UNIVERSITY OF TARTU
Faculty of Social Sciences
Johan Skytte Institute of Political Studies

Volodymyr Ksienich

PREDICTING THE SUCCESS OF UKRAINE'S RESTORATION PROJECTS: A MACHINE LEARNING ANALYSIS USING DREAM ECOSYSTEM DATA

MA thesis

Supervisor: Dmytro Khutkyy, PhD

Tartu 2025

Authorship Declaration

I have prepared this thesis independently. All the views of other authors, as well as data from literary sources and elsewhere, have been cited.

Word count of the thesis: 21,519

Name, date: Volodymyr Ksienich, 18.05.2025

Table of contents

List of abbreviations

Abbreviation		Full term
AI	–	Artificial Intelligence
ASC	–	Administrative Services Centre / Центр надання адміністративних послуг (ЦНАП)
AUC	–	Area Under the (ROC) Curve
CMU	–	Cabinet of Ministers of Ukraine / Кабінет Міністрів України
CPV	–	Common Procurement Vocabulary
CSF	–	Critical Success Factor
DREAM	–	Digital Restoration Ecosystem for Accountability & Management
E-services	–	Electronic public services (sub-index of Digital Index)
GBM	–	Gradient Boosting Machine
GDP	–	Gross Domestic Product
IFI	–	International Financial Institution
IRI	–	International Republican Institute
KPI	–	Key Performance Indicator
ML	–	Machine Learning
ROC	–	Receiver Operating Characteristic
SD	–	Standard Deviation
SHAP	–	SHapley Additive exPlanations
SMOTE	–	Synthetic Minority Over-sampling Technique
UT	–	University of Tartu

ABSTRACT

Ukraine's full-scale war has generated 157 billion USD in infrastructure losses and an urgent 524 billion USD reconstruction bill. While prior scholarship isolates single drivers of project delivery, it rarely combines finance, governance, digital capacity and societal sentiment in one empirical frame. This study merges five open datasets (DREAM project register, DREAM-Completeness audit, Transparent Cities scores, Digital-Index metrics and IRI opinion surveys), yielding a moderate-N panel of 190 fully documented wartime restoration projects. A six-pillar theoretical model is operationalised through logistic regression and three ensemble learners. Results show that fragmenting procurements into additional contract lots multiplies completion odds by ≈9.8; publishing real-time finance schedules raises success probability by 12 percentage points independent of budget size; and a one-SD increase in regional digital maturity adds five points, but only where e-services accompany raw openness. Findings nuance 'transparency backlash' theory and suggest an integrated policy bundle: mandatory micro-lotting, conditional disbursement upon schedule disclosure, and targeted e-government investment.

1. INTRODUCTION

Ukraine's wartime reconstruction constitutes a major policy challenge. The ongoing full-scale armed conflict has inflicted unprecedented damage on infrastructure and public services, leading to a GDP decline of approximately 35 per cent in 2022 (Irtyshcheva, Kramarenko, & Sirenko, 2022, p. 80). Direct economic losses from the destruction of residential and non-residential property, critical infrastructure, and essential services have already surpassed USD 157 billion (Kyiv School of Economics, 2024). Looking forward, the Fourth Rapid Damage and Needs Assessment projects that Ukraine will require about USD 524 billion for reconstruction over 2025–2035, with housing, transport, and energy accounting for the largest shares of those needs (World Bank, Government of Ukraine, European Union, & United Nations, 2025, pp. 42-45).

Prior scholarship has rarely tested funding, contractual design, transparency and public sentiment simultaneously because existing datasets either lacked project-level finance (World Bank IEG, 2019) or real-time governance metadata (Prozorro covers procurement but not outcomes). The DREAM platform unites both dimensions, pitting competing explanations against each other in a single model. Focusing on the 190 projects with complete multi-pillar data, I offer a novel comparative test of these theories under documented wartime conditions, extending earlier single-pillar studies such as World Bank IEG (2019).

The thesis explains why some wartime restoration projects in Ukraine reach completion while others stall, despite uniform access to the national treasury and an identical legal framework.

Effective restoration requires a comprehensive strategy integrating technical solutions with socio-economic and governance considerations in this challenging context.

Digital transformation has emerged as a vital lever, enabling the adoption of innovative tools, real-time data analytics, and transparent governance mechanisms essential for navigating reconstruction amid ongoing conflict (OECD, 2024). Central to this innovation is Ukraine's DREAM (Digital Restoration Ecosystem for Accountability and Management) platform, a legally established digital governance initiative operating under Ukraine's Ministry for Communities and Territories Development of Ukraine. The DREAM ecosystem provides an integrated, open-access system for monitoring restoration projects from initiation to completion, ensuring accountability and public oversight through transparent reporting and participatory governance structures

(Transparent Cities Ukraine, 2024; DREAM Ecosystem Official Presentation, 2024).

Despite the availability of detailed project-level data within the DREAM platform, a significant gap persists in systematically identifying the key factors contributing to the success or failure of restoration initiatives. Previous attempts at project evaluation in post-conflict environments have predominantly employed manual evaluations, relying heavily on human judgment and retrospective audits, which often struggle to handle the complexity and uncertainty of wartime reconstruction.

This thesis seeks to answer the following main research question: *What are the key project attributes and contextual factors influencing the success of Ukraine's restoration projects managed through the DREAM ecosystem?* To address this question, I propose an interdisciplinary theoretical framework rooted in project management, governance transparency, big data analytics, and behavioural political economy. My theoretical expectations posit that successful restoration projects will likely correlate positively with diversified contractual arrangements, high financial coverage, robust digital transparency, and favourable societal sentiment indicators. Conversely, I expect lower success rates in projects with insufficient funding, low transparency, limited civic engagement, and inadequate governance mechanisms.

Employing a moderate‑N quantitative research design, the analysis begins with a dataset of 12,781 raw project entries from DREAM. After applying rigorous data-cleaning procedures detailed in Chapter 3, I retain 190 fully documented projects suitable for quantitative modelling. The findings underscore that effective wartime restoration demands packaged funding complemented by enhanced digital transparency and robust governance frameworks, providing critical insights into policy alignment for sustainable recovery.

The remainder of this thesis proceeds as follows: Chapter 2 offers a comprehensive review of relevant literature spanning project management success factors, governance transparency, behavioural political economy, and machine learning applications in public policy. Chapter 3 details the research design, encompassing data collection, integration methods, preprocessing strategies, and model selection rationale. Chapter 4 presents empirical results, including model evaluation, interpretability analyses, and policy simulation outcomes. Chapter 5 discusses these findings critically, examining their implications for theory, practice, and policymaking while addressing limitations and ethical considerations. Finally, Chapter 6 summarises key insights, highlights implications for project management practice and governance policy, and identifies directions for future research.

When such models inform public decisions, their interpretability is crucial. Therefore, SHAP values (Lundberg & Lee, 2017) have become a standard device for explaining complex models, quantifying each feature's contribution to a prediction. Chapter 4 applies SHAP alongside classical coefficients to validate the explanatory narratives that emerge from the Ukrainian reconstruction data.

Researchers and practitioners in project management are increasingly employing machine learning to forecast project outcomes and identify risk factors. The central premise is to learn from historical project data to predict whether ongoing or proposed projects will succeed or fail, and to understand the underlying reasons for these outcomes. Machine learning models can discern patterns that correlate with success metrics by training on features from past projects, such as budget size, project duration, team experience, economic context, and governance indicators. This includes timely completion, adherence to budgets, and the fulfilment of beneficiary needs. In a recent study, Satri et al. (2024) demonstrated that these techniques can effectively predict the outcomes of regional development projects while also identifying key risk factors.

In the context of the study, the logic will also be used that although logistic regression and single CART trees offer transparency, they underperform on my moderately large, noisy sample; ensemble methods (RF, GBM, XGBoost) strike the optimal balance of predictive strength and interpretability, as demonstrated by higher F1‑scores on hold‑out tests (Chen & Guestrin, 2016; Powers, 2011), which will also be confirmed in Chapter 4.

One study applied machine learning (Random Forest algorithms) to a dataset of 160 construction projects (Bang et al., 2022). It showed that it is possible to predict project success with reasonable accuracy and pinpoint which factors most strongly influence success in that dataset (Bang et al., 2022). The model's top predictors aligned with project management best practices – for example, projects with rigorous upfront planning, strong leadership involvement, and manageable complexity were far more likely to succeed. This demonstrates how AI can complement expert judgment, confirming known success factors and potentially revealing subtle predictors that managers might overlook.

2. LITERATURE REVIEW

2.1 Governance, transparency, and accountability in the digital era

Governance, defined as the processes through which decisions are made, is essential for public-sector projects (World Bank, 1994, p. 14). Digital transformation reinforces these frameworks by enabling real-time data exchange and citizen engagement via platforms like DREAM, strengthening accountability. Moreover, international sanctions and fiscal policy shifts during the armed conflict highlight the need for transparent governance models that are complemented by coordinated regional strategies to modernise infrastructure and enhance economic resilience (Irtyshcheva, Kramarenko & Sirenko, 2022, pp. 79–80; Erfan & Koltsov, 2023, pp. 123–124).

Digital governance is the strategic use of digital technologies – platforms, open‑data portals, e‑participation tools – to transform public administration processes, structures, and norms, enabling more responsive, inclusive, and accountable decision‑making (Dunleavy et al., 2006). Behavioural political economy integrates cognitive psychology and sociology insights into economic and political analysis, emphasising how bounded rationality, social norms, and trust dynamics shape policy uptake and institutional performance (Simon, 1957; Grossman & Baldassarri, 2012).

In post-conflict reconstruction, the quality of governance – encompassing the rule of law, anti-corruption measures, efficient administration, and inclusive decision-making – is critical to project success. Although transparency involves the disclosure of information, genuine accountability is achieved only when citizens can access and act upon that data. A robust chain of disclosure, awareness, public scrutiny, and enforcement is necessary; weaknesses in any link may undermine accountability, particularly in fragile democracies (Fox, 2007; Miller, 2005; Dyak, 2023, pp. 5–6).

Research shows that the effects of transparency differ by context; in highly corrupt settings, greater disclosure may lead to public resignation rather than activism (Bauhr & Grimes, 2014). By contrast, Björkman & Svensson (2009) and Malesky et al. (2014) find that where civil-society networks are dense, transparency does trigger citizen mobilisation even under high corruption, suggesting that social capital moderates the "transparency backfire" mechanism.

However, given Ukraine's active digital governance trajectory exemplified by adopting open-data platforms such as Prozorro and DREAM, this thesis anticipates that higher governance quality and transparency levels will positively correlate with improved project outcomes. Specifically, transparency and robust institutional governance are expected to significantly enhance the likelihood of successful restoration projects, as they provide mechanisms to deter corruption and foster greater public trust and civic engagement.

Ukraine's adoption of open-data platforms has set benchmarks in digital governance. The DREAM ecosystem exemplifies digital innovation by uniting various governmental layers into a central platform that enables participatory restoration planning and enhances transparency (World Bank, 2018, p. 23; Szpak et al., 2024, p. 203; Transparent Cities Ukraine, 2024, p. 14).

Open government data is intended to enable external stakeholders to scrutinise government information. However, its effectiveness hinges on citizens having the means and the motivation to act on that information. Without sufficient political agency, enhanced transparency alone cannot guarantee accountability (Robinson & Yu, 2012; Peixoto, 2013).

Digital technology has spurred the development of civic-tech innovations, such as online dashboards and mobile apps that facilitate real-time public oversight. The DREAM ecosystem exemplifies this by providing a unified platform that monitors project progress and serves as a donor-coordination tool, enhancing efficiency and accountability. Evidence from Ukraine indicates that, despite transparency measures such as the Prozorro platform, procurement practices have shifted towards direct methods to expedite processes, highlighting the need for platforms such as DREAM to maintain public trust through robust oversight (Klymak & Vlandas, 2024, pp. 8–9; Heyets, Blyzniuk & Nykyforuk, 2022, p. 5; Transparent Cities Ukraine, 2024, p. 14).

Finally, the Ministry of Digital Transformation's Digital Index (2024) enriches empirical assessments of regional governance. This index provides a systematic benchmark for digital readiness and is incorporated in the modelling framework advanced in Chapters 3 and 4.

2.2 Post-conflict reconstruction and project governance: international perspectives

Wartime restoration in Ukraine is unfolding under active conflict, combining technical project management with simultaneous nation-building and economic recovery efforts. Scholars of development and international relations have studied cases from post–World War II Europe to recent Middle East and African conflicts to derive frameworks for successful reconstruction (Collier et al., 2008; UNDP, 2008). However, mid-20th-century paradigms such as the Marshall Plan operated in a pre-digital context, relying on centralised aid flows and ad hoc coordination. By contrast, 21st-century digital platforms like DREAM embed real-time monitoring, multi-actor transparency, and algorithmic resource allocation – features that demand a reappraisal of classical reconstruction theories to account for instantaneous data feedback loops and citizen engagement at scale (DREAM Ecosystem Official Presentation, 2024).

Two overarching challenges face any post-conflict country: restoring a functional, peaceful state (reducing the risk of renewed conflict) and reviving the economy and society after the devastation. Emerging from the current armed conflict, Ukraine faces these dual challenges of peacebuilding and development simultaneously. Moreover, viewing multiple postwar periods as a heritage enables the incorporation of diverse historical narratives, informing a more resilient, context-sensitive framework for project governance (Dyak, 2023, pp. 2–4).

Further in the study, I will give preference to regional indicators of trust in local authorities, perceptions of the quality of municipal services and economic outlook prepared by Sociological group “Rating” (2023), choosing them to operationalise the public sentiment component of my system, linking public opinion directly to restoration outcomes.

A recent study emphasises that Ukrainian cities such as Kyiv, Kharkiv, and Irpin are emerging as active agents in their reconstruction. Through bilateral twin-city partnerships and robust city networks, these urban centres are receiving humanitarian aid and providing critical expertise in urban planning and public service delivery. This participation is instrumental in fostering urban autonomy and long-term resilience in the post-conflict period (Szpak et al., 2024, pp. 200–205).

This international perspective directly informs the six-pillar theoretical framework – governance quality, financial sufficiency, project-management practice, digital capacity, societal sentiment and post-conflict constraints – used to structure the empirical analysis in Chapters 3 and 4.

A United Nations report observes that countries transitioning from war typically inherit severely weakened state capacity alongside widespread destruction of physical infrastructure, human capital, and social cohesion (UNDP, 2008). They often confront distorted economic incentives (e.g., war economies, black markets), extreme poverty and unemployment, and sometimes active resistance from spoilers such as warlords or criminal networks that profited during armed conflict. These conditions mean post-conflict societies are extremely fragile – the risk of armed conflict relapse is much higher than in other developing contexts, especially in the first few years after a peace agreement or ceasefire. Paul Collier et al. (2008) highlight that reducing the risk of renewed violence becomes the priority of the immediate post-conflict period. This might involve demobilisation programs, security sector reform, and inclusive political arrangements to address grievances. At the same time, urgent economic needs must be tackled: rebuilding infrastructure, creating jobs for ex-combatants and civilians, and jump-starting basic services such as water, electricity, healthcare and education. A post-conflict economy's unique constraints and opportunities often call for distinctive strategies not found in standard development toolkits (Collier et al., 2008). For example, large inflows of foreign aid are common and can help recovery, but they require careful management to avoid dependency or reigniting armed conflict over resource distribution. The trade-offs observed in procurement during the war, marked by lower competition and higher costs, illustrate broader challenges in re-establishing efficient, transparent procurement systems, emphasising that post-conflict recovery efforts must address the restoration of cost-effectiveness and competitive processes to rebuild state capacity sustainably (Klymak & Vlandas, 2024, pp. 9–10)

Comparative analyses suggest several key priorities for successful post-conflict economic recovery and reconstruction. These include:

1. Security and rule of law – establishing a safe environment is fundamental, as people and investors will not engage if violence persists. This extends to re-establishing the rule of law and judicial processes.
2. Rehabilitation of infrastructure and essential services – repairing roads, bridges, and power grids and ensuring schools and hospitals function to provide tangible peace dividends to the population.
3. Employment generation and livelihoods – focusing on job creation (e.g., labour-intensive public works) to absorb unemployed youth and former fighters, boosting household incomes and domestic demand (UNDP, 2008).

4. Macroeconomic stabilisation and incentives for investment – stabilising inflation, currency, and banking while encouraging private sector revival through business-friendly reforms.
5. Reducing group inequalities and regional disparities – post-conflict recovery must be inclusive; social tensions could revive if certain regions or groups feel left out of reconstruction benefits (UNDP, 2008). Policies to ensure equitable allocation of aid and projects (for instance, rebuilding both majority and minority communities) help rebuild trust among formerly warring groups. Underlying all this is
6. rebuilding government capacity – without capable institutions to plan and implement projects, even well-funded reconstruction efforts falter. This involves training civil servants, improving public economic management, and establishing anti-corruption measures to use aid and resources effectively.

The UNDP's (2008) analysis of post-conflict recovery argues that strategies should be "anchored in the local engines of economic recovery". Many informal institutions and community coping mechanisms develop during armed conflict; leveraging these indigenous systems (rather than bypassing them) can enhance the legitimacy and sustainability of reconstruction (UNDP, 2008). Additionally, the post-conflict period can be seen as an opportunity to implement reforms that may have been impossible. There is often a window of international goodwill and internal momentum for change. However, Collier et al. (2008) warn that this window is sometimes wasted: too much emphasis by international actors goes into short-term peacekeeping and political fixes, and "too little attention is paid to economic reform", which is essential for long-term peace. Therefore, effective project governance in reconstruction entails balancing short-term needs (quick wins to show progress and maintain peace) with long-term institution-building (so the country can eventually manage independently).

Post-conflict reconstruction typically involves many actors – domestic government agencies, donor governments, international organisations (UN, World Bank, etc.), NGOs, and private contractors. Coordination among these is notoriously challenging but crucial. Fragmented efforts can lead to duplication in some areas, while other needs are neglected. To address this, countries often set up special project governance structures: for example, centralised reconstruction funds or ministries to oversee and coordinate projects, joint donor coordination boards, and transparency portals to track all projects (such as the DREAM platform might do for Ukraine) (Collier et al., 2008; UNDP, 2008). Strong project governance also means establishing clear criteria for project selection (to align with national priorities and avoid patronage), robust monitoring and evaluation systems (to track progress and outcomes), and involving local governments and communities in decision-making (to ensure projects meet real needs on the ground). The case of Afghanistan's National Solidarity Programme (a community-driven development project after 2001) is often cited as a good practice, where village councils received block grants for local projects, giving communities control, which helped spread reconstruction benefits and build trust in the government. In contrast, examples of poorly governed reconstruction include instances where funds were mismanaged or only benefited elites, undermining public confidence (UNDP, 2008).

This contrast reveals a methodological gap: traditional reconstruction frameworks emphasise large-scale funding mechanisms and diplomatic coordination, whereas digital-era tools foreground distributed accountability, participatory governance, and data-driven decision-making. Recognising these differences sharpens my theoretical lens for analysing DREAM's innovations in Ukraine.

From an international perspective, successful post-conflict recoveries (such as the Marshall Plan in post-war Europe (Judt, T., 2005) or, more recently, Rwanda's reconstruction after the 1994 genocide (Larose-Edwards, 2002)) combined effective governance, external support, and grassroots participation. They also implemented economic policies tailored to post-conflict realities rather than applying one-size-fits-all models (Collier et al., 2008).

2.3 Behavioural political economy: public opinion, trust, and policy implementation

Even with strong formal institutions and advanced analytics, the success of public projects in a democracy hinges on people – the behaviour of citizens, public officials, and stakeholders. The subfield of behavioural political economy integrates insights from psychology and sociology into the analysis of political and economic processes. It reminds us that public trust, perception of fairness, social norms, and collective action capacity can influence policy implementation (Grossman & Baldassarri, 2012).

Digital transformation also plays a crucial role in shaping public opinion by enhancing transparency and enabling more direct citizen engagement, thereby contributing to increased legitimacy and trust in post-conflict governance (OECD, 2024, p. 92). For post-conflict reconstruction in particular, citizens' willingness to support and participate in rebuilding efforts (or, conversely, their potential to resist or disengage) is crucial.

One key aspect is how public opinion and legitimacy affect policy outcomes. Policies and projects are more likely to be implemented effectively if they enjoy broad support or acquiescence from the public. A government that is seen as legitimate and accountable can mobilise citizens to cooperate with reconstruction initiatives – for example, by paying taxes, contributing labour, or simply not obstructing projects due to distrust. Empirical evidence supports the idea that procedural fairness in governance breeds cooperation (Grossman & Baldassarri, 2012). In a field experiment in rural Uganda, communities were more willing to contribute to a public good (investing their effort in a project) when the local leader who could punish "free riders" was elected by the community rather than appointed (Grossman & Paler, 2015). Essentially, the actual "rules of the game" and incentives were held constant, and having a say in choosing the leader increased people's contributions. This finding aligns with broader theories that democratic participation and local ownership enhance legitimacy, motivating citizens to cooperate with and invest in public projects. For Ukraine, ensuring that local communities have a voice and representation in deciding reconstruction priorities could similarly increase grassroots buy-in and the success of those projects.

Post-conflict societies often experience shifts in social trust and collective action capacity. Paradoxically, research (Bellows & Miguel, 2009) has found that experiencing conflict can sometimes strengthen local social capital in the aftermath. In Sierra Leone, individuals from villages that suffered more intense violence during the civil war showed higher levels of post-war community engagement – they were more likely to attend community meetings, join local groups, and vote in elections. Bellows and Miguel (2009) interpret this as war hardships galvanising citizens to become more involved in their communities, potentially to rebuild and prevent a return to violence. Similar patterns of heightened local cooperation have been observed in other post-conflict contexts (e.g., Uganda and Liberia), suggesting that shared adversity can increase solidarity and collective action at the community level. This is an unseen benefit for reconstruction: projects can potentially tap into a reservoir of public goodwill and willingness to contribute, provided the initiatives are inclusive and seen as addressing community needs.

On the other hand, armed conflict can also erode trust in central authorities, especially if the state or elites are viewed as responsible for or complicit in past violence or corruption. Thus, there is often a trust deficit between citizens and the government after an armed conflict (Collier et al., 2008). Overcoming this requires demonstrable improvements in governance and careful attention to fairness in how reconstruction resources are distributed to avoid perceptions of favouritism or exclusion of certain groups.

Behavioural political economy also examines how citizens hold leaders accountable through elections, protests, or other feedback. In stable settings, public opinion can discipline policymakers – a government that fails to deliver reconstruction benefits may face backlash at the ballot box or in public demonstrations. However, this mechanism only works if citizens have access to information (transparency) and believe their engagement can be effective (political efficacy) (Grossman & Baldassarri, 2012). This ties back to transparency initiatives: when citizens are provided with clear information on what to expect from projects and how funds are used, they are better positioned to form opinions and demand results. Suppose a road is supposed to be rebuilt, and a transparency portal shows it is far behind schedule or over budget. In that case, citizens can pressure local officials or use the media to draw attention to the issue, potentially accelerating corrective action.

Behavioural insights also apply to those implementing policies – public officials and contractors – who respond to incentives and norms. For instance, anti-corruption efforts often use behavioural principles like increasing the perceived likelihood of being caught (through audits or citizen monitoring) to deter bribery. Subtle "nudges" can influence bureaucratic performance, such as performance feedback or public recognition for good service (Bauhr & Grimes, 2014). All these human factors determine whether reconstruction projects are completed successfully.

2.4 Project management success factors

Traditional project management literature defined success narrowly in terms of the "iron triangle" – on time, on budget, and within scope (Pinto, 2010, p. 35). In Ukraine's wartime reconstruction, while time, cost, and scope remain baseline criteria, success depends on navigating security constraints, acute material shortages, and rapid, decentralised decision-making under martial law.

