, W. K., & Pham, M. T. (2026b). Could we use correlation to examine panel data with I(0) and I(1) variables? *The International Journal of Finance, 38*. Advance online publication.

Wong, W. K., Pham, M. T., & Yue, M. (2024). Could regressing a stationary series on a non-stationary series obtain meaningful outcomes — A remedy. *The International Journal of Finance*, *36*, 1–20.

Wong, W. K., & Yue, M. (2024). Could regressing a stationary series on a non-stationary series obtain meaningful outcomes? *Annals of Financial Economics, 19*(3), Article 2450011.

Woode, J. K., Ayamba, S. B., Adjei, A. F., Musah, F., & Bambir, J. (2025). Conditional dependence between non-ferrous metals and global uncertainties: Insight from the nickel crash. *Annals of Financial Economics*, *20*(3), Article 2550019.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data* (2nd ed.). MIT Press.

Xiao, Z. (2009). Quantile cointegrating regression. *Journal of Econometrics*, *150*(2), 248–260. https://doi.org/10.1016/j.jeconom.2008.12.005

Xing, L., Chang, B. H., & Aldawsari, S. H. (2024). Green finance mechanisms for sustainable development: Evidence from panel data. *Sustainability, 16*(22), Article 9762. https://doi.org/10.3390/su16229762

Yang, J., & Hui, N. (2024). How digital finance affects the sustainability of corporate green innovation. *Finance Research Letters, 63*, Article 105314. https://doi.org/10.1016/j.frl.2024.105314

Zada, H., Mansoor, A., Khan, N., Wong, W. K., & Jibir, A. (2026). Monetary policy uncertainty and stock market returns in developed and emerging countries: Evidence from a quantile-on-quantile approach. *Advances in Decision Sciences, 30*(3), 1–24.