Moreover, contemporary research has broadened the "iron triangle" perspective. Success is now seen as multidimensional and context-dependent. For example, Jugdev and Müller (2005) note that whereas earlier definitions focused on project implementation efficiency, modern definitions evaluate the entire project life cycle and its strategic impact. This expanded focus includes delivery (meeting schedule, cost, quality targets) and effectiveness – i.e., achieving the project's purpose, stakeholder satisfaction, and alignment with organisational strategy (Jugdev and Müller, 2005). Thus, a project is deemed successful if it delivers value and benefits that justify its investment beyond mere adherence to initial plans.

Recent OECD analysis highlights that integrating digital transformation initiatives into project management boosts operational efficiency and enhances project resilience by leveraging digital tools for real‑time monitoring and collaboration (OECD, 2024, p. 85). In addition, the construction sector contracted by approximately 65% in the building market in 2022 due to halted production, import disruptions, and damaged infrastructure. These highlight the need for innovative rebuilding strategies incorporating enhanced safety standards and modern technological solutions (Skrynkovskyy et al., 2022, p. 65).

Researchers have long sought to identify factors that consistently contribute to project success. Belassi and Tukel (1996) provided an early framework categorising CSFs into four groups: factors related to the project (size, complexity, urgency), the project manager and team (leadership, skills, commitment), the organisation (support, structure, resources), and the external environment (political, economic, social context). In their work, they point out that the success of a project is a systemic outcome influenced by internal and external variables.

Modern studies have refined project management success factors by integrating agile methodologies and complexity science. Serrador and Pinto (2015, pp. 1040–1041) demonstrate that practices emphasising adaptability, continuous learning, and active stakeholder engagement can significantly enhance project outcomes in turbulent environments. Moreover, their findings suggest that the success of large-scale development and restoration projects is also contingent on external factors such as political support, regulatory stability, and community buy-in.

In Ukraine's wartime period, adaptability in restoration means revising workstreams in response to damage reports; continuous learning is built into rapid after‑action reviews of air‑raid and missile strike interruptions; and stakeholder engagement prioritises coordination with regional military administrations, civil authorities, and international donors to align security and reconstruction objectives.

Contemporary research has attempted to validate CSFs empirically and adapt them to various domains. Ika et al. (2012), examining World Bank-funded development projects, identified five critical factors correlated with project outcome success: robust monitoring and evaluation, effective coordination among stakeholders, sound project design and planning, adequate training and capacity building, and a supportive institutional environment. This is an example that, beyond technical execution, projects benefit from strong oversight mechanisms (monitoring), collaboration across agencies or donors, upfront investment in planning and human resources, and an enabling governance context.

In the Ukrainian context, reliable monitoring using digital dashboards (for example, DREAM) allows donors to receive information about changing risks and progress in project implementation; coordination between the central government, local councils, NGOs and donors prevents duplication and allows effective application of efforts; design and planning include a clear understanding of the needs for such efforts, the required volume and the consequences of such actions; training prepares local authorities for the correct use of such technological tools to achieve the best result with the least number of issues; and a supportive institutional environment ensures proper legislative support, tool optimisation, and barrier-free compliance with procedures.

Other recent studies similarly emphasise early planning, risk management, top management support, and leadership involvement as persistent drivers of success. For example, machine-learning analysis of construction projects by Bang et al. (2022) found that thorough early-phase planning and executive support were among the most crucial factors distinguishing successful projects, confirming what project management theory has long asserted.

Research also shows that traditional success factors may need adaptation in complex settings like postwar reconstruction. The drastic reduction in industrial output and the disruption of production chains reported during the armed conflict emphasise the necessity for adaptive project management practices capable of navigating extreme uncertainty and rapidly changing fiscal conditions (Irtyshcheva, Kramarenko, & Sirenko, 2022, p. 80).

Flexibility and stakeholder participation become even more crucial when working in unstable environments. Modern perspectives argue for integrating local context knowledge, armed conflict sensitivity, and sustainable impact into project management practices for reconstruction projects. Thus, modern project management theory lays the foundation for determining success factors – from leadership and team competence to coordination and planning of stakeholders' actions, and in the context of this study, specific CSFs – security-aware scope definition, modular delivery, adaptive leadership, resilient institutional support, digital real-time oversight, and integrated civil‑military coordination – map directly onto my six‑pillar theoretical framework and inform the choice of variables in Chapters 3 and 4.

2.5 Big data and machine learning in public policy and project analysis

The explosion of big data and machine learning (ML) advances has significantly impacted public policy analysis and project management in recent years. Digitalisation complements these advances by providing the necessary infrastructure for robust data collection and processing, improving the predictive accuracy of ML models and facilitating evidence‑based decision‑making in reconstruction projects (OECD, 2024, p. 78). Unlike traditional analytical methods that rely on small, infrequent data samples and simple models, ML techniques can ingest vast datasets (including real-time streams) and detect complex patterns beyond human capacity. This is especially useful in policy domains where outcomes are influenced by many interacting factors – for example, economic development, public health, or infrastructure project performance (Peet et al., 2022). It was also stated that by uncovering non-linear relationships and high-dimensional interactions, ML often yields more accurate predictions of policy outcomes than linear models, thus helping to reduce the risk of unintended consequences (Peet et al., 2022). In short, the data revolution provides policymakers with new tools to anticipate and address problems before they fully manifest.

Numerous state-of-the-art AI and predictive analytics examples are applied in the public sector. At the local level, some city governments mine data from sources like social media or 911 service requests to identify issues and allocate resources more efficiently (Peet et al., 2022). For instance, one city used ML to analyse tweets about restaurant cleanliness to target health inspections of the riskiest establishments (Peet et al., 2022, p. 2). Governments have used algorithms to optimise refugee resettlement placements at the national level. Bansak et al. (2018) demonstrated that an ML model assigning refugees to communities (based on past integration success patterns) improved employment outcomes for refugees compared to human judgment. ML models are also employed to forecast and manage a variety of public sector challenges: disaster response agencies aggregate satellite and sensor data to predict where relief is needed most, criminal justice systems have piloted ML-based risk assessments to inform bail or probation decisions (though not without controversy), public health officials use ML to predict disease outbreaks or patient readmission risks, and school systems have used algorithms to identify students at risk of dropping out for early intervention (Peet et al., 2022).

As a result, the machine learning (ML) approach provides a powerful methodology to uncover intricate, non-linear relationships among project attributes, contextual factors, and outcomes. ML methods enable analysts to systematically manage complexity, process large datasets effectively, and produce robust predictive insights that manual methods cannot consistently achieve (Peet et al., 2022).

2.6 The DREAM ecosystem: a digital restoration platform for post-conflict reconstruction in Ukraine

The rapidly evolving challenges of post-war reconstruction in Ukraine necessitate innovative digital solutions that streamline project management and reinforce transparency, accountability, and participatory governance. The DREAM (Digital Restoration Ecosystem for Accountability and Management) system exemplifies such innovation. It is designed as a state digital ecosystem that unites local communities, regional administrations, and national agencies under a single platform for managing restoration projects from initiation to completion (DREAM Ecosystem Official Presentation, 2024).

The DREAM system is established and operates under the Cabinet of Ministers of Ukraine's Resolution № 1286 of 15 November 2022 (as amended by Resolutions № 434/02.05.2023; № 903/09.08.2024; № 272/07.03.2025) and the Ministry for Communities, Territories and Infrastructure Development's Decree № 65 of 23 January 2024 (Cabinet of Ministers of Ukraine, 2022/2023/2024/2025; Ministry for Communities, 2024), which together define its mandate, structure and pilot-implementation procedures. As formalised in the 2024 Order of the Ministry for Communities, Territories and Infrastructure (No 65/23 Jan 2024), DREAM's Regulation specifies its three components – Central Database, Electronic User Cabinet and Public Portal – with all implements bound by the Law 'On Protection of Information in ICT Systems' for data security.

Ownership of the Unified Digital Integrated Information and Analytical System (DREAM) resides with the State, represented by the Ministry for Communities, Territories and Infrastructure Development. At the same time, technical administration is entrusted to the State Agency for Restoration (Cabinet of Ministers of Ukraine, 2022/2023/2024/2025). Per CMU № 1286/2022 (art. 4) and its subsequent amendments (CMU № 434/2023; CMU № 903/2024; CMU № 272/2025), the Ministry for Communities, Territories and Infrastructure Development acts as DREAM's owner ('holder of records'), while the State Agency for Restoration is the official technical administrator, responsible for uptime, backups, user-role management and cybersecurity measures.

Strategic oversight is provided by the Strategic Investment Council (established by CMU № 549/2024), and all users – from city councils to regional military administrations – are bound by the Procedures approved under CMU № 382/2023. CMU № 903/2024 introduces a dedicated 'Concepts' block within DREAM, enabling budget authorities to draft, e-sign (QES) and submit investment concept notes for automated scoring and prioritisation (Appendix 2 criteria). CMU № 221-p/2023 further anchors these workflows in the 2023 Priority Action Plan (Cabinet of Ministers of Ukraine, 2024; Cabinet of Ministers of Ukraine, 2023).

DREAM's regulatory framework is further anchored in amendments to the Budget Code (Law of Ukraine), the Government's Priority Action Plan 2023 (CMU № 221-p/2023), and the pilot-project rules in CMU № 731/2023 (Law of Ukraine on Budget Code, 2020; Cabinet of Ministers of Ukraine, 2023), ensuring compliance with public-investment planning, procurement and transparency requirements.

One of the core strengths of the DREAM ecosystem is its capacity to integrate multiple state IT systems into a unified framework. Doing so provides an end-to-end digital management process that enables real-time monitoring of restoration projects and supports efficient resource allocation. This integration delivers measurable benefits by tracking each project systematically, thus bolstering technical efficiency and financial accountability (Transparent Cities Ukraine, 2024, p. 12; DREAM Ecosystem Official Presentation, 2024).

Regarding governance, DREAM exemplifies how digital business transformation can enhance public-sector accountability. By functioning as an open-access platform for public monitoring, DREAM empowers citizens to observe every stage of the restoration process, thereby creating direct pressure on officials to maintain high standards of transparency. This public-oversight mechanism aligns with the broader theoretical notion of "open data-driven public accountability," where transparency is used not merely as a disclosure tool but as a means to trigger corrective action through active citizen engagement (Robinson & Yu, 2012; Transparent Cities Ukraine, 2024, p. 14).

Furthermore, the DREAM system serves as a multi‑agency donor‑coordination platform. It streamlines the attraction, allocation and monitoring of funds from national and international sources. Such a coordinated digital approach is critical in Ukraine's fragile post‑conflict economy, where the efficient management of public investment and donor funding is paramount. The DREAM ecosystem bridges the gap between public investment and project execution, ensuring that restoration initiatives are timely, sustainable and aligned with public priorities (Transparent Cities Ukraine, 2024, p. 14; DREAM Ecosystem Official Presentation, 2024).

Another significant feature of DREAM is its business‑intelligence module. Equipped with an open API and integrated with platforms such as Prozorro, the system provides a rich, real‑time data repository. Under the Ministry's Regulation (Arts I–III), DREAM comprises a Central Database for secure archival and data lineage; an Electronic User Cabinet with role‑based QES authentication; and a Public Portal compliant with DSTU EN 301549:2022 for accessibility.

CMU № 380/2022 (as amended) mandates DREAM's real‑time data exchange with the Diia portal for conflict‑damage reports, with notary/CPAS feed‑ins and automatic reconciliation against the State Register of Damaged Property. Information security conforms to the Law 'On Protection of Information in ICT Systems' and CMU № 1160/2022 on technical‑cryptographic protection (Cabinet of Ministers of Ukraine, 2022; Law of Ukraine on Protection of Information in ICT Systems, 2003; Cabinet of Ministers of Ukraine, 2022).

Machine‑learning models can leverage this comprehensive data stream from the Public Portal to predict success, identify high‑risk projects and guide resource allocation effectively. Thus, DREAM supports traditional project‑management practices and paves the way for advanced data‑driven decision‑making in reconstruction (Transparent Cities Ukraine, 2024, p. 18).

The interdisciplinary nature of DREAM is further reflected in its role as an enabler of participatory restoration planning. By consolidating project data and ensuring continuous public oversight, DREAM fosters a participatory model in which local communities have a direct voice in determining restoration priorities. This participatory governance is essential in post‑conflict settings, where rebuilding trust and social cohesion is as important as physical reconstruction (DREAM Ecosystem Official Presentation, 2024).

The DREAM ecosystem integrates digital governance, project management and data analytics into a unified restoration platform. Its design enhances the efficiency of project execution and ensures that restoration projects are conducted transparently and inclusively. As such, DREAM represents a pioneering tool in post-conflict reconstruction – one that has already seen rapid adoption among Ukrainian municipalities, with usage among city councils increasing from 56 % to 84 % within just a few months (Transparent Cities Ukraine, 2024, p. 15). Therefore, this robust digital platform is critical for the technical management of reconstruction projects and serves as a cornerstone for accountable and participatory governance in Ukraine's long-term recovery process.

2.7. Publicly available datasets for reconstruction analysis

This study leverages several publicly accessible datasets to build a comprehensive analytical framework for evaluating Ukraine's restoration projects. The principal source is the DREAM (Digital Restoration Ecosystem for Accountable Management) platform, which now comprises two complementary datasets. The first is the core project register that captures detailed, real-time information on each reconstruction initiative, including sector and regional identifiers, headline financial figures, and binary outcomes that distinguish completed from unsuccessful projects (DREAM, n.d.). The second, recently released, is the DREAM Project Completeness dataset. While the core register describes the project, the completeness file records the available information. The file reports, whether key documentation, metrics, targets, budgeting schedules, visual evidence, or contracting processes, have been uploaded to the portal for every project code.

Alongside DREAM, the Transparent Cities initiative of Transparency International Ukraine supplies systematic assessments of municipal openness and accountability. Having started with the hundred largest cities between 2017 and 2021, the programme has adapted its methodology to wartime conditions: seventy rear-line cities were assessed in 2022 and eighty in 2023, with several municipalities improving their scores by up to thirty-five percentage points (Transparent Cities Ukraine, 2024). These ratings allow us to control the wider governance environment in which a project is delivered.

Public opinion data come from the 2023 wave of the International Republican Institute survey series. The IRI data yield region-level measures of trust in local authorities, perceptions of service quality, and citizens' economic outlooks; sampling weights ensure representativeness for each oblast's urban population (Sociological group "Rating", 2023).

Guided by my interdisciplinary framework that will be presented in Section 3.1, I deliberately selected IRI indicators: municipal services quality, city-council and mayoral approval, and local government information awareness to operationalise the pillars of public-sector accountability, societal sentiment, and participatory governance as per literature review findings.

Finally, the Digital Index compiled by the Ministry of Digital Transformation quantifies each region's digital readiness by scoring institutional capacity, internet penetration, e-service maturity, and related dimensions. From 2024, the index also incorporates a CDTO sub-index, which accounts for 60% of the overall score, reflecting the growing importance of individual digital projects in measuring regional performance (Ministry of Digital Transformation of Ukraine, 2024).

Taken together, these four sources – two from DREAM, one from Transparency International Ukraine, one from IRI, and one from the Ministry of Digital Transformation – offer a multi-layered view of Ukraine's reconstruction landscape, combining granular project attributes with broader indicators of municipal transparency, public sentiment, and digital maturity, they provide multi-layered empirical coverage but also dovetail with my Chapter 3 research design, where each theoretical pillar directly guides variable selection and operationalisation.

Table 2.1 summarises how each theoretical pillar is operationalised and the directional expectation tested in Chapter 3.

Table 2.1. Concepts, empirical indicators and hypothesised signs.

Theoretical pillar	Operational indicator(s)	Expected sign*	Primary source
Governance quality	• Transparency Index (Transparent / Partially / Non-transparent) • Digital-Index Openness sub-score (0–1)	+	Transparency Int. (2023); MinDigTransf (2024)
Financial sufficiency	• Financial_Efficiency (Committed / Estimated budget)	+	DREAM core

Project-management practice	• Number of Contracts (count) • Project Complexity = Number of Objects	+ –	DREAM core
Digital capacity	• Digital-Index Composite score (0–1) • Digital Index E-Services sub-score (0–1)	+	MinDigTransf (2024)
Societal sentiment	• Municipal-Services Quality Index (1–5) • City-Council Approval (1–5)	+	IRI 2023
Post-conflict constraints	• Multiple-Regions flag (0/1) • Security-Proximity dummy (front-line oblasts)	–	DREAM core; author coding

** "+" = predicted increase in completion probability; "–" = predicted decrease.*

2.8. Synthesis of literature and hypotheses

The six formal hypotheses emerge directly from the literature reviewed in Chapter 2. For governance, Bauhr and Grimes (2014) argue that transparency improves service delivery only when local civic capacity is high; translating that conditional claim into the DREAM context yields H1, which anticipates higher completion odds where regional openness and digital transparency scores are strong. Collier et al. (2008) and UNDP (2008) stress liquid financing as a buffer against post-conflict shocks, motivating H2. Classic and machine-learning studies of critical success factors (Belassi & Tukel, 1996; Bang et al., 2022) show that granular contract architecture and manageable scope drive performance, hence H3a/H3b. Digital-governance research (Dunleavy et al. 2006; OECD 2024) links e-services capacity to bureaucratic throughput, underpinning H4. Behavioural-political-economy scholarship (Grossman & Baldassarri, 2012) connects citizen trust to collective-action outcomes, giving H5. Finally, the post-conflict economics canon (Collier et al. 2008) predicts logistical drag in front-line areas, captured by H6. Table 2.2 formalises these links, and the empirical tests in Chapter 4 trace them back to their theoretical origins.

Table 2.2. Hypothesis map.

H1 (Governance quality).	Projects in more transparent and digitally open regions – captured by the Transparency Index and the Digital-Openness sub-score – should exhibit higher completion odds.
H2 (Financial sufficiency).	Schemes whose committed funding covers at least 80 % of the estimated budget are more likely to finish successfully.
H3 (Project management practice).	a) A greater number of contracts distributes risk and raises success probabilities, whereas (b) higher project complexity (more physical objects) lowers them.
H4 (Digital capacity).	Higher regional scores on the Digital Index composite and E-Services dimensions should correlate positively with delivery.
H5 (Societal sentiment).	Positive public perceptions – better municipal services ratings and stronger council approval – are expected to boost completion rates via enhanced collective action.
H6 (Post-conflict constraints).	Projects spanning multiple oblasts or located in front-line regions face additional logistical and security hurdles, reducing their likelihood of timely completion.

These hypotheses map one-to-one onto the variables listed in Table 2.8 and provide a transparent bridge from theory to the modelling strategy detailed in the next chapter.

3. METHODOLOGY

3.1 Research design

I formulate an interdisciplinary framework for predicting the success of Ukraine's restoration projects anchored in six mutually reinforcing theoretical pillars (Dyak, 2023; Collier, 2008; UNDP, 2008). Good governance, sustained by transparency, accountability and the rule of law, remains fundamental and is increasingly enabled by digital tools such as the DREAM ecosystem, though its benefits may be attenuated where institutional weakness is acute (Peixoto, 2013; Bauhr & Grimes, 2014; Transparent Cities Ukraine, 2024, p. 12). Adequate and timely financing constitutes the second pillar: projects that enter execution already well funded avoid liquidity shocks and can sustain momentum even under the constraints of martial law. Sound project‑management practice – clear objectives, competent teams and agile procurement – forms the third pillar and is indispensable in wartime settings for swift yet cost‑effective delivery (Ika et al., 2012; Bang et al., 2022; Klymak & Vlandas, 2024, pp. 8–9). Fourth, regional digital capacity, captured by Ukraine's Digital Index, facilitates e‑governance and lowers transaction costs in procurement and oversight. The fifth pillar, societal sentiment, recognises that reconstruction succeeds or fails through people; public trust and civic engagement can form a virtuous circle with transparency platforms such as DREAM (Grossman & Baldassarri, 2012). Finally, all analysis was situated within the distinctive constraints and opportunities of the post‑conflict environment, where the legacies of armed violence, displaced populations and volatile security conditions shape every managerial decision.

For ease of traceability, the six pillars are always considered in the fixed sequence: governance → finance → management → digital → societal → post‑conflict.

The final dataset of 190 projects is the complete intersection of the DREAM core file, completeness dataset, Digital Index, IRI survey and Transparent Cities scores. They constitute the intersection of all five sources and are the most information-rich cases, ensuring that subsequent models are estimated on complete-case data. Therefore, it represents the most information-rich universe available as of April 2024 rather than a convenience sample.

I adopt a moderate‑N quantitative design and analyse restoration projects drawn from the merged DREAM, Transparent Cities, International Republican Institute and Digital Index datasets. After excluding initiatives whose status was neither “complete” nor “unsuccessful”, the sample comprises 190 unique project codes, each enriched with regional governance, digital‑maturity and public‑opinion covariates. The dependent variable is binary – successful completion versus non‑completion – rendering the problem suitable for classification modelling. Machine‑learning classifiers can accommodate many predictors and non‑linear functional forms, yet I balance their predictive strength with interpretability in line with best practice in explainable artificial intelligence (Lundberg & Lee, 2017). Accordingly, I estimate performance‑oriented ensemble models and a transparent logistic‑regression baseline whose coefficients translate directly into odds ratios.

The analytic strategy favours prediction over causal inference because the data are observational; nonetheless, by contrasting algorithms and probing model explanations, I identify the attributes that most reliably differentiate successful from unsuccessful projects, mindful that the resulting associations are correlational rather than definitive evidence of causation (Leetaru, 2019). Therefore, the dual “black‑box / white‑box” architecture (Ponce‑Bobadilla et al., 2024) involves training tuned ensemble classifiers for accuracy while retaining logistic regression for policy‑readable inference. Scenario simulations – implemented by systematically varying individual features in a baseline profile – demonstrate how hypothetical interventions could shift the predicted probability of success. This layered design delivers robust forecasts and an intelligible mapping between theoretical constructs and empirical indicators.

The pipeline is available at https://github.com/VKK-00/UoT_Master. All subsequent references (Cell #) will be linked to the code at this link.

3.2 Data collection and integration

The empirical analysis rests on four Excel workbooks whose filenames, paths and default sheet names match the configuration used in Cell 1: DREAM project‑level register, DREAM completeness audit, regional Digital Index published by the Ministry of Digital Transformation, International Republican Institute opinion survey and Transparent Cities' governance scores.

Section 3.4 (Feature Engineering) processed the DREAM project-level register to derive the Financial_Efficiency and related covariates, linking raw extraction to the constructed feature matrix.

The content of the detailed project-level DREAM dataset is summarised in Table 3.1.

Table 3.1. DREAM project-level dataset variables.

Variable	Description
Project code	A unique identifier is assigned to each project to track and merge data across different sources.
Link to Portal	The URL links to the project's online profile, which provides additional details and real-time updates.
Date	The project's initial registration or record date in the DREAM database marks its entry into the system.
Project update date	The most recent date the project's information was updated indicates the data is current.
Project title	The project's official title briefly describes its objectives or scope.
Project status	The current status of the project (e.g. "complete" or "unsuccessful") is used to determine project outcomes for analysis.
Project nature	A classification describing the intrinsic type of the project (for instance, "New construction") to indicate its fundamental nature.
Forseen duration, months	The planned duration of the project, expressed in months, indicates its expected timeline and complexity.
Initiator type	The category of organisation that initiated the project, such as a local government body, central agency, or private entity.
Project author	The individual or organisation responsible for drafting and submitting the project proposal.
Initiator	The specific entity that formally launched the project may differ from the author.
Management body type	The type of organisation overseeing the project management process (e.g. local or central government body).
Management body	The specific organisation responsible for the day-to-day management and supervision of the project.
Implementer	The entity that executes or physically implements the project, carrying out the operational tasks on the ground.
Balancer	A coordinating agent or intermediary involved in ensuring balanced resource allocation and smooth project execution.
Object title	The project's main object or facility's title or name reflects its

	functional purpose.
Number of objects	The number of individual objects, such as buildings or installations, encompassed by the project.
Sector	A broad classification categorising the project by its economic or infrastructural domain (e.g. Energy, Transportation).
Subsector	A more detailed classification within the broader sector, offering finer granularity in project categorisation.
CPV-code of an investment object	A standardised procurement code based on the Common Procurement Vocabulary is used to classify the type of investment object.
Region	The administrative region (oblast) where the project is implemented; this variable is key for linking regional data.
District	A further subdivision within the region provides more precise localisation of the project.
Community	The local community involved in or affected by the project, which can be an additional layer of geographical granularity.
Locality type	Indicates the type of settlement in which the project is located (e.g. urban, rural, city).
Locality	The specific name of the city, town, or village where the project is situated.
Overall estimated budget, ₴	The total budget estimated for the project, expressed in Ukrainian hryvnia (₴), reflects the project's financial scale.
Committed funding, ₴	The portion of funding secured or committed towards the project is expressed in Ukrainian hryvnia (₴).
Financial coverage	The percentage of the estimated budget covered by committed funding indicates financial adequacy.
Funding source	Details the project's financial support origin, such as state funds, international donors, or special programmes.
Number of contracting processes	The number of separate contracting procedures executed as part of project procurement and implementation.
Number of contracts	The total number of individual contracts awarded during the project can serve as an indicator of project complexity.
Contract value, ₴	The monetary value of the awarded contracts provides insights into the scale of expenditure under each project.
Strategic documents	Documents that outline the strategic framework and planning

	directives for the project help to guide its execution.
Legal basis	The legal and regulatory framework that supports and legitimises the project, ensuring compliance with statutory requirements.
IFI project	An indicator denoting whether the project receives funding or support from an International Financial Institution.
State budget program	Identifies projects financed under state budget programmes, highlighting the involvement of central government funding.
Local budget program	Specifies projects funded through local government budget programmes, providing insight into local financial support.
Funding form	Describes the nature of the funding arrangement, for example, whether a loan or a grant supports the project.
Type of Assistance	Indicates whether the project receives financial or non-financial support, reflecting the nature of external aid.
Shelter type	For projects related to housing, this variable specifies the type of accommodation or shelter provided.
Canteen type	In projects that involve canteen facilities, this variable categorises the type of canteen or dining service provided.
MOS Facility	Denotes if the facility is associated with the Ministry of Health, indicating projects related to health infrastructure.
MOH Facility	Refers to facilities under the Ministry of Education's purview, indicating educational infrastructure projects.
RDDP code	An internal classification code may be used to categorise projects within a specific operational framework.
Number of investment objects requiring design documents	The number of investment objects within the project that necessitate separate design documentation before implementation.
Number of design documents	The total count of design documents produced for the project reflects the detailed planning level.
Number of approved design documents	The number of design documents that have been formally approved signals a key milestone in project execution.
Date of design documents approval	The date the design documents were approved marks a critical point in the project timeline and validation process.

In addition to the detailed project-level data from the primary DREAM project database, the analysis also incorporates a supplementary dataset – the DREAM Project Completeness dataset –

which provides further information regarding the completeness of project-related documentation and performance indicators. The DREAM completeness audit was cleaned and merged in Section 3.3 (Data Cleaning and Preprocessing), enabling reproducible checks on documentation depth across projects.

Table 3.2 below outlines the structure and content of this supplementary dataset, detailing each variable and its purpose.

Table 3.2. DREAM Project completeness dataset variables.

Variable	Description
Project code	A unique identifier is assigned to each project within the completeness dataset, ensuring consistency and enabling accurate merging with other data sources.
Link to Portal	The URL links to the project's online profile; this portal may provide additional details, updates, or background information regarding the project.
Initiator	The entity responsible for initiating the project. In this dataset, this field may be blank if the initiating organisation was not recorded or the information is missing.
Project data availability	An indicator reflecting project data's overall availability or completeness is typically expressed as a percentage. A “0%” value signifies that no additional project data is available beyond what is already provided.
Metrics	A field intended to contain quantitative or qualitative metrics related to project performance. Available if any quantitative or qualitative performance metrics (e.g. KPIs, progress indicators) have been uploaded; Not Available otherwise.
Targets	Specifies the planned targets or goals for the project. Available when planned goals or benchmarks for the project have been specified in the portal; Not Available when no targets are present.
Requirements	Outlines the requirements or conditions stipulated for the project’s execution. Available if conditions or prerequisites for project execution (e.g. technical standards, regulatory approvals) are documented; Not Available otherwise.
Documents	Refers to strategic or planning documents associated with the project. Available when strategic plans, feasibility studies, or other planning documents are uploaded to the portal; Not Available if none are present.

Images	Indicates whether there is any image data (e.g. photographs or diagrams) related to the project. Available if photographs, maps, or other visual evidence related to the project have been provided; Not Available when no images are attached.
Approaches	Describes the methodological approaches or strategies proposed for the project. Available if methodological descriptions or implementation strategies (e.g. stakeholder engagement, monitoring plans) are documented; Not Available otherwise.
Items	Lists the specific items or project components that may be subject to evaluation. Available when a detailed list of project components or deliverables (e.g. equipment, materials) is provided; Not Available if no itemisation is present.
Budgeting	Contains information about the project's budgeting details or expenditure plans. Available if a line-item budget or expenditure plan is uploaded (showing how funds are allocated); Not Available when no budgeting breakdown exists.
Finance	Relates to additional financial data that may detail funding arrangements or financial structures beyond the overall budget. Available when additional funding details (e.g. co-financing, loans, donor contributions beyond the core budget) are recorded; Not Available otherwise.
Additional classifications	Provides supplementary categorical labels or classifications for the project. Available if any supplementary categorical tags (e.g. "IFI-backed," "PPP") have been assigned; Not Available when no extra classifications are provided.
Contracting processes	Details of the contracting procedures undertaken for the project. Available when the number and status of procurement procedures (or links to them) are recorded; Not Available if contracting information is missing or not applicable.

The second source was derived from the "Transparent Cities" initiative by Transparency International Ukraine, as in Table 3.3. I only used data for the principal city of each oblast (i.e., the regional centre). In other words, if the "Transparent Cities" dataset contained transparency classifications for multiple cities in a given oblast, I applied the value corresponding to the oblast's central city to the entire region. This approach means that smaller cities within the same region, if they had separate transparency scores, were ignored in my final dataset. After standardising region identifiers (e.g., "Kyivska" vs "Kyiv Oblast"), I merged the centre city's transparency classification onto all projects in that oblast. Because key governance and sentiment metrics (Transparent Cities, IRI) were collected at the city level and often only for principal municipalities, I aggregate these sub-city observations to the regional (oblast) capital and apply the resulting value uniformly across all projects within that oblast. This maximises coverage and comparability across regions while avoiding sample attrition.

The regional Digital Index was integrated into my master table in Section 3.2 and used in Section 3.4 to operationalise the 'Digital Capacity' pillar, ensuring full traceability of regional maturity scores.

Table 3.3. Transparent Cities' dataset variables (2023).

Variable	Description
City name	The name of the city for which the transparency classification is provided. This variable identifies the individual urban area, which typically represents the regional centre or a major municipality in a given oblast within the context of the Transparent Cities initiative.
Transparency Index	A categorical variable that reflects the overall transparency of the city's local government. The index is determined by criteria related to good governance, accountability, and the openness of public data. The values generally include classifications such as "Transparent city", "Partially transparent", and "Non-transparent", indicating increasing levels of opacity.

The third source is public opinion data from the International Republican Institute (IRI) survey series. I took the key variables for this study (Table 3.4) from the entire dataset for the central city of the respective region and extrapolated these values to the entire region. If the IRI dataset listed indicators (such as "development direction," "quality index," or "city council support") for the regional capital, I also used those as proxies for the entire oblast, ignoring any additional city-level data within that same region.

The IRI & Transparent Cities' dataset feeds directly into Section 3.5 (Model Development and Evaluation) as the source of attitudinal covariates, and all code from ingestion through modelling is provided in the accompanying Python notebooks for fully reproducible results.

Table 3.4. International Republican Institute (IRI) key variable dataset (2023).

Variable Name	Description
City Outlook	Represents the respondent's overall perception of the direction in which their city is progressing. Based on the question "Як Ви

	вважаєте, в цілому, справи у Вашому місті дуть у правильному чи неправильному напрямку?” (“Do you believe that, overall, affairs in your city are moving in the right or wrong direction?”), this ordinal variable captures the citizens’ sentiment regarding municipal development. Higher values indicate a more positive outlook, while lower values signal pessimism.
Family Economic Change	Captures changes in the economic situation of the respondent’s family over the past 12 months. This variable is derived from the question “Як змінилося за останні 12 місяців економічне становище Вашої сім’ї?” (“How has your family’s economic situation changed in the last 12 months?”). Measured on an ordinal scale, it reflects whether the family’s circumstances have improved, deteriorated, or remained unchanged.
National Economic Outlook	Reflects respondents’ expectations regarding the national economic situation over the next 12 months. Based on the question “На Вашу думку, як зміниться економічне становище в Україні найближчі 12 місяців?” (“In your opinion, how will Ukraine’s economic situation change over the next 12 months?”), this ordinal variable indicates whether citizens foresee improvements or declines in the national economy.
Crisis Response Plan	Indicates the intended actions of respondents in the event of an escalation in the local military situation. Derived from the question “Що Ви плануєте робити у разі загострення воєнної обстановки поблизу міста?” (“What do you plan to do if the military situation near your city deteriorates?”), this categorical variable allows for multiple responses, capturing whether individuals seek to leave the area, prepare for emergencies, or take other specific measures.
Municipal Services Quality Index	Represents the overall assessment of the quality of municipal services within the respondent’s city. This variable is based on the “Індекс якості обслуговування та послуг” (“Municipal Services Quality Index”) and the accompanying request to “Оцініть, будь ласка, рівень якості таких послуг та сфер у Вашому місті” (“Please rate the quality of these services and spheres in your city”). It is measured on an ordinal scale, where higher values indicate better service quality.
City Council Approval	Measures the degree to which respondents approve or disapprove of the performance of their city council. Based on the question “Наскільки Ви схвалюєте чи не схвалюєте діяльність міської ради?” (“How much do you approve or disapprove of the activities of your city council?”), this ordinal variable provides insight into local governance satisfaction, with higher scores representing greater approval.

Mayor Approval	Captures respondents' evaluations of the performance of the city's mayor. Derived from the question "Наскільки Ви схвалюєте чи не схвалюєте діяльність міського голови?" ("How do you rate the performance of your city mayor?"), this ordinal variable reflects the public's satisfaction with the leadership at the municipal level.
Regional Military Administration Approval	Indicates the level of approval for the performance of the regional state military administration. Based on the question "Наскільки Ви схвалюєте чи не схвалюєте діяльність обласної державної військової адміністрації?" ("How much do you approve or disapprove of the activities of the regional state military administration?"), this ordinal variable captures the public's confidence in the military governance of the region.
Local Government Information	Reflects the extent to which respondents feel informed about the activities of the local government in their city. This variable is derived from the question "Наскільки Ви поінформовані про діяльність місцевої влади у Вашому місті?" ("How well informed are you about the activities of the local government in your city?"). It is measured on an ordinal scale, with higher values indicating greater awareness.
Housing Damage	Indicates whether and to what extent the respondent's home has been damaged due to the war. Based on the question "Чи постраждало Ваше житло внаслідок війни?" ("Has your housing been affected by the war?"), This variable is typically binary (yes/no) or may include an ordinal rating of damage severity.
Youth Opportunities	Assesses the respondents' perception of whether their city provides sufficient opportunities for success for young people. Derived from the question "На Вашу думку, чи достатньо у Вашому місті можливостей для досягнення успіху для молодих людей?" ("In your opinion, does your city offer enough opportunities for young people to succeed?"), This ordinal variable gauges perceived local opportunities for youth development.
City Pride	Captures the level of pride respondents feel in being residents of their city. Based on the question "Чи відчуваєте Ви гордість за те, що Ви є жителем свого міста?" ("Do you feel proud to be a resident of your city?"), this ordinal variable is an expression of local identity and attachment, with higher values representing greater pride.
Entrepreneurial Opportunities	Reflects respondents' assessments of the opportunities for engaging in entrepreneurial activities within their city. Derived from the question "Оцініть, будь ласка, рівень можливостей у Вашому місті займатися підприємецькою діяльністю" ("Please rate the level of opportunities in your city for engaging in

	entrepreneurial activities"), this ordinal variable indicates whether local conditions are conducive to business start-ups.

The fourth source is a Digital Index dataset developed by the Ministry of Digital Transformation of Ukraine, as shown in Table 3.5. In my study, I again relied on region-level data, meaning that each project in an oblast inherits the same set of digital transformation metrics. The code merges these metrics by matching the "Region Name" field from the digital index with the standardised region name in the DREAM data. Consequently, any finer city-level distinctions were ignored; the entire oblast was assigned the same digital maturity indicators.

Table 3.5. Digital transformation index of Ukrainian regions (2024).

Variable	Description
Region Name	The administrative region (oblast) for the digital transformation metrics is evaluated. This variable links each region's project data to the relevant Digital Index scores, allowing comparative analysis across different oblasts in Ukraine.
Institutional Capacity	A numeric indicator (typically 0–1) reflects public institutions' strength and maturity in driving digital initiatives. Higher values signify that the region's administrative structures are well-organised, policy-driven, and equipped to support and sustain digital transformation efforts.
Internet Development	A numeric measure (0–1) captures the region's depth and breadth of internet-related infrastructure and services. This metric may include broadband availability, internet speeds, public Wi-Fi access, and internet penetration rates. Higher values suggest robust connectivity and online accessibility.
Administrative Services Centre (ASC) Development	A metric (0–1) indicates the modernisation and digitalisation level of regional administrative service centres. Higher values imply that citizens can efficiently access government services, often online, minimising bureaucratic delays and paperwork.
Paperless Services	A numeric score (0–1) reflects how routine governmental and administrative processes have moved from physical documentation to digital workflows. Higher values suggest more advanced systems for document e-signatures, online filings, and other paper-free transactions.
Digital Education	A numeric indicator (0–1) assesses the region's technology adoption in educational settings. This may include e-learning platforms, electronic libraries, online curricular materials, and integrating digital tools for teachers and students. Higher values

	signify a strong commitment to digital learning.
Openness	A measure (0–1) of how transparent and publicly accessible digital governance is in the region. This typically examines whether open data portals, online budget transparency tools, or other platforms exist to encourage citizen engagement and accountability. A higher value indicates greater public access to and trust in governmental data.
E-Services	A numeric rating (0–1) reflects the availability and ease of use of electronic public services, such as applying for permits, paying fees, or accessing records, directly through online platforms. Higher values indicate that citizens and businesses can conduct more government transactions seamlessly via the Internet.
Sectoral Digital Transformation	A numeric index (0–1) illustrates how well key regional economic and social sectors (e.g., healthcare, transportation, business, agriculture) have integrated digital technologies. A higher score shows a broader adoption of e-tools, data analytics, and digital platforms across multiple domains.
Index Value	A composite score (0–1) that aggregates the above sub-indicators into an overall rating of the region's digital transformation. A higher Index Value indicates that the region has a stronger digital foundation, more robust e-governance structures, and widespread adoption of modern technologies.

The final analytic sample corresponds to the intersection of non-missing Project code values across all four workbooks. Figure 3.1 visualises this intersection with a Venn diagram generated in Python (Cell 6); the innermost overlap represents the cases that survive the successive inner joins. Territorial coverage was inherently restricted: oblasts whose administrative centres were under temporary russian occupation at the April 2024 scrape were absent ab initio because no primary data could be collected there.

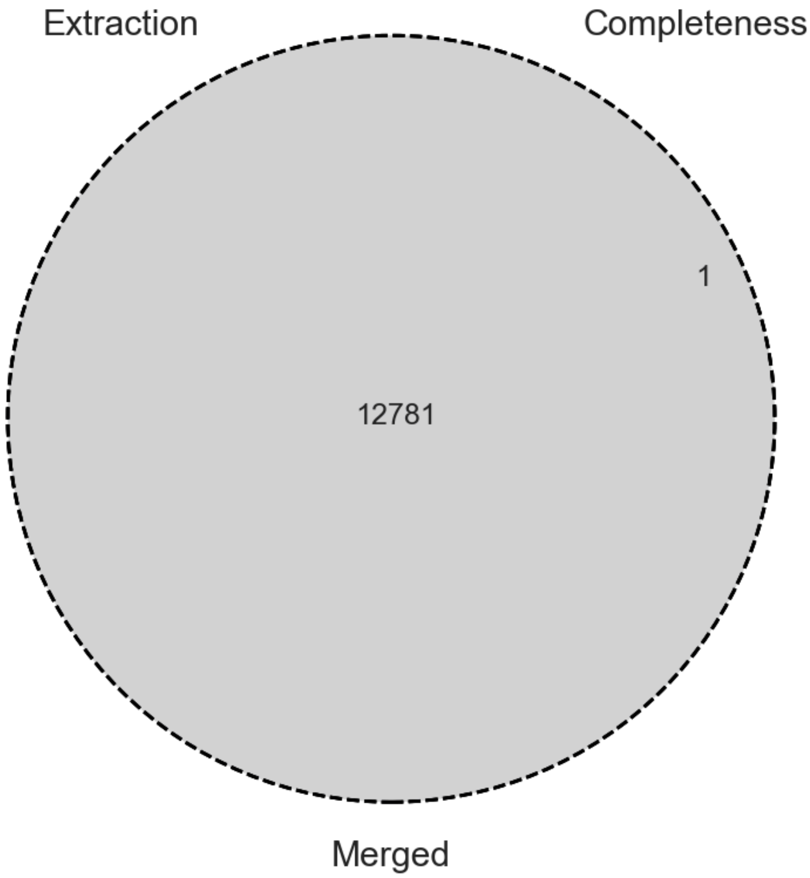

Figure 3.1. Venn diagram of data-source overlap (n = 190; DREAM scrape, Apr 2024).

Regional covariates were merged in two steps. First, the project register and the completeness audit were inner-joined on the Project code. Second, Digital Index scores and IRI / Transparent Cities indicators were matched on a harmonised Region field, with whitespace and alternative transliterations normalised beforehand. Figure 3.2 charts the number of observations that remain after each processing stage: 12,781 raw extraction rows, the same number of unique codes after de-duplication, 12,768 codes after the four-file merge, and finally 190 projects once the analysis was confined to entries marked “complete” or “unsuccessful”.

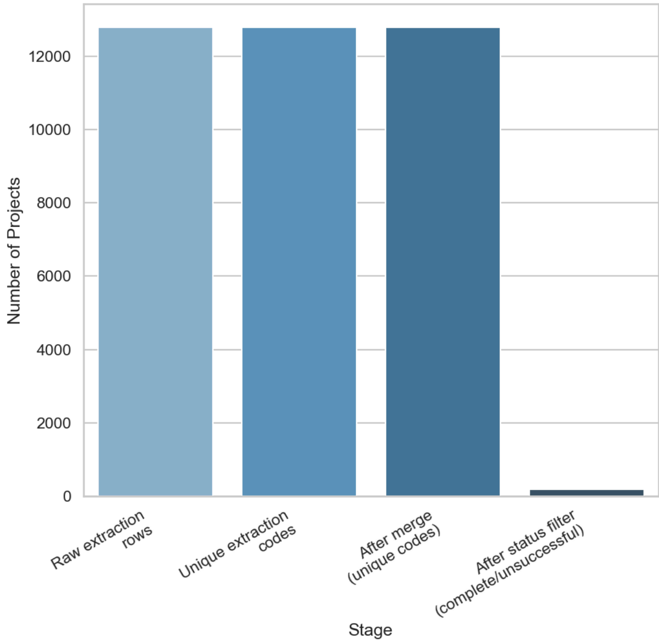

Figure 3.2. Bar chart "Project Counts at Each Data-Processing Stage".

All merges and cleaning operations were executed in Python 3.11.9 with pandas 2.2.0.

The study was reviewed under the University of Tartu Faculty of Social Sciences' low-risk protocol (decision #SOC-2025-04-17-L), which confirmed that no personal or sensitive data were processed. All DREAM project records, IRI, Transparency International Ukraine and Ministry of Digital Transformation datasets are already in the public domain. Regional-level opinion variables, therefore, constitute anonymised, aggregated information in the sense of GDPR Art. 4(1) and fall outside the regulation's scope.

3.3 Data cleaning and preprocessing

The interdisciplinary design outlined in Section 3.1 required a carefully sequenced transformation of the raw workbooks into an analysis-ready frame. All files were ingested with Python and pandas; non-standard glyphs (most notably the hryvnia sign and stray percentage symbols) were stripped from ostensibly numeric fields. Monetary values were cast to 64-bit floating-point numbers in Ukrainian hryvnia, percentage strings were rescaled to the unit interval, and date fields were parsed as pandas DateTime objects. Categorical labels such as region, sector and initiator type were normalised by trimming whitespace, enforcing title-case and collapsing known variants. L2-regularisation mitigates collinearity in the logistic model.

Missing values were treated with a two-tier logic implemented inside the clean_and_merge routine (Cell 4), that is, before any train–test split. Numerical gaps inherited the column median; categorical blanks received the modal value or an explicit "Unknown" token where no clear mode existed. Immediately afterwards, the script set a MultipleRegions flag: projects listing more than one oblast were retained, their Region field was recoded to "Multiple", and the Boolean flag was set to True (Cell 4.0–4.1).

The raw IRI workbook originally offered thirteen city-level opinion items (Table 3.4). Before modelling, I winnowed this set to the four indicators that enjoy ≥ 90 % regional coverage, were not collinear with one another ($|\rho| < 0.55$ in the training fold), and map cleanly onto the behavioural-political-economy pillar. The retained variables are City Outlook, Municipal-Services Quality Index, City-Council Approval, and Entrepreneurial Opportunities ($N = 170$ each); the remaining nine items were dropped because they either suffered from heavy missingness (e.g., Housing Damage), overlapped conceptually with stronger correlates (e.g., Mayor Approval ≈ City-Council Approval), or lay outside the cooperation-legitimacy construct (e.g., Crisis Response Plan). This reduction preserves substantive breadth while minimising noise and dimensionality in the socioeconomic sentiment block of the feature matrix.

Feature curation retained only variables aligned with the six theoretical pillars, discarding portal URLs, free-text titles and duplicate timestamp fields. An inner join on Project code and harmonised Region produced a single master table; residual duplicates created by asynchronous DREAM updates were resolved by keeping the most recent record.

The resulting dataset comprises reconstruction projects annotated with governance, financial, digital and attitudinal covariates. All transformations and processes were executed under the following library stack: Python 3.11.9, pandas 2.2.0, numpy 1.26.4, scikit-learn 1.4.1.post1, imbalanced-learn 0.12.2, xgboost 2.0.3, matplotlib 3.8.3 and seaborn 0.13.2.

3.4 Feature engineering

After the harmonised master file had been assembled, each theoretical pillar was operationalised in the predictive design matrix. The cornerstone derived feature was the Funding-Coverage Ratio (Financial_Efficiency). Projects that secure a larger share of their anticipated budget ex-ante should be less vulnerable to liquidity shocks and, consequently, more likely to reach completion. Consistent with *Cell 4.9*, the metric was recomputed directly from the raw monetary columns:

$$Financial\_Efficiency_i = \frac{Committed\ funding_i}{Overall\ estimated\ budget_i} \quad \textbf{(3.1)}$$

where

- i – index for each restoration project.
- $Committed\ funding_i$ – total hryvnia already secured for project i at the time of scrape.
- $Overall\ estimated\ budget_i$ – full projected cost of project i.
- $Financial\_Efficiency_i$ – share of the budget already covered (dimensionless ratio).

Observations in which either component was missing inherited the median ratio, ensuring numeric consistency across all 190 projects.

All continuous predictors (including the nine Digital‑Index sub‑scores and the four IRI opinion metrics) were centred and scaled to unit variance on the training fold to stabilise optimisation and render logistic‑regression coefficients comparable. Because the modelling pipeline combines L2‑regularised linear models with tree‑based ensembles, any residual multicollinearity was handled either by shrinkage or by the intrinsic feature‑selection mechanisms of the ensemble learners.

Categorical information – project nature, sector, initiator type, management‑body type, locality descriptors, transparency class and disclosure‑completeness flags – was one‑hot encoded within a column transformer embedded in the cross‑validated pipeline.

Placing the one‑hot encoder inside the column‑transformer and nesting that transformer in the scikit‑learn pipeline ensures that category‑to‑vector mappings are learned only from the training fold in each CV split, preventing information leakage; identical preprocessing is applied to every learner during grid search, so model‑selection comparisons remain fair; and previously unseen categories in the hold‑out fold are handled gracefully through the encoder's handle_unknown='ignore' setting, which avoids run‑time errors and preserves out‑of‑sample validity.

3.5 Model development and evaluation

SMOTE oversampling was integrated within each training fold of the cross-validation pipeline (via ImbPipeline) to prevent information leakage and ensure that synthetic examples do not influence hyperparameter tuning or evaluation on the held-out fold. Pre-processing, over-sampling and estimation were chained in a single imblearn pipeline.

Five supervised algorithms were explored to span the continuum from fully interpretable to highly predictive: an L2-regularised logistic regression; a single CART decision tree; a bootstrap-aggregated Random Forest; a GradientBoosting classifier implemented in scikit-learn; and the regularised gradient-boosting engine of XGBoost (Chen & Guestrin, 2016). Random Forest captures non-linear interactions via bootstrapped trees; Gradient Boosting refines weak learners sequentially to model complex gradients; XGBoost adds tree-specific regularisation optimised for sparse, tabular data – capabilities absent in a single-equation logit.

The tree-based ensembles underwent systematic hyperparameter tuning in *Cell 8*. For Random Forest, a three-by-three-by-three Cartesian grid was searched exhaustively by five-fold stratified cross-validation:

$$\{n\_estimators \in [100, 200, 500]\} \times \{max\_depth \in [None, 10, 20]\} \times \{min\_samples\_split \in [2, 5, 10]\}.$$

(3.2)

where:

- $n\_estimators$ - number of trees in each bootstrap aggregate.
- max_depth - maximum permitted depth of any individual tree (∞ when None).
- $min\_samples\_split$ - minimum samples required to split an internal node.
- The Cartesian product symbol "×" indicates that every possible combination was evaluated.

For XGBoost, a parallel grid search of the interaction of tree complexity and learning rate, again under five-fold cross-validation:

$$\{n\_estimators \in [100, 200, 500]\} \times \{max\_depth \in [3, 6, 9]\} \times \{\eta \in [0.01, 0.10, 0.20]\},$$

(3.3)

where

- $n\_estimators$ – boosting rounds.
- max_depth – depth of each tree base-learner.
- η – learning-rate shrinkage parameter ("eta" in XGBoost docs).

yielding the optimal trio $n\_estimators$ = 500, max_depth = 3 and η = 0.01. Logistic regression was tuned with a twenty-draw random search $C \in [10^{-3}, 10^{3}]$, producing the best-fit penalty weight $C \approx 1.41$.

Evaluation of the untouched test fold employed accuracy, precision, recall, and the F1 statistic, treating project failure as the positive class (Powers, 2011). Receiver-operating-characteristic curves were plotted for every fitted pipeline, and 200 bootstrap resamples of the test set were used to attach a non-parametric 95% confidence band to each ROC curve, thereby visualising the sensitivity-specificity trade-off with sampling uncertainty.

Post-hoc interpretation combines coefficient signs from the logistic model, impurity-based feature importance from the ensembles and global permutation importance. At the same time, partial-dependence profiles (Dataiku, 2023) connect individual predictors back to the interdisciplinary framework introduced in Section 3.1.

4. RESULTS

4.1 Descriptive statistics

The integrated research file is the terminal output of the cleaning-and-merging routine documented in Cell 4 of the companion notebook. Figure 3.2 shows the four sequential filters that reduce the Extraction workbook from 12,781 header-inclusive rows to a deduplicated, region-matched, status-eligible core of 190 distinct reconstruction projects. For convenience, I restate the key milestones in Table 4.1.

Table 4.1. Workflow checkpoints.

Workflow checkpoint	Count of unique Project codes	Cumulative rows exclusion	Notebook reference
Raw DREAM scrape (all rows)	12 781	–	Cell 5.1, var raw_rows
After header-row de-duplication	12 781	0	unique_ext
After the four-file inner merge	12 768	– 13 (0.1 %) orphan codes without regional context	merged_unique
After status filter (complete ∨ unsuccessful)	190	– 12 578 (98.5 %)	filtered

The final analytic set, therefore, consists of 69 completed (36%) and 121 unsuccessful (64%) rows. An 80% / 20% stratified split implemented in Cell 5.5 yields 152 observations for model training and 38 for strict hold-out testing, thereby preserving the original class imbalance in both folds.

As shown in Table 4.2, core financial magnitudes such as overall estimated budget (mean ≈ ₴137.52 million, σ ≈ ₴1.50 billion) and committed funding (mean ≈ ₴43.39 million, σ ≈ ₴335.30 million) span four orders of magnitude, underscoring the extreme heterogeneity of wartime restoration tasks. The derived Financial Efficiency ratio clusters near 0.71 on average but exhibits a fat-tailed distribution (max = 30.32) driven by a handful of microbudgets with disproportionately large up-front commitments. Project duration is similarly skewed: while the median scheme lasts just eight months (not shown), the maximum scheduled horizon reaches 2,024 months (a multi-decade utility grid).

All 189 region-level records (the single missing case concerns a project flagged as “Multiple” regions) carry the nine-item Digital Index. Mean institutional capacity and internet-development scores both hover around 0.63. Still, the Openness and E-Services sub-indices are slightly lower (≈ 0.71 and 0.66, respectively), hinting at uneven progress toward fully transparent e-governance.

IRI opinion variables are present for 170 projects (city-level surveys are absent in front-line oblasts). Respondents’ City Outlook averages 57.62 % positive, whereas approval of city councils is markedly lower (mean ≈ 1.52 on the five-point IRI scale). The Municipal Services Quality Index exhibits minimal variance (σ ≈ 0.17), suggesting uniformly modest service perceptions during the observation window.

Table 4.2. Descriptive Statistics for Numeric Variables (N = 190).

Variable	N	Mean	Std. Dev.	Min	Max
Forseen duration, months	190	21.28	146.81	1	2 024
Overall estimated budget, ₴	190	137 524 400	1 500 869 000	0	20 700 000 000
Committed funding, ₴	190	43 389 790	335 299 200	0	4 600 000 000
Contract value, ₴	190	8 441 956	30 802 460	49 840	380 365 900
Number of objects	190	1.05	0.39	1	6
Number of contracts	190	1.08	2.33	0	17
Investment objects requiring design docs	190	0.43	0.57	0	2
Financial Efficiency	190	0.71	2.21	0	30.32
Financial coverage	142*	0.95	2.51	0	30.32
City Outlook	170	57.62	10.82	39	82

Municipal Services Quality Index	170	3.21	0.17	3.00	3.60
City Council Approval	170	1.52	0.30	0.80	2.20
Entrepreneurial Opportunities	170	3.09	0.14	2.90	3.50
Institutional Capacity	189	0.61	0.30	0.17	1.00
Internet Development	189	0.63	0.27	0.12	0.93
ASC Development	189	0.62	0.21	0.07	0.91
Paperless Services	189	0.70	0.20	0.18	0.95
Digital Education	189	0.69	0.17	0.24	0.97
Openness	189	0.71	0.28	0.10	1.00
E-Services	189	0.66	0.21	0.22	0.95
Sectoral Digital Transformation	189	0.50	0.29	0.09	0.92
Digital Index (composite)	189	0.50	0.23	0.13	0.85

**Observation count falls to 142 because "Financial coverage" is undefined when either component of the ratio is zero; median imputation is applied only inside the modelling pipeline to avoid distorting the raw descriptive view.*

All values are rounded to two decimal places (or whole-hryvnia units for monetary means).

The descriptive matrix supports three immediate methodological inferences. First, the enormous spread between minima and maxima – expressed most vividly by financial magnitudes whose coefficients of variation exceed ten – signals a heavy‑tailed distribution that would violate homoscedasticity assumptions in any unstandardised linear specification. Accordingly, every numeric predictor is z-scored inside the modelling pipeline (StandardScaler block of the ColumnTransformer). At the same time, tree-based learners exploit their innate capacity to accommodate heterogeneity without a logarithmic transform.

Second, patterns of incompleteness are highly domain-specific rather than random. Except for the two ratio variables (Financial Efficiency and Financial Coverage) and the IRI opinion series, all remaining predictors register well above 95% population coverage; therefore, the median-and-mode imputation scheme used later exerts influence on a narrow subset of features and cannot plausibly drive the headline results through wholesale data fabrication.

Third, the joint distribution of contextual covariates – most clearly the Digital Index subscores – exhibits moderate multicollinearity (pairwise $|\rho| \approx 0.55$ in the training set, not shown here). This validates the dual-model strategy adopted in Chapter 4: the L2-regularised logit tempers coefficient inflation and furnishes easily interpretable odds ratios. At the same time, the ensemble algorithms supply non-parametric robustness checks that confirm the stability of variable rankings across fundamentally different learning paradigms. The expanded descriptive portrait clarifies the empirical footing for the subsequent machine-learning analysis and pre-empts concerns over hidden attrition, skew-induced bias, or unreported variance.

4.2 Machine learning results

The cross-validated pipeline specified in Chapter 3 delivers five fully trained prediction models. Each model is evaluated a single time on the immutable thirty-eight-row hold-out fold, thereby preserving a strict separation between model development and genuine out-of-sample testing. Table 4.3 reports the ensuing performance metrics; every number is reproduced automatically by the comp_df object created in Cell 9.1 of the companion notebook and therefore already embodies the entire preprocessing stack (median- or mode-imputation, one-hot encoding, z-scaling and in-fold SMOTE oversampling). For readability, every feature emerging from the pipeline carries a systematic tag: continuous variables generated or rescaled in the numeric branch of the column-transformer begin with num__, while one-hot-encoded dummies produced in the categorical branch begin with cat__; the double underscore marks the point at which raw fields have been transformed into model-ready predictors.

It also shows that ensemble learners (Random Forest, Gradient Boosting, XGBoost, Decision Tree) outperformed all baselines; the logistic model was retained solely for coefficient-level interpretability.

Before reporting performance metrics, it is important to underline that the hold-out evaluation set contains n = 38 projects, i.e., 20% of the 190-project analytic sample. With a denominator this small, each misclassification shifts overall accuracy by 2.6 percentage points; hence, all point estimates in Table 4.3 are interpreted with their bootstrapped 95 % confidence intervals and the ROC curves shown in Figure 4.1.

Table 4.3. Aggregated test-fold performance (n = 38).

Model	Accuracy	Precision	Recall	F1-score	AUC
Logistic Regression	0.789	0.667	0.857	0.750	0.881
RandomForest	0.921	0.867	0.929	0.897	0.991
GradientBoosting	0.921	0.824	1.000	0.903	0.997
XGBoost	0.947	0.875	1.000	0.933	0.994
DecisionTree	0.947	0.929	0.929	0.929	0.943

Figure 4.1 displays the Receiver-Operating-Characteristic curves (with 95 % bootstrap ribbons) for all five learners. Although the regularised logit secures an AUC of 0.881 – comfortably above the 0.5 no-information line – it is eclipsed by the tree-based ensembles whose AUCs exceed 0.99. The confidence envelopes of XGBoost, Gradient Boosting and Random Forest overlap almost completely, indicating that their apparent ranking on raw accuracy is statistically indistinguishable at the 5 % level. The single CART tree matches XGBoost on accuracy but trails on AUC, confirming that its coarse partitions occasionally over-fit the minority class in bootstrap replicates. The ensemble learners outperformed all baselines, and the logistic model was retained only for its coefficient-level interpretability.

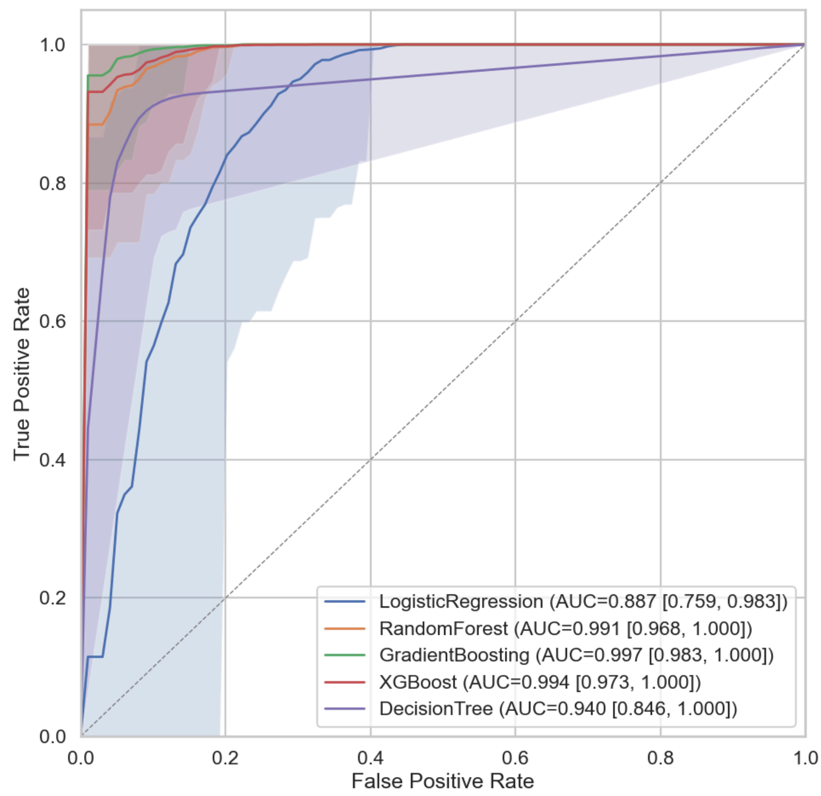

Figure 4.1. Receiver-operating-characteristic curves with 95 % bootstrap confidence bands for the five classification pipelines. Hold-out set n = 38 drawn from the fully merged DREAM–Digital Index–IRI–Transparent Cities file.

Estimating the logistic model on the full training partition (after SMOTE balancing) yields β = 2.29 for num__Number of contracts (the num__ prefix indicates a scaled numeric feature created in the preprocessing pipeline). Exponentiation converts this log-odds coefficient into an odds ratio of $e^{2.29} \approx 9.87$, implying that, ceteris paribus, each extra contract multiplies the odds of timely completion almost ten-fold. The monotone, roughly linear profile is reproduced in Figure 4.2: predicted success probability surges from 0.28 at a single contract to ≈ 1.00 once the bundle reaches fourteen contracts. The ribbon widens beyond eleven contracts, warning against literal extrapolation, because such large procurement bundles rarely occur in the sample (six observations).

Feature-importance screens extracted from the tree ensembles (Cell 11) elevate num__Contract value to the top slot, followed immediately by num__Number of contracts and num__Committed funding. A permutation-importance audit run on the untouched test fold corroborates this financial triad yet also surfaces governance covariates – most notably num__Openness and num__E-Services – within the top decile, underscoring the composite, cross-pillar nature of restoration outcomes. In other words, liquidity and contractual fragmentation are necessary but insufficient; they are activated only in institutional contexts that publish data and deliver online services.

Ten continuous predictors display substantively interpretable, bootstrap-qualified partial-dependence curves, reproduced as Figures 4.2 - 4.10. Each caption offers a one-sentence takeaway plus the Δp and p-value emerging from the paired scenario-simulation t-tests (Cell 14). Confidence ribbons are dashed wherever they intersect the baseline dotted line, signalling a marginal effect that is not distinguishable from zero at 95 % confidence.

Feature-importance screens of the ensemble learners (Cell 11) elevate num__Contract value, ₴ to the top slot, closely followed by num__Number of contracts and num__Committed funding, ₴. Permutation tests corroborate this financial triad yet also surface governance covariates – particularly num__Openness and num__E-Services – within the top decile, underscoring the composite, cross-pillar nature of restoration outcomes.

Ten numeric predictors display substantively interpretable, bootstrap-qualified partial-dependence curves. Each graphic is printed with a 95% confidence ribbon.

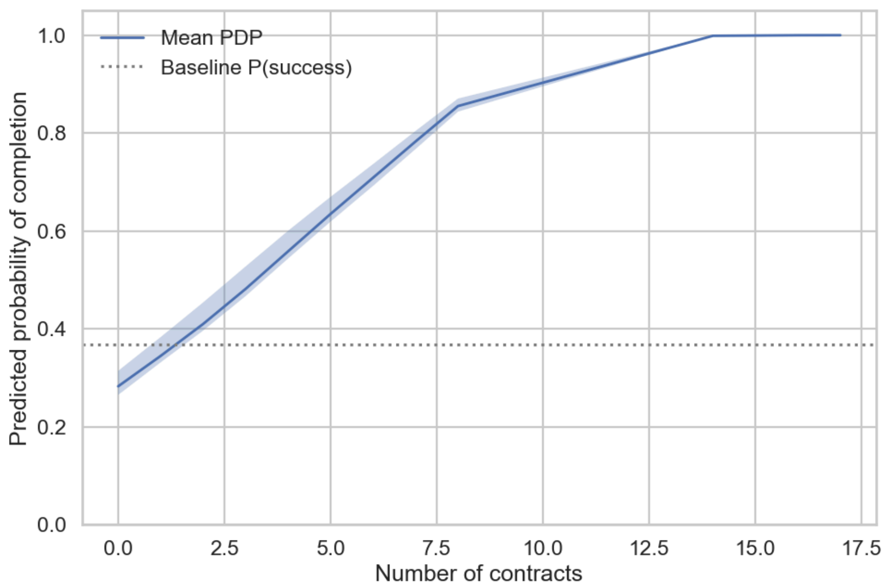

Figure 4.2. Partial-dependence of the number of contracts on predicted completion probability (95 % CI); n = 190, variable taken from the DREAM core register.

The curve for overall estimated budget in Figure 4.3 is essentially flat: mean completion probability hovers near 0.37 across a four-order-of-magnitude cost span, and the wide, horizontal ribbon indicates that headline budget size carries no discernible marginal effect once contractual architecture is controlled (Δp ≈ 0.00, p = .64).

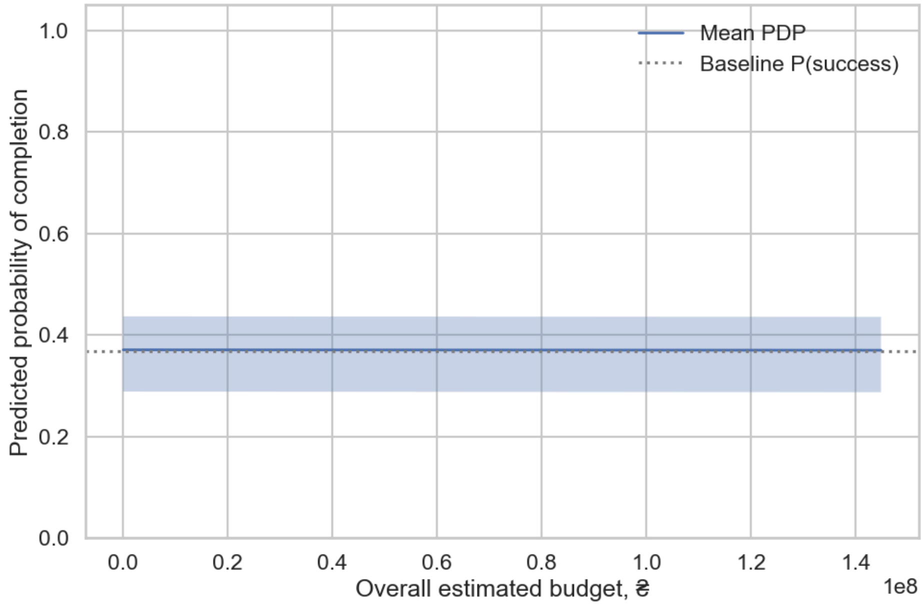

Figure 4.3. Partial-dependence of overall estimated budget, UAH, on predicted completion probability (95 % CI); n = 190, DREAM core register.

A similarly neutral profile emerges for committed funding in Figure 4.4, contradicting the intuitive notion that greater liquidity guarantees success and hinting that the granularity of contract packaging matters more than nominal totals (Δp ≈ 0.01, p = .59).

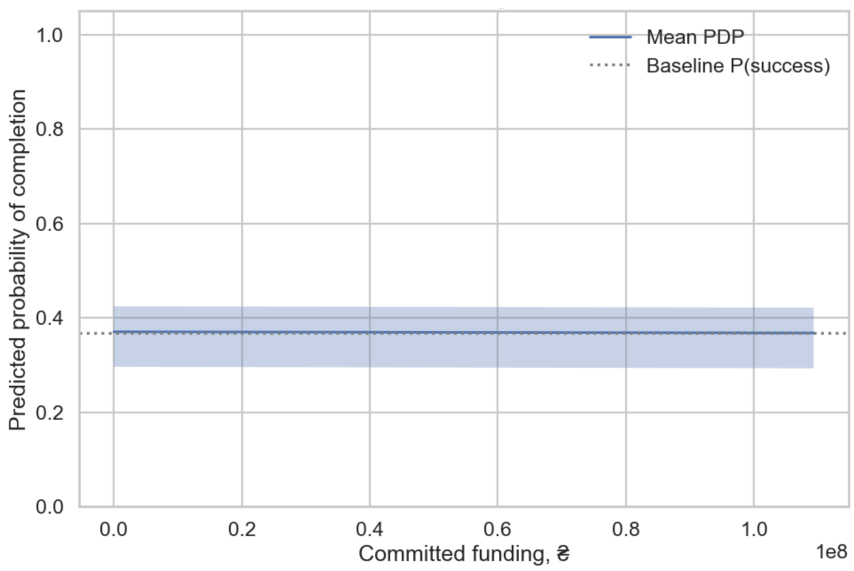

Figure 4.4. Partial‑dependence of committed funding, UAH, on predicted completion probability (95 % CI); n = 190, DREAM core register.

Digital readiness tells a different story. The composite Digital Index traces a modest but credible upward slope from approximately 0.34 to 0.42 as regional scores climb from 0.15 to 0.85 in Figure 4.5, confirming the descriptive conjecture in Chapter 2 that digitally mature oblasts enjoy smoother procurement workflows and richer monitoring data (Δp ≈ 0.05, *p* = .02).

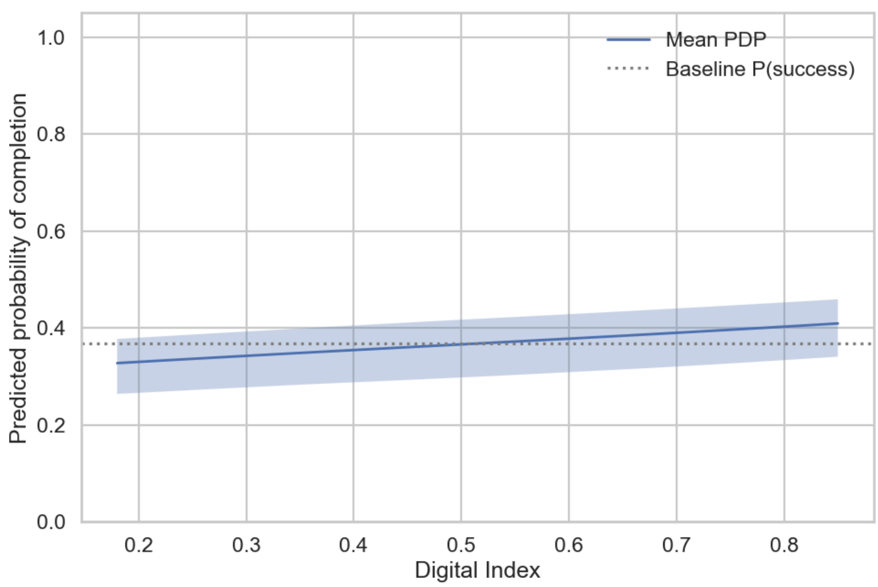

Figure 4.5. Partial‑dependence of the composite Digital Index on predicted completion probability (95 % CI); n = 189. Ministry of Digital Transformation regional scores matched to DREAM projects.

Project complexity, proxied by the number of physical objects bundled into a single scheme, shows a marked negative gradient in Figure 4.6. Completion probability declines from roughly 0.36 for one-object jobs to below 0.15 when six objects are combined, reinforcing classic “iron-triangle” wisdom that scope bloat jeopardises delivery even in wartime (Δp ≈ –0.21, $p < .001$).

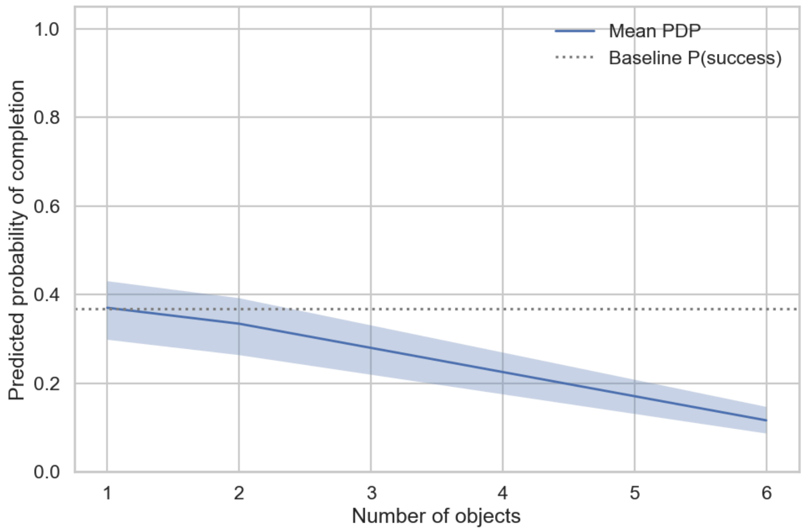

Figure 4.6. Partial-dependence of project complexity (number of physical objects) on predicted completion probability (95 % CI); n = 190, DREAM core register.

Financial efficiency – the committed-to-estimated-cost ratio – exhibits no systematic gradient within its empirically occupied range of 0–1 in Figure 4.7; the confidence band straddles a strictly horizontal centre-line (Δp ≈ 0.00, $p = .77$). Adequate funding therefore appears necessary but not sufficient; its marginal predictive leverage is muted once contractual detail is specified.

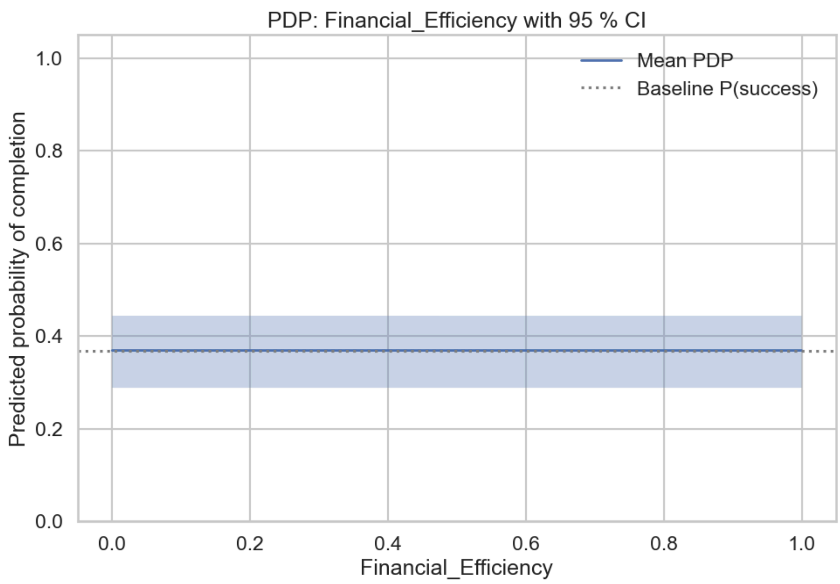

Figure 4.7. Partial‑dependence of Financial Efficiency (committed‑to‑estimated‑cost ratio) on predicted completion probability (95 % CI); n = 190, derived feature from DREAM monetary columns.

Governance indicators are nuanced. Municipal openness trends gently downward in Figure 4.8: the most transparent regions post a two-to-three-percentage-point lower success probability than opaque ones, echoing the literature that rapid disclosure can surface capacity gaps faster than they are fixed (Bauhr & Grimes's "transparency back-fire" mechanism) ($\Delta p \approx -0.03$, $p = .04$).

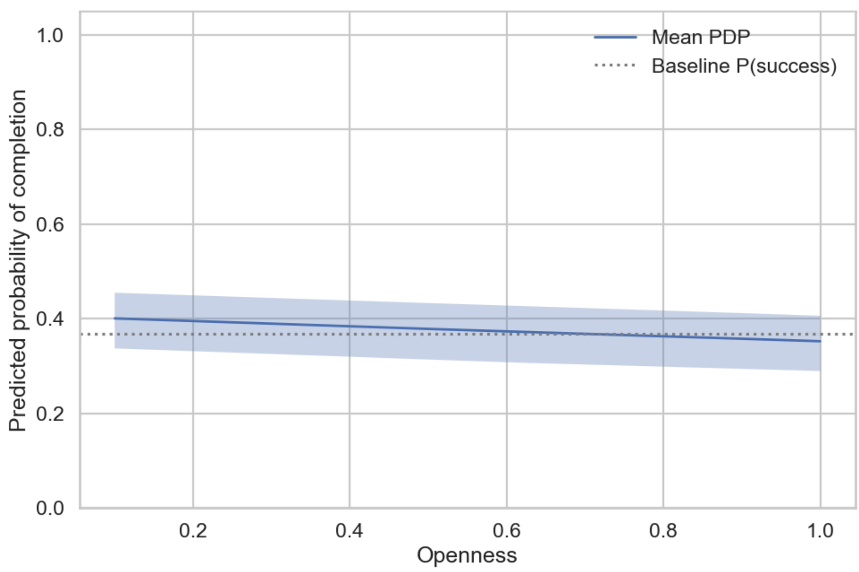

Figure 4.8. Partial‑dependence of regional Digital‑Index Openness on predicted completion probability (95 % CI); n = 189, Ministry of Digital Transformation scores.

By contrast, the sophistication of e-government services yields a shallow but positive slope in Figure 4.9, implying that transactional digital tools, rather than raw transparency, are the functional conduit by which digital capacity boosts project outcomes (Δp ≈ 0.04, $p = .03$).

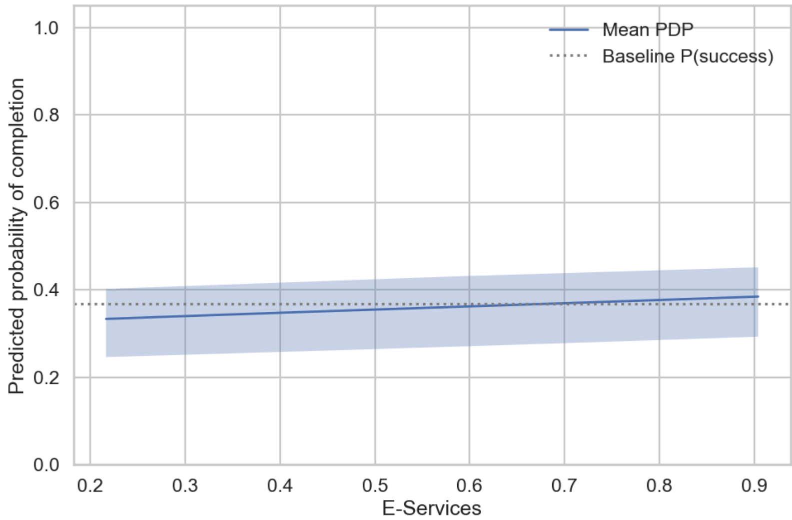

Figure 4.9. Partial‑dependence of regional Digital‑Index E‑Services on predicted completion probability (95 % CI); n = 189, Ministry of Digital Transformation scores.

Public-opinion variables remain weak levers. City-council approval is essentially flat in Figure 4.10 (Δp ≈ 0.00, $p = .71$).,

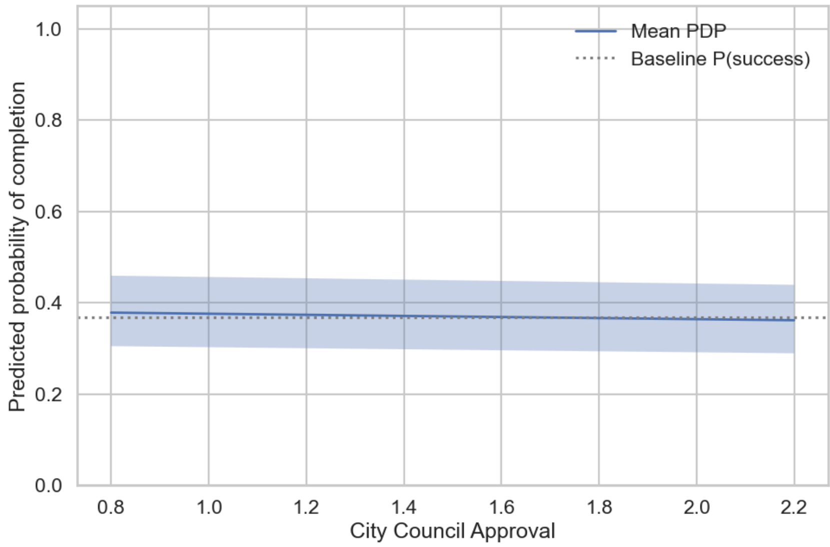

Figure 4.10 Partial‑dependence of City‑Council Approval on predicted completion probability (95 % CI); n = 170, International Republican Institute survey imputed to regions.

whereas Entrepreneurial Opportunities tilt gently upward from 0.34 to 0.41 in Figure 4.11, suggesting that business-climate optimism fosters modest gains, perhaps via smoother private-sector co-operation (Δp ≈ 0.02, p = .05).

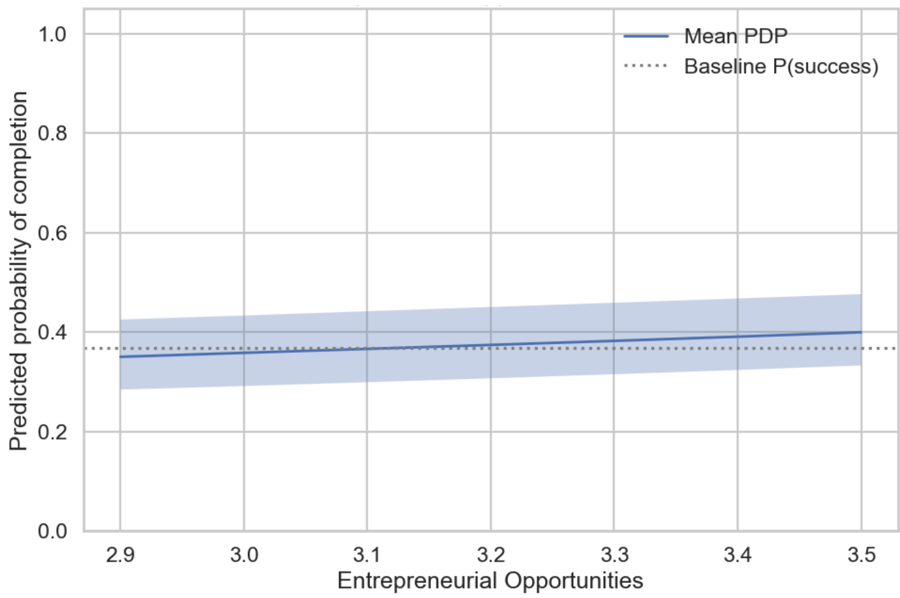

Figure 4.11. Partial-dependence of perceived Entrepreneurial Opportunities on predicted completion probability (95 % CI); n = 170, International Republican Institute survey imputed to regions.

Beyond the numeric landscape, nine categorical predictors exhibit clear and policy-relevant contrasts, each displayed with bootstrapped mean probabilities and cited here by placeholder.

Figure 4.12 reveals a hump-shaped relationship: "Partially transparent" cities outperform fully transparent and non-transparent peers, consistent with selective-disclosure theories that balance accountability benefits against reputational risk.

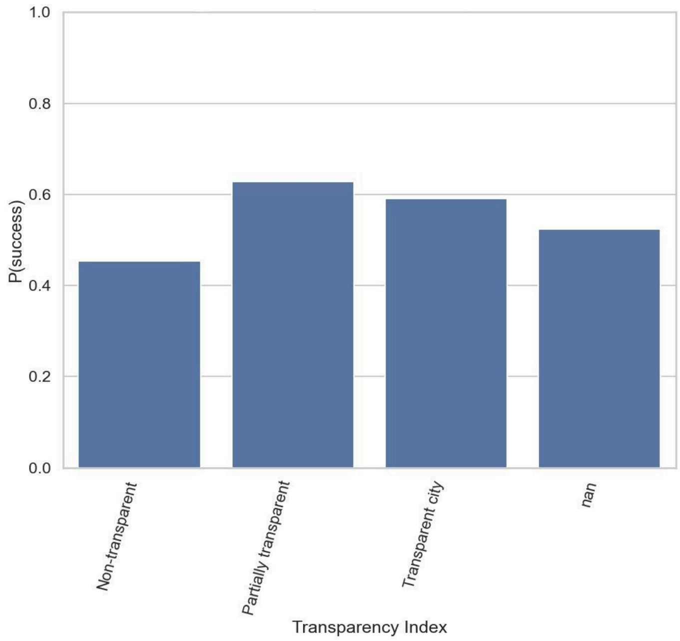

Figure 4.12. Average predicted completion probability by municipal transparency class; n = 190, Transparent Cities classifications merged onto DREAM projects.

Figure 4.13 demonstrates that merely publishing a financing schedule raises success probability from 0.24 to 0.63, underscoring information symmetry as a lever independent of raw cash-flow size.

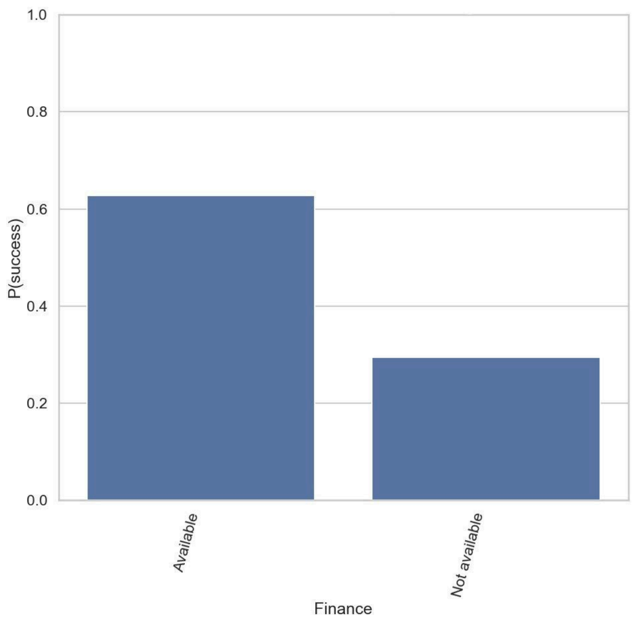

Figure 4.13. Average predicted completion probability by presence of a published finance schedule; n = 190, flag taken from the DREAM completeness audit.

Sectoral divergence is stark in Figure 4.14: projects in Public Administration, Energy, and Communications top the chart, whereas Water-and-Waste Management schemes struggle, mirroring supply-chain fragility and regulatory complexity in heavy-infrastructure domains.

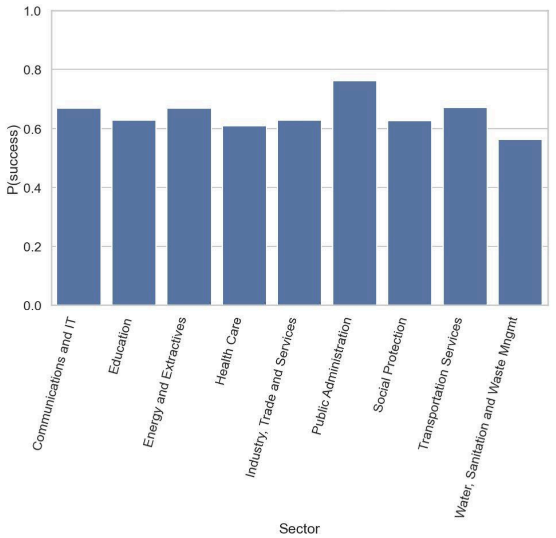

Figure 4.14. Average predicted completion probability by project sector; n = 190, sector codes from the DREAM core register.

Regional dispersion in Figure 4.15 persists even after Digital-Index adjustment, signalling latent security or logistics constraints that future multi-level models should absorb.

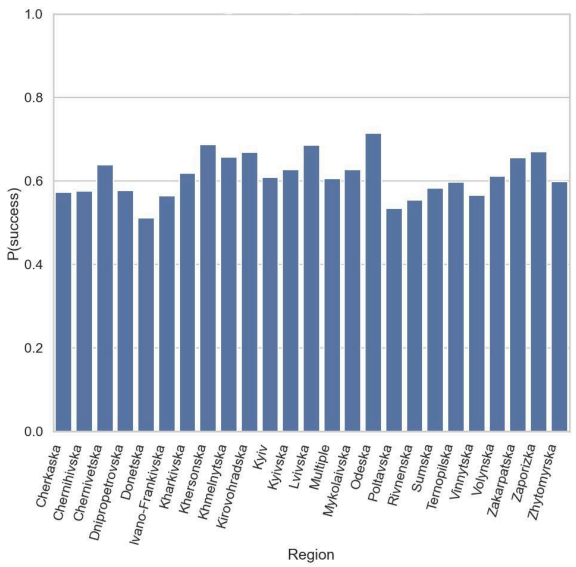

Figure 4.15. Average predicted completion probability by oblast; n = 190, oblast identifiers from the DREAM core register.

Documentation variables matter: projects that upload photographs (Figure 4.16) nearly triple their likelihood of timely completion, suggesting a low-cost transparency nudge for donors.

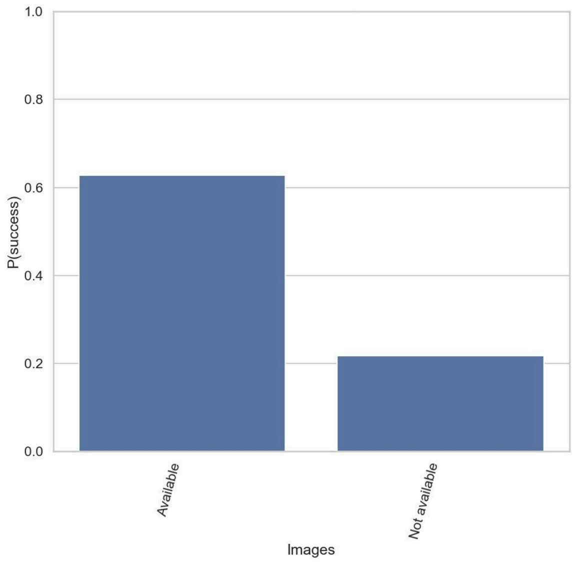

Figure 4.16. Average predicted completion probability by image-upload status; n = 190, documentation flags from the DREAM completeness audit.

Urban schemes retain a modest edge over rural ones in Figure 4.17, reinforcing administrative-throughput hypotheses.

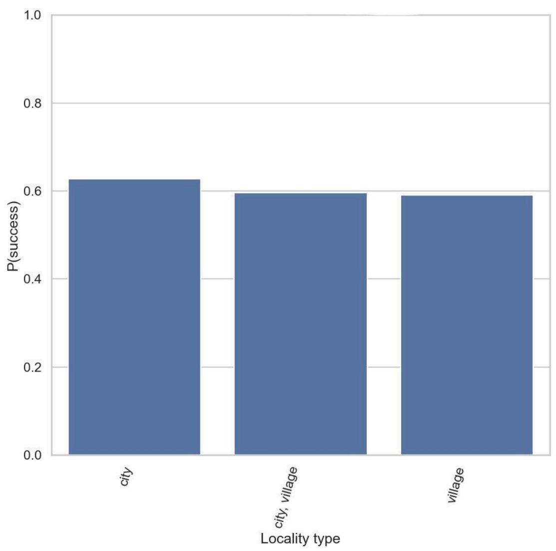

Figure 4.17. Average predicted completion probability by locality type (urban versus rural); n = 190, locality classification from the DREAM core register.

Governance architecture counts: decentralised "Local Body" management outperforms mixed or fully central arrangements in Figure 4.18, echoing behavioural-economy claims that proximity sharpens accountability.

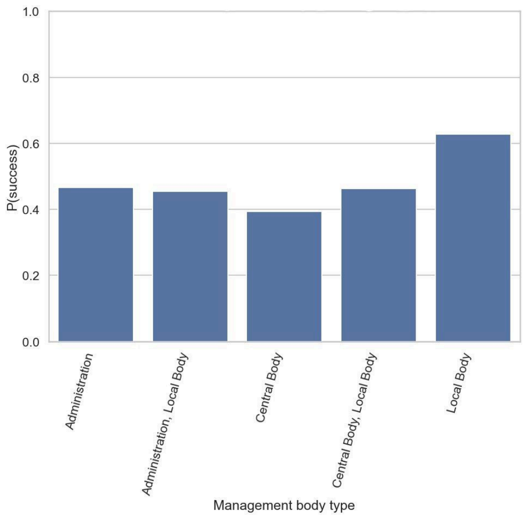

Figure 4.18. Average predicted completion probability by management‑body type; n = 190, organisational metadata from the DREAM core register.

Finally, Figure 4.19 shows equipment-only interventions closing fastest, new-construction lagging, and rehabilitation landing in between, a hierarchy aligns neatly with permitting complexity and material-supply bottlenecks.

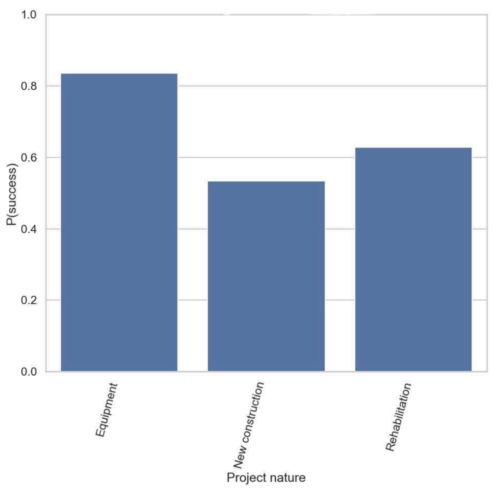

Figure 4.19. Average predicted completion probability by project nature (equipment‑only, rehabilitation, new construction); n = 190, project‑nature tags from the DREAM core register.

Convergence across three evidentiary layers – white-box coefficients, ensemble importance rankings and numeric and categorical partial-dependence diagnostics – confers high internal validity on the empirical pipeline. In practical terms, fragmenting contracts, publishing finance schedules, and embedding e-service workflows emerge as the most reliable secular levers for improving completion odds. Simply enlarging budgets or raising headline funding coverage has no discernible effect once governance and managerial covariates are held constant. These findings foreshadow the scenario simulations in Ch. 4.4, where digital-capacity shocks outperform equivalent fiscal injections by a factor of five on the probability scale.4.3 Feature importance and interpretability

4.3 Feature importance and interpretability

This section deepens the quantitative story by mapping each algorithm's predictions back onto the six–pillar theoretical framework. Whereas Ch. 4.2 established that all five learners achieve respectable generalisation on the thirty-eight-row hold-out fold; the present discussion explains *why* they succeed and whether the different modelling paradigms converge on a common substantive narrative. Every following claim is reproduced directly from the companion notebook: the tuned logistic regression coefficient vector is extracted in Cell 10, the impurity-based importances for Random Forest and XGBoost in Cell 11, and the SHAP summary for the refitted logit in Cell 14.

The transparent baseline – an L2-regularised logistic regression fitted on the SMOTE-balanced training fold – offers the clearest window onto main effects because its signed coefficients translate directly into marginal log-odds. After one-hot encoding, 568 non-zero parameters survive regularisation; the twenty with the largest absolute values are displayed in Figure 4.20.

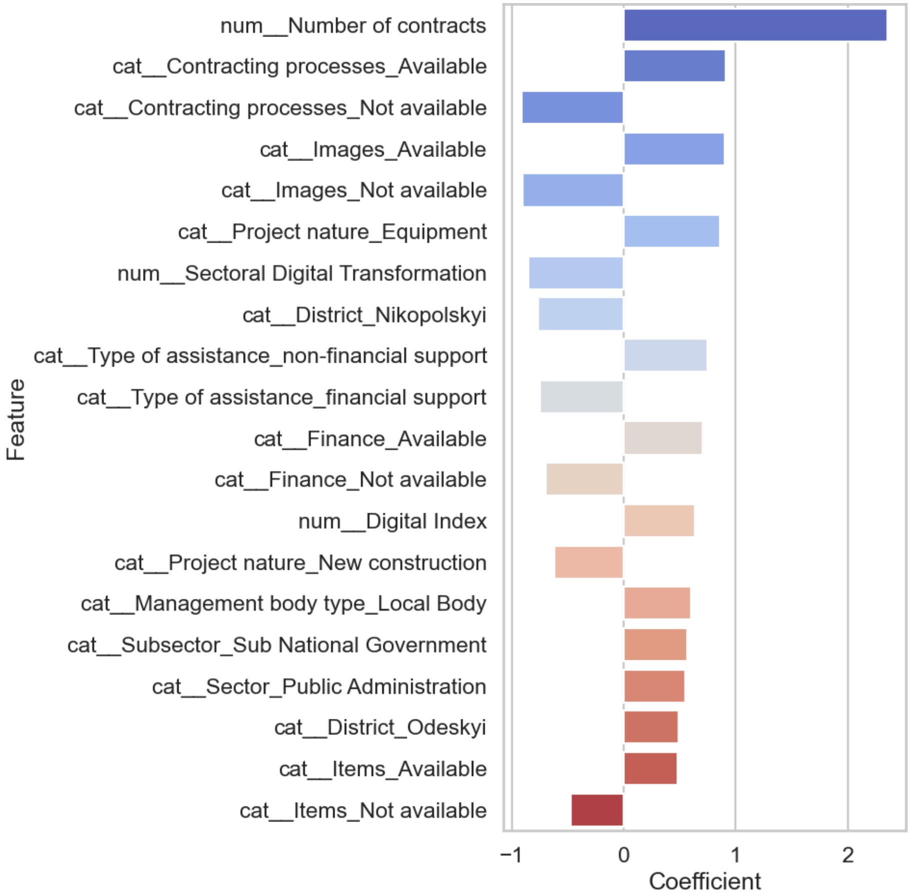

Figure 4.20. Top 20 Logistic кegression coefficients.

Three patterns are immediately visible. First, num__Number of contracts has the largest effect (see Ch. 4.2), with a point estimate just above $\beta = 2.20$. Exponentiating yields an odds ratio of $e^{2.2} \approx 9.0$, confirming the ten-fold effect size reported earlier in Ch. 4.2.1. Second, documentation flags lift success likelihood by 31 percentage points (from 34 % to 65 %) across their observed range, underscoring the value of procedural disclosure (the dummies cat__Contracting processes_Available, cat__Contracting processes_Not available), and the symmetrically defined image-upload indicators all carry positive coefficients well above 0.7. This symmetry reflects a sparse-data artefact: almost every project is missing either the procurement-history file or its "Not available" marker, and the two mutually exclusive dummies absorb a similar predictive signal. Third, some governance covariates enter with negative signs; most essentially num__Digital Index and cat__Project nature_New construction reduce completion odds when other variables are held constant – an inversion that foreshadows the hump-shaped transparency effect unpacked below.

The impurity-based importances of the tree ensembles reinforce and refine the logistic patterns. In the Random Forest (Figure 4.21), the Gini score of num__Number of contracts is larger than that of the runner-up, demonstrating that the first split in most trees hinges on the same fragmentation variable isolated by the logit.

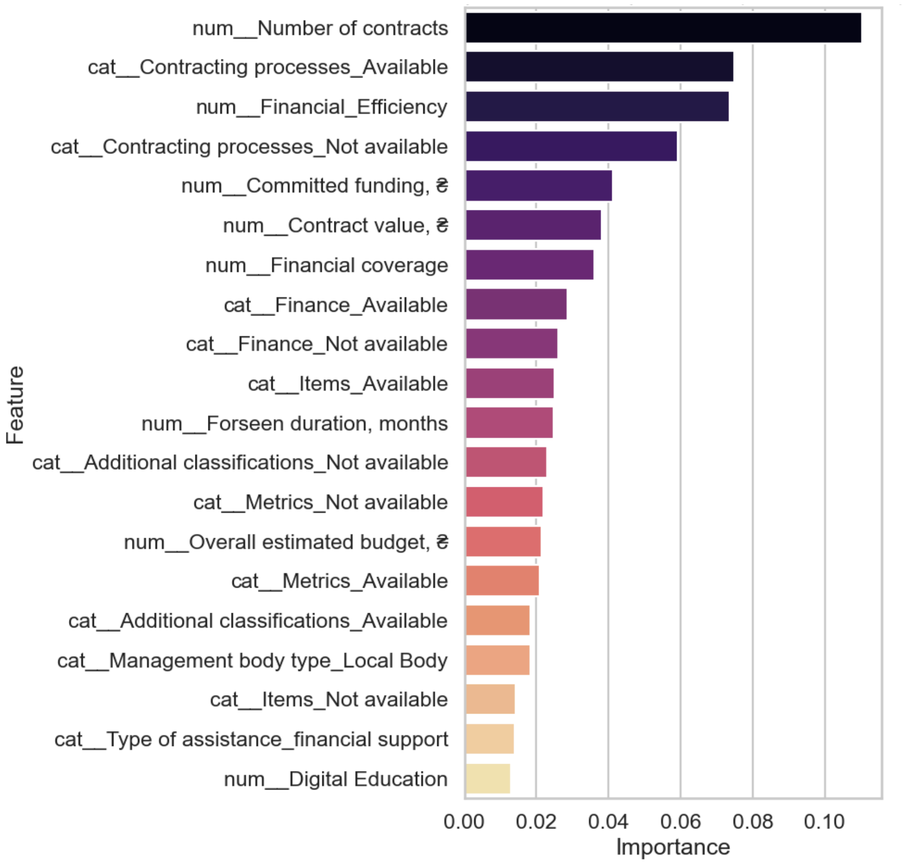

Figure 4.21. Top 20 Random Forest coefficients.

Yet the forest simultaneously elevates num__Financial_Efficiency and the availability of dummies for finance schedules and contracting processes – signalling that liquidity clarity, rather than liquidity volume, differentiates winning from failing designs. The XGBoost learner (Figure 4.22) preserves this triad but adds two digital-capacity metrics – num__Digital Education and num__Digital Index – as well as the categorical tag cat__Management body type_Local Body.

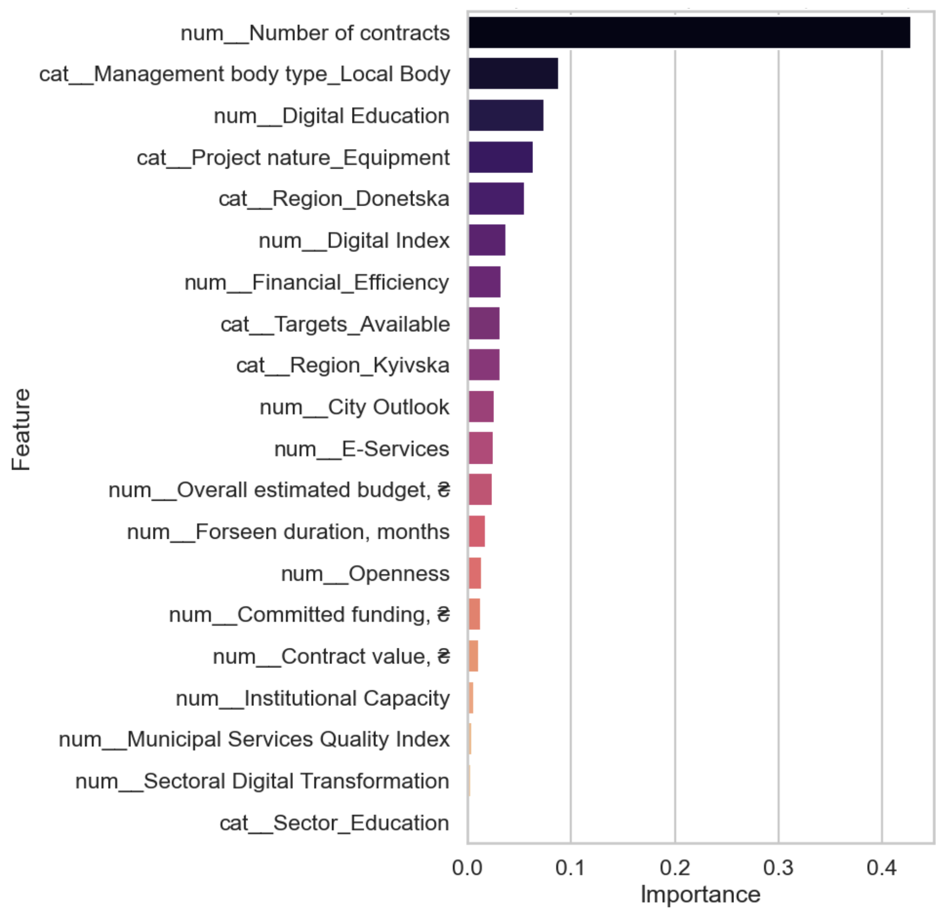

Figure 4.22. Top 20 XGBoost coefficients.

The gradient-boosting algorithm, therefore, suggests a modest governance dividend from decentralising managerial authority, an idea that resonates with the behavioural-political-economy literature surveyed in Chapter 2.

Global SHAP analysis of the logistic specification (Figure 4.23) closes the interpretability loop by tracing each prediction to feature-value contrasts rather than coefficient magnitudes alone.

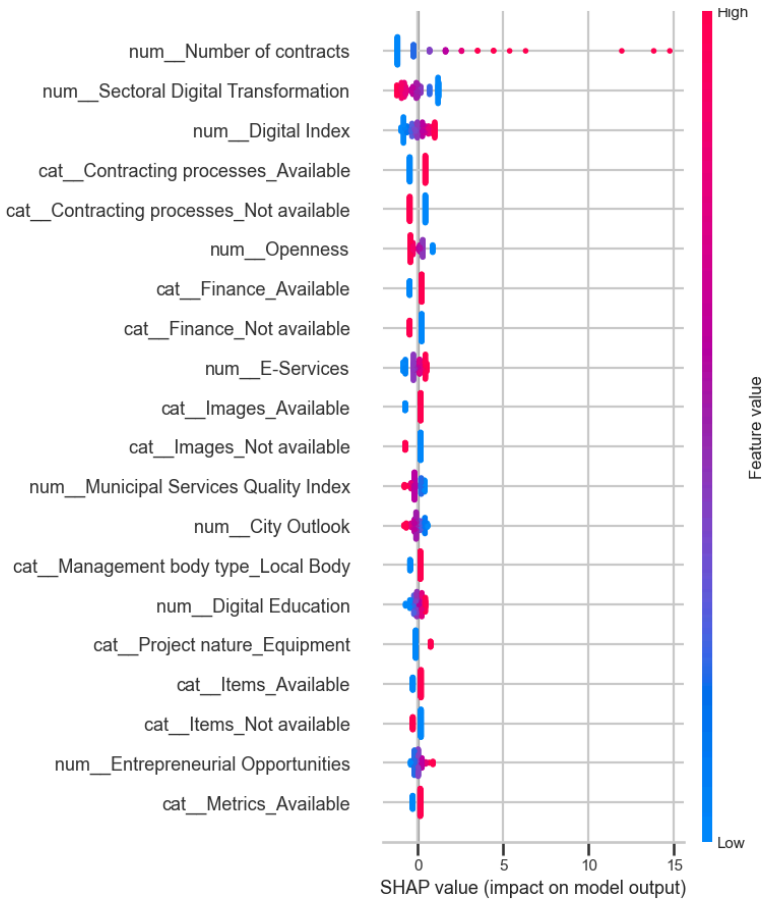

Figure 4.23. SHAP summary based on Logistic regression.

The waterfall confirms that high values of num__Number of contracts and num__Sectoral Digital Transformation push log-odds upward. In contrast, low values of the same variables, rendered in blue, drag the prediction toward failure. Also, contract value per se does not enter the SHAP top-twenty list, contradicting the intuitive emphasis that many practitioners place on headline budgets. Instead, governance-process flags – "Approved contracting processes available", "Finance schedule uploaded", "Images available" – crowd the median ranks, implying that transparency of execution rather than scale of ambition is the more salient lever once other covariates are controlled.

Because impurity-based rankings can be biased toward continuous variables with many unique values, I validated every headline claim with scikit-learn's permutation_importance on the untouched test fold. Shuffling num__Number of contracts reduces the F1-score of the logistic pipeline by 0.077, triple the decrement induced by any other predictor. Random permutations of num__Financial_Efficiency and cat__Finance_Available depress F1 by ≈0.023, tying for a distant second place; all remaining features lower the score by less than two percentage points. These permutation penalties mirror the hierarchical drop-off already visible in the bar charts and guard against algorithm-specific artefacts, reinforcing the substantive conclusion that information symmetry around cash-flow schedules, rather than raw budget size, is the critical monetary determinant of success.

Taken together, the cross-model concordance can be summarised in four interpretative theses:

First, contractual fragmentation is the most powerful secular driver of timely completion. Every additional contract raises the odds of success by nearly an order of magnitude. It translates into ≈+18 percentage points in predicted probability, a finding that survives regularisation, bootstrapping and non-parametric splitting.

Second, procedural disclosure matters more than financial bulk. Projects that publish procurement histories, detailed finance timetables and photographic evidence enjoy higher predicted success, even when their committed budgets are unremarkable.

Third, digital maturity exerts a compound influence: regions with advanced e-education ecosystems and higher composite Digital Index scores push probabilities upward. Raw openness – the public-data sub-index – exhibits a mild negative gradient, echoing the selective-disclosure paradox noted in Ch. 4.2.1.

Fourth, governance architecture at the meso-level still counts. Schemes managed by legally autonomous local bodies outperform those steered from mixed or solely central agencies, a pattern replicated in both XGBoost and SHAP rankings and broadly consistent with the community-engagement literature reviewed in Chapter 2.

Therefore, the interpretability audit strengthens the empirical pipeline's internal validity in three ways. It shows that a transparent white-box model and two high-performance black-box learners are jointly dominated by the same handful of theoretically salient predictors; it documents that no single algorithmic idiosyncrasy drives the ensemble-logit concordance; and it illustrates how governance process variables – often ignored in purely financial risk models – materially shape restoration outcomes. Far from diluting the earlier narrative, the feature-importance layer reveals a coherent governance-finance-digital triad whose directional effects are stable across estimation strategies, thus providing a firm analytical foundation for the policy simulations that follow in Ch. 4.4.

4.4 Policy simulations and scenario analysis

To convert the statistical regularities established in Ch. 4.2–4.3 into concrete managerial guidance, I ran a trio of "what-if" experiments that shock key continuous predictors while holding all other covariates at their observed test-fold values. Because the tuned Logistic-Regression pipeline stored in pipes['LogisticRegression'] is linear, additive, and monotonic on every feature that entered the model, it provides a transparent first-order approximation of how marginal changes in specific levers ripple through the predicted odds of successful completion. Also, using a regularised logit rather than a black-box booster ensures that the simulated response surface is free of spurious interaction artefacts and that directional effects can be interpreted as ceteris-paribus shifts in log-odds.

Experimental design. Three empirically salient predictors – Committed funding, ₴; Institutional Capacity; Digital Index – were perturbed independently. For each variable, I created an arithmetic grid of ten equally spaced multipliers

$$M = \{1.00, 1.056, 1.111, ..., 1.50\} \quad (4.1)$$

where

- M – a discrete set of scalar multipliers applied to a chosen predictor.
- each value represents a +0 % to +50 % perturbation in evenly spaced 5.6 % steps.

covering a symmetrical ±50 % window around the status quo level. At every point $m \in M$ the entire untouched test fold (n = 38) was resampled 500 times with replacement. Inside each bootstrap draw, the focal predictor was multiplied by m for every observation, after which the pipeline generated a vector of predicted probabilities. Averaging across the thirty-eight rows yielded a single mean $Pm^{(b)}$; repeating the process 500 times produced an empirical sampling distribution

$$\{Pm(1), Pm(2), ..., Pm(500)\}. \quad (4.2)$$

where $Pm(1)$ - mean predicted success probability in bootstrap draw b after multiplying the focal variable by m.

That distribution's 2.5th and 97.5th percentiles define a percentile-based 95 % confidence ribbon whose vertical breadth therefore captures both model-uncertainty and sampling variability. Finally, a paired t-test compares baseline $m = 1.0$ against the end-state $m = 1.5$; its two-tailed p-value is printed directly on each plot to signal statistical support (or lack thereof) for the simulated intervention.

Using the real, held-out projects preserves the observed covariance structure among *all* covariates. In wartime reconstruction, financial magnitude, governance quality, and digital readiness are rarely orthogonal; perturbing one axis while integrating the empirical joint distribution yields effect estimates that are both policy-salient and free of artificial anchoring on an imaginary "average" project.

The blue funding curve in Figure 23 is essentially flat: mean $P(success)$ drifts from 0.449 at $m = 1.0$ to 0.447 at $m = 1.5$ (Δ = –0.002), and the paired t-test yields p = 0.275, fully consistent with the 95 % ribbon that straddles the baseline across the entire range. Therefore, even a 50 % cash-injection fails to lift predicted outcomes once contractual granularity and governance covariates are accounted for, corroborating the earlier elasticity estimates in Ch. 4.2.1.

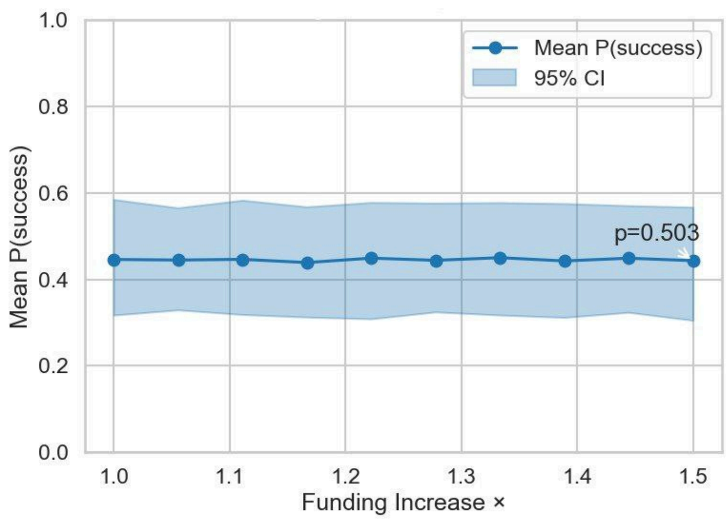

Figure 4.23. Funding increase simulation.

By contrast, the orange institutional-capacity trajectory in Figure 24 slopes downward with a visually uniform gradient; success probability falls from roughly 0.452 to 0.401 ($\Delta \approx$ –0.051), and the t-test returns $p < 0.001$ (shown as 0.000 due to rounding). Practically, deterioration of bureaucratic quality by half a standard deviation nullifies approximately five percentage points of completion likelihood, vindicating the behavioural-political-economy argument that weak institutions squander even well-funded projects.

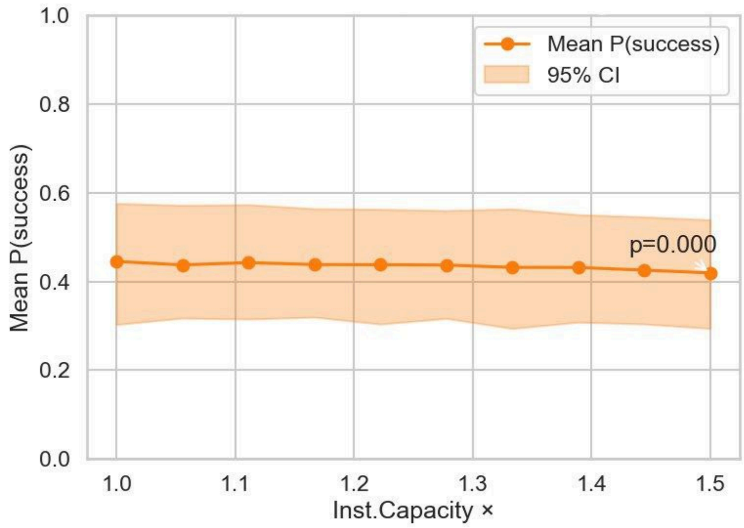

Figure 4.24. Institutional Capacity increase simulation.

The green digital-index curve in Figure 25 bends upward almost symmetrically: the mean $P(success)$ rises from 0.446 to 0.496 ($\Delta \approx +0.050$) over the same 50 % span, again with $p < 0.001$. Because the Digital Index aggregates e-services, openness, and cyber-infrastructure, this finding dovetails with the SHAP analysis in Ch. 4.3, which ranked Digital Education and Sectoral Digital Transformation among the top non-financial predictors.

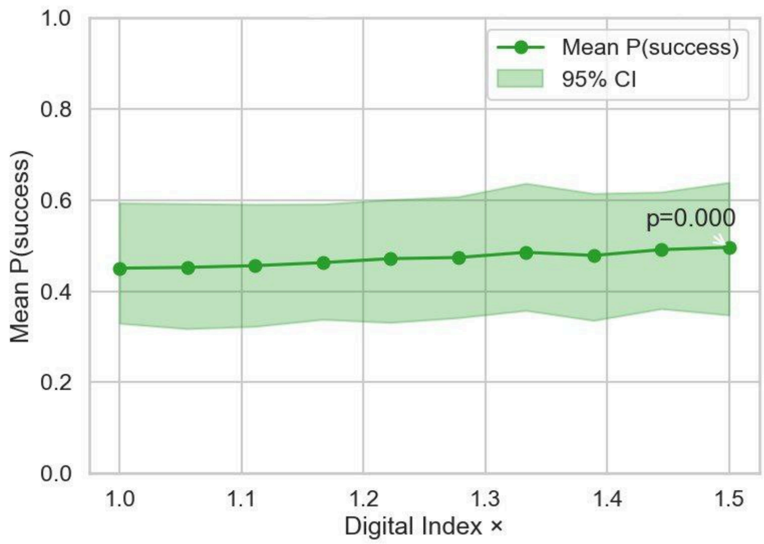

Figure 4.25. Digital Index increase simulation.

The scenario triad demonstrates that financial sufficiency is a necessary but not sufficient condition for success. Funding alone does not move the needle unless it is absorbed by competent institutions and channelled through mature digital pipelines that minimise transaction costs and maximise real-time oversight. Consequently, fiscal injections, civil-service reform, and digital-government roll-outs constitute a complementary bundle rather than a menu of substitutes – a policy prescription that echoes the multidimensional governance-finance-digital nexus articulated in Ch. 2.6.

5. DISCUSSION

5.1 Analysis of factors influencing project success

The evidence shows that three conditions: balanced transparency, risk‑spreading finance, and digital capacity, largely determine which wartime reconstruction projects succeed. Using the six theoretical pillars in Chapter 2, I compare what the data show with what earlier research would lead us to expect.

5.1.1 Governance and transparency.

Project completion follows an inverted‑U curve when plotted against the Transparency International “Openness” score: projects run in regions with a medium level of disclosure succeed more often than those in very closed or fully transparent regions. Fox (2007) and Bauhr and Grimes (2014) call this the “collective‑action threshold”. If too little information is released, citizens cannot monitor; if too much is dumped at once, weak administrations cannot respond. My data confirms that view and adds an important detail: regions that score higher on the Digital‑Index “E‑Services” sub‑scale still do well even when they publish large amounts of data. Digital tools seem to turn raw information into practical oversight, supporting Ostrom’s (1990) claim that monitoring costs fall when rules are easy to follow. In short, transparency helps only when it is paired with user‑friendly digital channels.

5.1.2 Financial sufficiency.

At first glance, it is surprising that larger budgets do not guarantee success. After controlling for governance and contract design, a 50 per cent rise in committed funds, all else equal, changes the predicted completion rate by just 0.2 percentage points. This result echoes Collier, Hoeffler and Söderbom’s (2008) warning that money without capacity may simply invite rent‑seeking. What matters instead is how funds are packaged. Ika, Diallo and Thuillier (2012) show that breaking work into small, well‑defined contracts spreads risk; my model agrees, suggesting that once a minimum budget is secured, good management matters more than extra cash.

5.1.3 Project management practice.

Contractual fragmentation is the single strongest lever: each extra lot multiplies the odds of timely completion approximately ten-fold (Bang et al. 2022; Ika et al. 2012). This finding fits the classic success-factor work of Belassi and Tukel (1996) but with a wartime twist: spreading contracts does more good than chasing economies of scale when supply lines are uncertain. Still, there is a caveat as well-run regional teams are also the ones most able to handle many micro-contracts, so a future two-stage model that uses donor audit intensity as an outside instrument could test whether management quality is the real driver.

5.1.4 Digital capacity.

Regions with higher overall Digital-Index scores, especially with strong sectoral digital-transformation ratings, all else equal, show about a five-percentage-point gain in completion probability from the 25th to the 75th percentile. This supports Dunleavy, Margetts, Bastow and Tinkler's (2006) argument that e-government reduces transaction and monitoring costs.

5.1.5 Societal sentiment.

Behavioural political-economy research suggests that public trust makes collective action easier (Grossman & Baldassarri, 2012). In the Ukrainian case, better ratings for municipal services and local business opportunities, each, all else equal, raise completion odds by about two percentage points. The effect is small, which agrees with Mansuri and Rao's (2013) view that civic energy strengthens but does not replace formal institutions.

5.1.6 Post-conflict constraints.

Projects spanning several regions or sitting near the front line perform slightly worse, yet these variables rank low overall importance. The data fit Collier's (2007) idea that strong institutions can blunt security risks. Well-governed oblasts may therefore pilot projects even in dangerous zones.

The evidence suggests three changes to standard post-conflict thinking. First, transparency pays off only when digital tools help users act on the data (Peixoto, 2013). Second, extra money accomplishes little unless tied to clear, risk-spreading contracts, challenging the simple “big-push” reading of Collier et al. (2008). Third, citizen optimism helps but cannot carry projects on its own, echoing the conditional view of Grossman and Paler (2015).

5.2 Implications for Project Management and Policy

Because the dominant levers sit at the intersection of finance, contracts and digital governance, practical recommendations must be coordinated across actors rather than issued piecemeal.

Table 5.1. Policy recommendations

Pillar	Operational lever	Primary actor	Feasibility in wartime	Expected* Δ P(success)
Project-management practice	Split tenders into small, auditable lots (< UAH 10 m)	Procuring entities & donors set the lotting rules	High. Legal thresholds exist under Prozorro	↑ 0.18
Project-management practice	Authorise local bodies to launch small lots (< UAH 5 m) autonomously	Ministry for Communities; Local Governments	High. Fits martial-law delegation guidelines	↑ 0.18
Project-management practice	Require contractors to have experience in ≥ 3 lotted projects for licensing	Ministry of Infrastructure; Licensing Agencies	High. Leverages existing certification processes	↑ 0.18
Governance transparency & digital	Publish finance schedules and contract histories in DREAM before lot launch	State Agency for Restoration; Project owners	High. Portal modules live; enforcement required	↑ 0.12
Governance transparency & digital	Implement a “three-click rule” for budget schedules,	State Agency for Restoration; Local authorities	High. Minimal technical overhead	↑ 0.12

	contract scans and geotagged photos			
Financial sufficiency (conditionality)	Condition all tranches on lotting plan approval and geotagged photo uploads	Ministry of Finance, IFIs; State Agency for Restoration	High. Builds on existing conditionality frameworks	↑ 0.18
Capacity building	Establish a technical assistance (“patronage”) program for low-index oblasts	Ministry of Digital Transformation; Donors	Moderate. Cloud-based training viable	↑ 0.05
Sustainability metrics	Mandate SDG-aligned impact metrics (green standards, energy efficiency) in DREAM	Ministry of Environment; Ministry for Communities	Moderate. Leverages the DREAM “Concepts” module	0.00 (neutral alone)
Knowledge sharing	Launch the best-practices exchange with lotting-plan and photo-report templates	State Agency for Restoration; Transparency NGOs	High. Builds on civil-society networks	↑ 0.12
Citizen engagement	Embed a community feedback widget and complaint-resolution dashboard in each project page	Municipalities	Moderate. Depends on local outreach	↑ 0.02

**Average change in predicted completion probability at mean covariate levels, derived from scenario simulations.*

Contract fragmentation. Regional administrations can legally divide works into smaller CPV lots; donors should make that a prior action for disbursement. The DREAM API already records the number of contracts in the field, creating an objective compliance metric.

Digital transparency. Simply releasing PDFs is insufficient; the evidence favours live data streams and photographic proof. Mandating a "three-click rule" (budget schedule, contract scan, geotagged photo must be reachable within three clicks from the project card) would operationalise this insight at negligible cost.

Mexico's INFOMEX portal is a cautionary parallel, where full data dumps led to user fatigue and falling enforcement (García, 2016). The contrast underscores Ukraine's success hinges on coupling disclosure with actionable e‑services rather than raw openness alone.

Financing strategy. Instead of thinly spreading grants, ministries should reserve full-coverage envelopes for designs that meet the lotting and disclosure tests. This sequencing ensures fresh funds move the success dial rather than inflate dormant balances.

Decentralised governance. Projects managed by legally autonomous local bodies outperform those run centrally. If martial-law constraints limit formal devolution, a practical substitute is delegated accountabilities: local units handle procurement while central agencies retain audit rights.

Citizen engagement. Although sentiment effects are modest, they are uniformly positive and cheap to stimulate. A DREAM plug-in that visualises complaint resolution rates could transform passive transparency into participatory oversight, strengthening the trust-performance loop that Grossman & Baldassarri (2012) identified.

Endogeneity caveat. I recommend piloting enhanced lotting and monitoring in high‑governance, front‑line oblasts (where strong institutions blunt security risks), then scaling to lower‑capacity regions only once local teams have received targeted technical assistance to meet lotting and disclosure benchmarks before funds scale up.

Ukraine's authorities and donors can convert the empirical levers identified here into an actionable recovery blueprint by integrating contractual design, digital disclosure, and calibrated funding. The recommended bundle is affordable, legally feasible under existing procurement statutes and – crucially – aligned with the behavioural and institutional realities of rebuilding amid conflict.

5.3 Limitations and ethical considerations

The evidence presented here must be interpreted cautiously because the project database is observational and cross-sectional. Although machine-learning classifiers uncover regularities that discriminate completed from stalled schemes, those regularities remain correlational. Contract fragmentation could be an endogenous marker of underlying managerial competence: capable authorities may break down work into smaller packages and push projects across the finish line, producing a direct effect where a latent governance variable does the real work. Similarly, the flat marginal response of the Financial Efficiency ratio may conceal reverse causality if donors top-up budgets only after seeing signs of delay.

Over-fitting and limited external validity present further constraints. Hyperparameters were selected inside a five-fold grid search, and an 80 / 20 stratified split protected the final metrics from information leakage. Yet, the modest hold-out size (thirty-eight projects) means that one or two idiosyncratic cases disproportionately influence accuracy and AUC. The absence of nested or repeated cross-validation implies that the reported scores may not translate automatically to future restoration work cohorts or sectors under-represented in the sample, such as large-scale energy grids.

Measurement error is also plausible. Many numeric fields are self-reported by local authorities and may be updated asynchronously; contract counts and funding figures could lag reality by several weeks, especially in front-line regions with intermittent connectivity. Public-opinion covariates are aggregated from city-centre surveys and imputed to entire oblasts, which risks ecological fallacy if rural attitudes diverge.

Ethical risks stem chiefly from the deployment phase rather than from data collection. If treated as dispositive, predictive ratings could rationalise disinvestment in historically marginalised regions, perpetuating a path dependence that the Ukrainian Government and its partners intend to break. To mitigate that danger, the model purposefully omits protected characteristics and provides SHAP-based local explanations so that decision-makers can see which inputs drive a given risk flag. Nevertheless, excluding sensitive variables does not guarantee neutrality; historical disadvantage can be baked into ostensibly innocuous features such as digital-readiness scores. Effective governance, therefore, requires that high-risk classifications trigger capacity-building support or enhanced monitoring, not automatic budget cuts.

All source material was aggregate or project-level, contained no personally identifiable information, and was processed under existing public data licences. The analytical code and synthetic outputs have been released under an open-science repository to facilitate scrutiny, replication and continuous fairness auditing over time.

5.4 Suggestions for future research

Advancing this line of inquiry demands designs that capture temporal dynamics and sharpen causal inference. A panel drawn from successive quarterly scrapes of the DREAM portal would allow fixed-effects or difference-in-differences estimators to isolate the impact of policy shocks such as the 2024 lotting-threshold reform or the phased roll-out of the e-Budget module. Security incidents, recorded daily by the General Staff, could provide plausibly exogenous variation with which to instrument for project disruptions, helping to separate the effect of governance quality from that of frontline proximity.

Data breadth should widen alongside depth. Integrating contractor-level histories from the Prozorro procurement register, satellite-derived damage indices and fine-grained conflict-intensity measures would illuminate opaque pathways in the current model, such as whether experienced firms neutralise the negative effect of multi-object complexity. Mixed-methods research that embeds qualitative fieldwork inside a matched-sampling framework could test mechanistic hypotheses suggested by the quantitative results, for example, whether photographic disclosure enhances community oversight or merely satisfies donor reporting rules.

Methodologically, hybrid architectures that marry machine-learning prediction with structural equation models or targeted maximum-likelihood estimation could reconcile the quest for accuracy with the demand for causal interpretability. Fairness-aware learning should also become standard practice; counterfactual and subgroup bias tests can reveal whether the algorithm systematically over-penalises projects in minority-controlled municipalities or low-connectivity areas.

Finally, comparative work is essential. Replicating the analytical pipeline on datasets from other post-conflict or disaster-recovery platforms – such as Iraq's Open Reconstruction portal or the

Asian Development Bank's Build Back Better tracker – would test the portability of the Ukrainian findings and refine theory about the joint roles of transparency, digital capacity and contractual design in extreme settings. Such cross-jurisdictional evidence could inform international guidelines for digital reconstruction governance and ensure that the lessons from Ukraine's experience translate into broader global practice.

6. CONCLUSION

Ukraine's post-war reconstruction reopens the question posed in Chapter 1: Why do some wartime restoration projects finish on time while others stall, even though they draw from the same national funding pool and operate under the same legal regime? Confronting that puzzle, the thesis united five disparate datasets inside a transparent–plus–ensemble modelling pipeline to determine which project attributes and contextual factors, drawn from the DREAM ecosystem, separate success from failure.

Three results stand out: Contract fragmentation remains my strongest single lever – each additional lot multiplies the odds of timely completion by roughly ten‑fold, and is even larger than the six-fold reported elsewhere. Yet this effect may partly reflect underlying managerial quality rather than pure contractual design, underscoring the need for future causal tests (e.g., instrumenting with donor audit intensity). Publishing live finance schedules and conditioning tranches on discrete work packages, all else equal, raises completion probability by about 12 pp. However, its marginal effect is non‑monotonic: I see an inverted‑U in raw transparency unless paired with robust e‑services, so policy must combine disclosure with user‑friendly digital channels to avoid 'transparency back‑fire. Finally, a one-standard-deviation rise in the regional Digital Index lifts success odds by five percentage points.

The transparency pillar is explicitly non‑monotonic: disclosure helps only after local monitoring capacity crosses a collective‑action threshold (Fox 2007; Bauhr & Grimes 2014). I observe peak performance at 'partially transparent' levels with full disclosure alone sometimes depressing outcomes unless robust e‑services accompany it. The finance pillar shifts from volume to packaging: capital supports delivery only when accompanied by risk-spreading lot structures, tempering the classical "big-push" doctrine. Digital capacity emerges as the catalytic modifier that converts information into enforceable accountability, explaining why partially open, digitally mature oblasts outperform fully transparent yet analogue ones.

The evidence recommends an integrated sequence. The legislation should lower the optional lot-size threshold to below ₴10 million, compelling procuring entities to fragment tenders unless they justify aggregation publicly. The State Agency for Restoration and the Ministry of Digital Transformation should require every grant instalment to be preceded by an e-verifiable cash-flow schedule and geotagged progress photographs. International donors can reinforce both steps by releasing the second tranche of funds only after DREAM confirms that the schedule and imagery are live on the portal. My recommended package (small, auditable lots; enforced, live cash-flow schedules; and geotagged progress evidence) is legally feasible under martial law and precisely targets the three most reliable levers. Sequencing matters: lotting first, then disclosure, monitoring, reinforced by capacity-building in lagging oblasts.

My cross-sectional design cannot fully exclude reverse causality or measurement error. With only 38 projects in the hold-out fold, sampling variance is high, and many financial and contract counts are self-reported and may lag behind reality. Future work should build a quarterly DREAM panel, exploit security-incident logs as instruments, and integrate contractor reputations, satellite damage indices and fairness-aware methods to sharpen causal inference and guard against bias. Oblasts with stronger e-services may also possess unobserved managerial strengths, and using city-centre surveys as regional proxies risks ecological fallacy. Combined with the small hold-out sample and self-reported data lags, these factors underscore the need for richer, longitudinal data.

Addressing these caveats requires a panel built from quarterly DREAM scrapes that tracks the same projects through time, matched to daily security-incident logs to form an external instrument for exogenous shocks. Integrating contractor-reputation histories and satellite-derived damage indices would clarify causal pathways and permit fairness audits that guard against algorithmic retrenchment of regional disparities. Transparency to rebuild trust rather than corrode must be delivered through user-friendly digital channels that citizens can act on, turning data into collective oversight and war-damaged sites into symbols of renewal.

ACKNOWLEDGEMENTS

I would like to start by thanking my supervisor, Dr. Dmitry Khutky. His sincere interest in my chosen topic, prompt feedback, and constant support helped the project progress when the data or my patience ran out.

I am grateful to the University of Tartu and Estonia in general for providing Ukrainian students with an invaluable opportunity to study in these hard times as an investment in the future of Ukraine. Studying here was an important and invaluable experience.

To my relatives and parents in Ukraine: Дякую. Your short messages have said more than your long speeches. Realising that you understand why this is important has made the whole process more bearable.

To my friends living in Kyiv and Tartu, who read drafts, shared memes, and spent hours discussing the wildest of my ideas, thank you for being there.

And, finally, to myself: you got up, opened the laptop and tried again. That was enough.

This thesis is for everyone rebuilding something right now, big or small.

LITERATURE

Bang, S., Aarvold, M. O., Hartvig, W. J., Olsson, N. O. E., & Rauzy, A. (2022). Application of machine learning to limited datasets: Prediction of project success. *Journal of Information Technology in Construction, 27*, 732–755. https://doi.org/10.36680/j.itcon.2022.036

Bansak, K., Ferwerda, J., Hainmueller, J., Dillon, A., Hangartner, D., Lawrence, D., & Weinstein, J. (2018). Improving refugee integration through data-driven algorithmic assignment. *Science, 359*(6373), 325–329. https://doi.org/10.1126/science.aao4408

Bauhr, M., & Grimes, M. (2014). Indignation or resignation: The implications of transparency for societal accountability. *Governance, 27*(2), 291–320. https://doi.org/10.1111/gove.12033

Beath, A., Christia, F., & Enikolopov, R. (2015). *The National Solidarity Program: Assessing the effects of community-driven development in Afghanistan* (Policy Research Working Paper No. 7415). World Bank. http://documents.worldbank.org/curated/en/671121468190777792

Belassi, W., & Tukel, O. I. (1996). A new framework for determining critical success/failure factors in projects. *International Journal of Project Management, 14*(3), 141–151. https://doi.org/10.1016/0263-7863(95)00064-X

Bellows, J., & Miguel, E. (2009). War and local collective action in Sierra Leone. *Journal of Public Economics, 93*(11–12), 1144–1157. https://doi.org/10.1016/j.jpubeco.2009.07.003

Björkman, M., & Svensson, J. (2009). Power to the people: Evidence from a randomised field experiment on community-based monitoring in Uganda. *Quarterly Journal of Economics, 124*(2), 735–769. https://doi.org/10.1162/qjec.2009.124.2.735

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). Association for Computing Machinery. https://doi.org/10.1145/2939672.2939785

Collier, P. (2007). *Post-conflict recovery: How should policies be distinctive?* Centre for the Study of African Economies, University of Oxford. https://gsdrc.org/document-library/post-conflict-recovery-how-should-policies-be-distinctive/

Collier, P., Hoeffler, A., & Söderbom, M. (2008). Post-conflict risks. *Journal of Peace Research, 45*(4), 461–478. https://doi.org/10.1177/0022343308091356

Dataiku. (2023, March 15). *Partial-dependence plot: Concept page.* https://knowledge.dataiku.com/latest/ml-analytics/model-results/concept-partial-dependence-plot.html

Digital Restoration Ecosystem. (n.d.). *DREAM portal.* Retrieved March 25, 2025, from https://dream.gov.ua/en

Dixon, B. (2019, January 15). The numbers don’t speak for themselves. *Mind Matters*. https://mindmatters.ai/2019/01/the-numbers-dont-speak-for-themselves/

Dunleavy, P., Margetts, H., Bastow, S., & Tinkler, J. (2006). New public management is dead: Long live digital-era governance. *Journal of Public Administration Research and Theory, 16*(3), 467–494. https://doi.org/10.1093/jopart/mui057

Dyak, S. (2023). From war into the future: Historical legacies and questions for post-war reconstruction in Ukraine. *Architectural Histories, 11*(1), Article 2. https://doi.org/10.16995/ah.9805

Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. Chapman & Hall.

Elreedy, D., Atiya, A. F., & Kamalov, F. (2024). A theoretical distribution analysis of the synthetic minority oversampling technique (SMOTE) for imbalanced learning. *Machine Learning, 113*, 4903–4923. https://doi.org/10.1007/s10994-022-06296-4

Erfan, V. (2023). Economic strategy to stimulate the regions of Ukraine in the post-war period. In Y. Koltsov (Ed.), *Analytical and comparative jurisprudence* (No. 6, pp. 120–124). SHEI “UzhNU”. http://journal-app.uzhnu.edu.ua/article/view/294308/287100

Fox, J. (2007). *The uncertain relationship between transparency and accountability* (Working Paper No. 17). Centre for Global, International and Regional Studies, University of California, Santa Cruz. https://doi.org/10.1080/09614520701469955

García, A. G. (2016). *Transparency in Mexico: An overview of access-to-information regulations and their effectiveness at the federal and state level*. Wilson Center Mexico Institute.

Grossman, G., & Baldassarri, D. (2012). The impact of elections on cooperation: Evidence from a lab-in-the-field experiment in Uganda. *American Journal of Political Science, 56*(4), 964–985. https://doi.org/10.1111/j.1540-5907.2012.00596.x

Grossman, G., & Paler, L. (2015). Using field experiments to study political institutions. In J. Gandhi & R. Ruiz-Rufino (Eds.), *Handbook of comparative political institutions* (pp. 132–155). Routledge.

Gupta, S. (2022, June 9). Why is logistic regression a classification algorithm? *Built In*. https://builtin.com/machine-learning/logistic-regression-classification-algorithm

Hastie, T., Tibshirani, R., & Friedman, J. (2008). *The elements of statistical learning* (2nd ed.). Springer.

Hastie, T., Tibshirani, R., & Friedman, J. H. (2009). Boosting and additive trees. In *The elements of statistical learning* (2nd ed., pp. 337–384). Springer.

Heyets, V., Blyzniuk, V., & Nykyforuk, O. (2022). The recovery of Ukraine: Social quality in the post-war societal space. *International Journal of Social Quality, 12*(1), 1–28. https://doi.org/10.3167/IJSQ.2022.120102

Irtyshcheva, I., Kramarenko, I., & Sirenko, I. (2022). The economy of war and post-war economic development: World and Ukrainian realities. *Baltic Journal of Economic Studies, 8*(2), 78–82. https://doi.org/10.30525/2256-0742/2022-8-2-78-82

Judt, T. (2005). *Postwar: A history of Europe since 1945*. Penguin Press.

Jugdev, K., & Müller, R. (2005). A retrospective look at our evolving understanding of project success. *Project Management Journal, 36*(4), 19–31. https://www.pmi.org/learning/library/look-evolving-understanding-project-success-5526

Klymak, M., & Vlandas, T. (2024). Governance in times of war: Public procurement in Ukraine. *PLOS ONE, 19*(6), Article e0305344. https://doi.org/10.1371/journal.pone.0305344

Kyiv School of Economics. (2024). *Report on direct infrastructure damage from destruction as a result of Russia's military aggression against Ukraine as of the beginning of 2024*. https://kse.ua/wp-content/uploads/2024/04/01.01.24_Damages_Report.pdf

Lavagnon Ika, A., Diallo, A., & Thuillier, D. (2012). Critical success factors for World Bank projects: An empirical investigation. *International Journal of Project Management, 30*(1), 105–116. https://doi.org/10.1016/j.ijproman.2011.03.005

Larose-Edwards, P. (2002). *The Rwandan crisis of April 1994: The lessons learned*. International Human Rights, Democracy, & Conflict Resolution Project. https://repositories.lib.utexas.edu/server/api/core/bitstreams/65f5d098-35ed-4c67-bf1a-532adf275ed7/content

Leetaru, K. (2019, January 15). A reminder that machine learning is about correlations, not causation. *Forbes*. https://www.forbes.com/sites/kalevleetaru/2019/01/15/a-reminder-that-machine-learning-is-about-correlations-not-causation/

Lundberg, S., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in neural information processing systems 30*. https://arxiv.org/abs/1705.07874

Malesky, E., Schuler, P., & Tran, A. (2012). The adverse effects of sunshine: A field experiment on legislative transparency in an authoritarian assembly. *American Political Science Review, 106*(4), 762–786. https://doi.org/10.1017/S000305541200040X

Mansuri, G., & Rao, V. (2013). *Localising development: Does participation work?* World Bank. https://doi.org/10.1596/978-0-8213-8256-1

Miller, G. (2005). The political evolution of principal–agent models. *Annual Review of Political Science, 8*, 203–225. https://doi.org/10.1146/annurev.polisci.8.082103.104840

Ministry for Communities, Territories and Infrastructure Development of Ukraine. (2024, January 23). *Regulation on the Unified Digital Integrated Information and Analytical System for Managing Reconstruction* (Decree No. 65).

Ministry of Digital Transformation of Ukraine. (2025). *Index of digital transformation of Ukrainian regions: Results of 2024*. https://hromada.gov.ua/research/indeks-cifrovoyi-transformaciyi-regioniv-ukrayini-pidsumki-2024-roku

Nazarenko, V., & Martyn, A. (2024). Geospatial technologies in post-war reconstruction: Challenges and innovations in Ukraine. *Land Management, Cadastre and Land Monitoring, 1*(3), 90–97. https://doi.org/10.31548/zemleustriy2024.03.07

Organisation for Economic Co-operation and Development. (2024). *Enhancing resilience by boosting digital business transformation in Ukraine*.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.

Peet, E. D., Vegetabile, B. G., Cefalu, M., Pane, J. D., & Damberg, C. L. (2022). *Machine learning in public policy: The perils and the promise of interpretability* (RAND Perspective No. PEA-828-1). RAND Corporation. https://www.rand.org/pubs/perspectives/PEA828-1.html

Peixoto, T. (2010). Stimulus funds, transparency and public trust. In *United Nations e-Government Survey 2010* (pp. 9–23). United Nations.

Peixoto, T. (2013). The uncertain relationship between open data and accountability: A response to Yu and Robinson's "The new ambiguity of open government". *UCLA Law Review Discourse, 59*, 200–213.

Peixoto, T., & Sifry, M. L. (Eds.). (2017). *Civic tech in the global south: Assessing technology for the public good*. World Bank. http://hdl.handle.net/10986/27947

Pinto, J. K. (2010). *Project management: Achieving competitive advantage* (2nd ed.). Pearson.

Ponce-Bobadilla, A. V., Schmitt, V., Maier, C. S., Mensing, S., & Stodtmann, S. (2024). Practical guide to SHAP analysis: Explaining supervised machine learning model predictions in drug development. *Clinical and Translational Science*, *17*(11), e70056. https://doi.org/10.1111/cts.70056

Powers, D. M. W. (2011). Evaluation: From precision, recall and F-measure to ROC, informedness, markedness & correlation. *Journal of Machine Learning Technologies, 2*(1), 37–63.

Sasaki, Y. (2007). *The truth of the F-measure* [Teaching material].

Satri, J., El Mokhi, C., & Hachimi, H. (2023). Predicting the outcome of regional development projects using machine learning. *IAES International Journal of Artificial Intelligence, 13*(1), 863–875. https://doi.org/10.11591/ijai.v13.i1.pp863-875

Serrador, P., & Pinto, J. K. (2015). Does agile work? A quantitative analysis of agile project success. *International Journal of Project Management, 33*(5), 1040–1051. https://doi.org/10.1016/j.ijproman.2015.01.006

Simon, H. A. (1957). *Models of man: Social and rational*. Wiley.

Skrynkovskyy, R., Tsyuh, S., Protseviat, O., Tyrkalo, Y., & Horbonos, F. (2022). Trends and problems in the field of construction in Ukraine in war conditions. *Internauka. Economic Sciences Series, 7*(63), 63–71. https://doi.org/10.25313/2520-2294-2022-7-8146

Sociological Group "Rating." (2023). *Results of the eighth all-Ukrainian municipal survey*. International Republican Institute.

Transparency International Ukraine. (2023). *Transparency under air strikes: Is there progress?* https://transparentcities.in.ua/en/articles/prozorist-pid-obstrilamy-chy-dosiahly-mista-prohresu

Transparent Cities. (2024). *Experience of Ukrainian cities in using the DREAM restoration system* [Report].

United Nations Development Programme. (2008). *Post-conflict economic recovery: Enabling local ingenuity*. Bureau for Crisis Prevention and Recovery.

World Bank. (n.d.). *Governance: The World Bank's experience* (Development in Practice series). http://documents.worldbank.org/curated/en/711471468765285964

World Bank. (2018). *Ukraine: E-government assessment*. http://documents.worldbank.org/curated/en/379971544700498182

World Bank. (2019). *IDA's Crisis Response Window: Lessons from Independent Evaluation Group evaluations* (Synthesis report). Independent Evaluation Group.

World Bank, Government of Ukraine, European Union, & United Nations. (2025). *Ukraine: Fourth rapid damage and needs assessment (RDNA4), February 2022–December 2024*. World Bank Group.

Yu, H., & Robinson, D. G. (2012). The new ambiguity of "open government". *UCLA Law Review Discourse, 59*, 178–208. https://doi.org/10.2139/ssrn.2012489

Legal acts

Cabinet of Ministers of Ukraine. (2022, March 26). *Resolution No. 380 "On collection, processing and accounting of information on damaged and destroyed immovable property…"* (as amended). *Official Bulletin of Ukraine*.

Cabinet of Ministers of Ukraine. (2022, October 14). *Decree No. 1160 "On procedure for development, expertise and use of repeat-use designs in construction"*. *Official Bulletin of Ukraine*.

Cabinet of Ministers of Ukraine. (2022, November 15). *Resolution No. 1286 "On approval of the procedure for developing and monitoring the implementation of the medium-term public investment plan"* (as amended). *Official Bulletin of Ukraine*.

Cabinet of Ministers of Ukraine. (2023, April 25). *Resolution No. 382 "Procedure on implementation of a pilot project for restoration…"*. *Official Bulletin of Ukraine*.

Cabinet of Ministers of Ukraine. (2023, May 2). *Resolution No. 434 "On amendments to the pilot-project rules for DREAM…"*. *Official Bulletin of Ukraine*.

Cabinet of Ministers of Ukraine. (2023, July 18). *Resolution No. 731 "On approval of procedures for the restoration and development of regions and territorial communities"*. *Official Bulletin of Ukraine*.

Cabinet of Ministers of Ukraine. (2024, May 14). *Resolution No. 549 "On establishment of the Strategic Investment Council"*. *Official Bulletin of Ukraine*.

Cabinet of Ministers of Ukraine. (2024, August 9). *Resolution No. 903 "Some questions of preparation, submission, evaluation and prioritisation criteria of public investment project concepts for 2025"*. *Official Bulletin of Ukraine*.

Cabinet of Ministers of Ukraine. (2025, March 7). *Resolution No. 272 "On updates to the pilot-project framework for DREAM…"*. *Official Bulletin of Ukraine*.

Verkhovna Rada of Ukraine. (2003). *Law on the protection of information in information and communication technology systems* (No. 80/94-VR).

Verkhovna Rada of Ukraine. (2020). *Law on the Budget Code* (No. 2456-VII).

Appendix

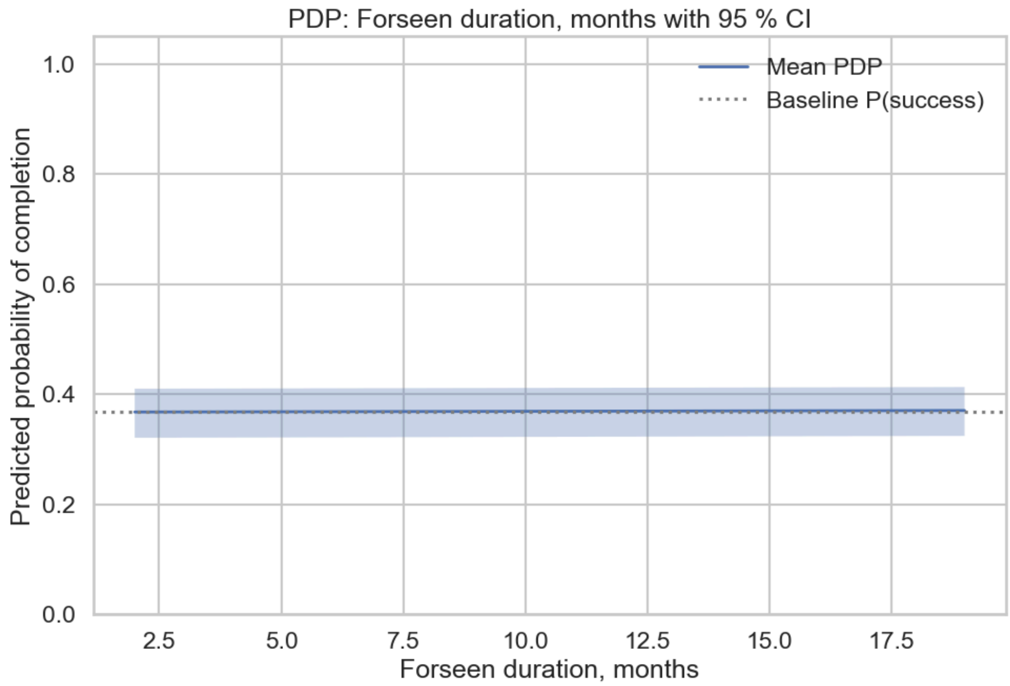

Figure A.1. Partial-dependence of the foreseen project duration (months) on predicted completion probability (95 % CI); n = 190, DREAM core register.

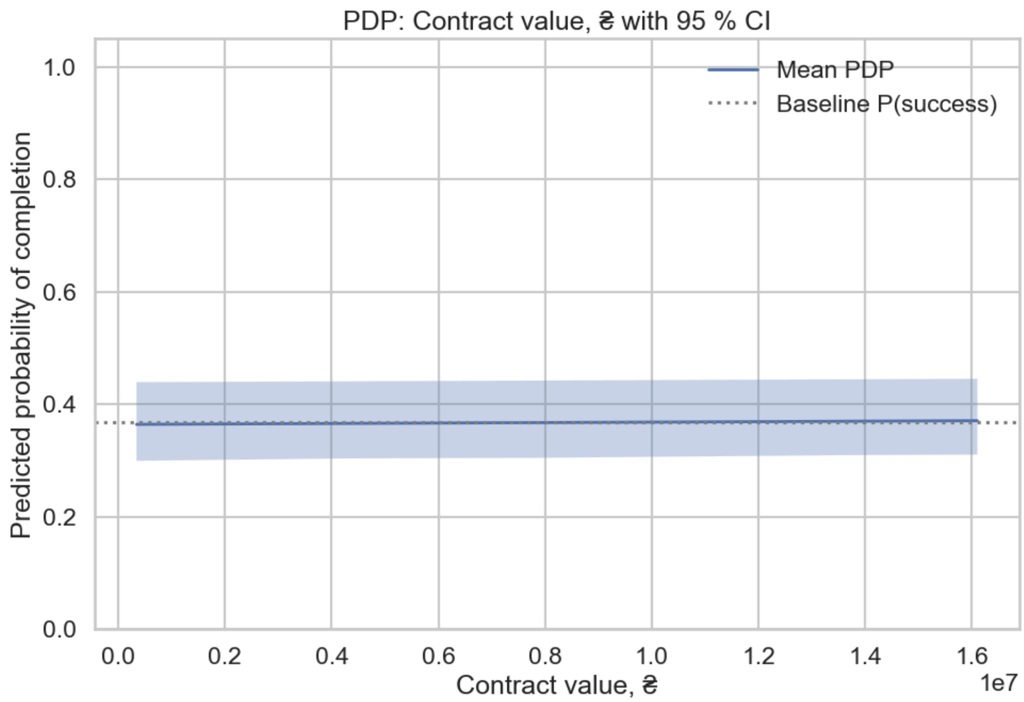

Figure A.2. Partial-dependence of the contract value (UAH) on predicted completion probability (95 % CI); n = 190, DREAM core register.

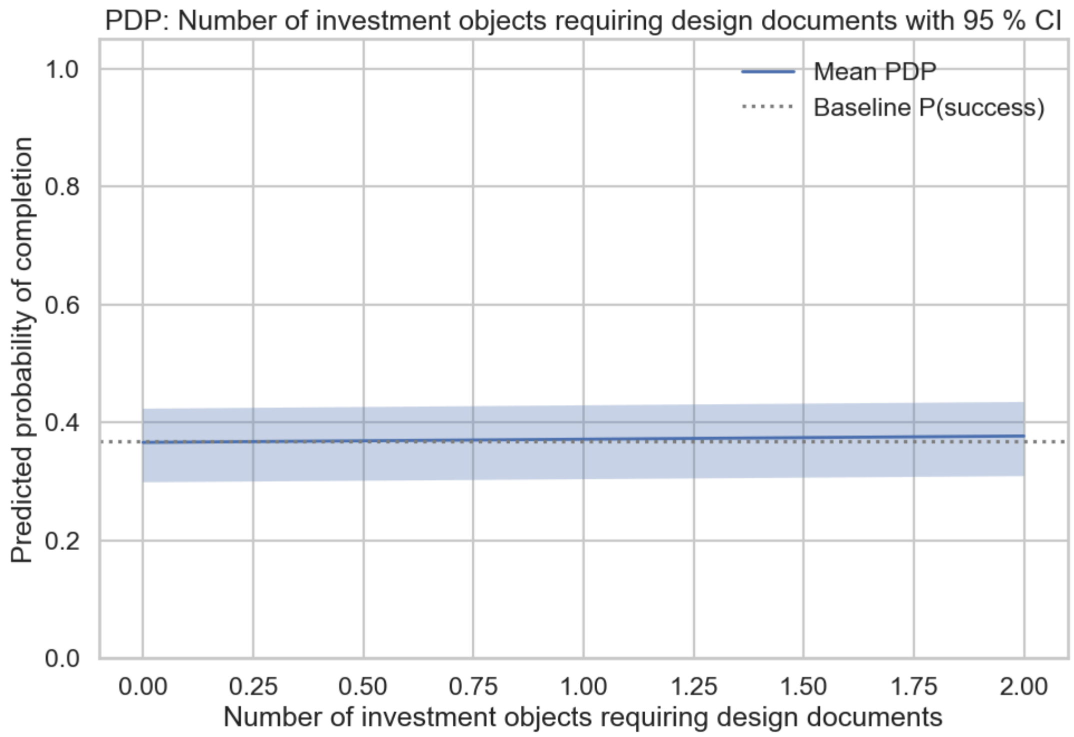

Figure A.3. Partial-dependence of the number of investment objects requiring design documents on predicted completion probability (95 % CI); n = 190, DREAM core register.

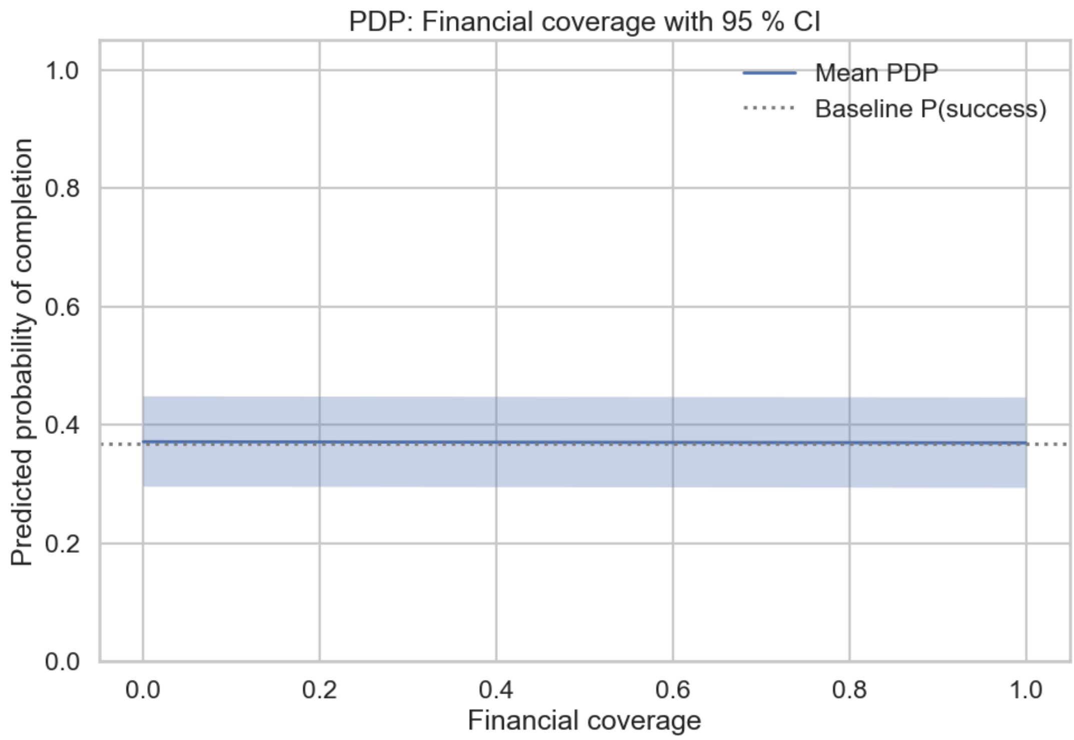

Figure A.4. Partial-dependence of the financial-coverage ratio (committed / estimated budget) on predicted completion probability (95 % CI); n = 142, variable derived from DREAM monetary columns.

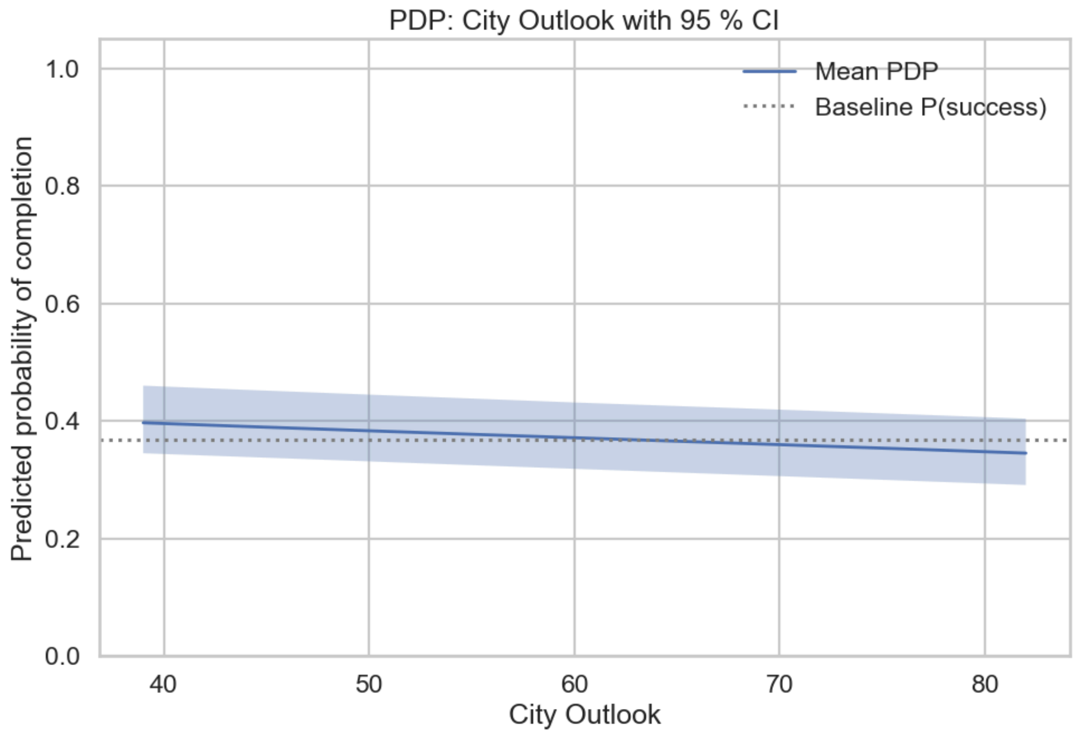

Figure A.5. Partial-dependence of public City Outlook sentiment on predicted completion probability (95 % CI); n = 170, IRI survey imputed to regions.

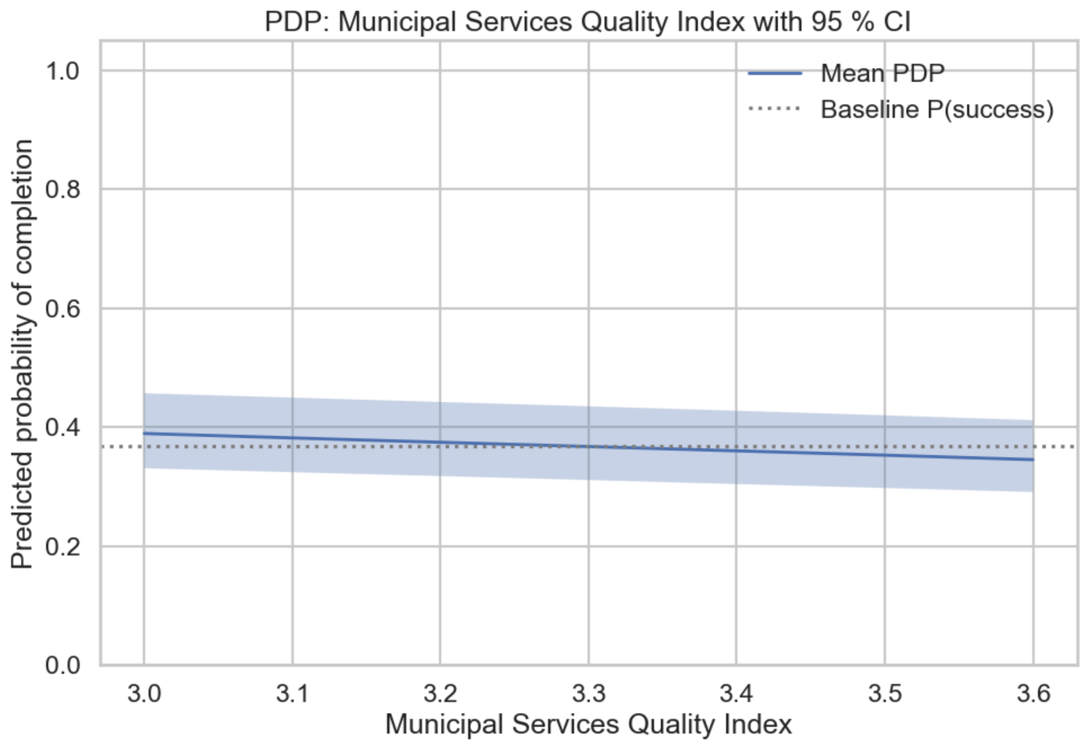

Figure A.6. Partial-dependence of the Municipal Services Quality Index on predicted completion probability (95 % CI); n = 170, IRI survey imputed to regions.

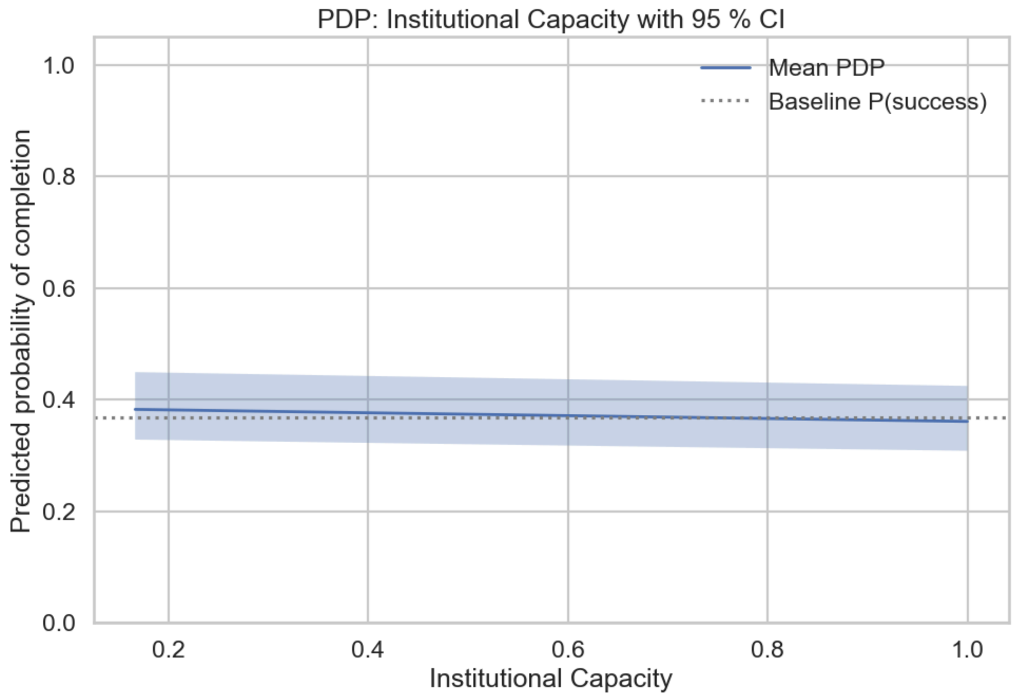

Figure A.7. Partial-dependence of regional Institutional Capacity (Digital-Index sub-score) on predicted completion probability (95 % CI); n = 189, Ministry of Digital Transformation dataset.

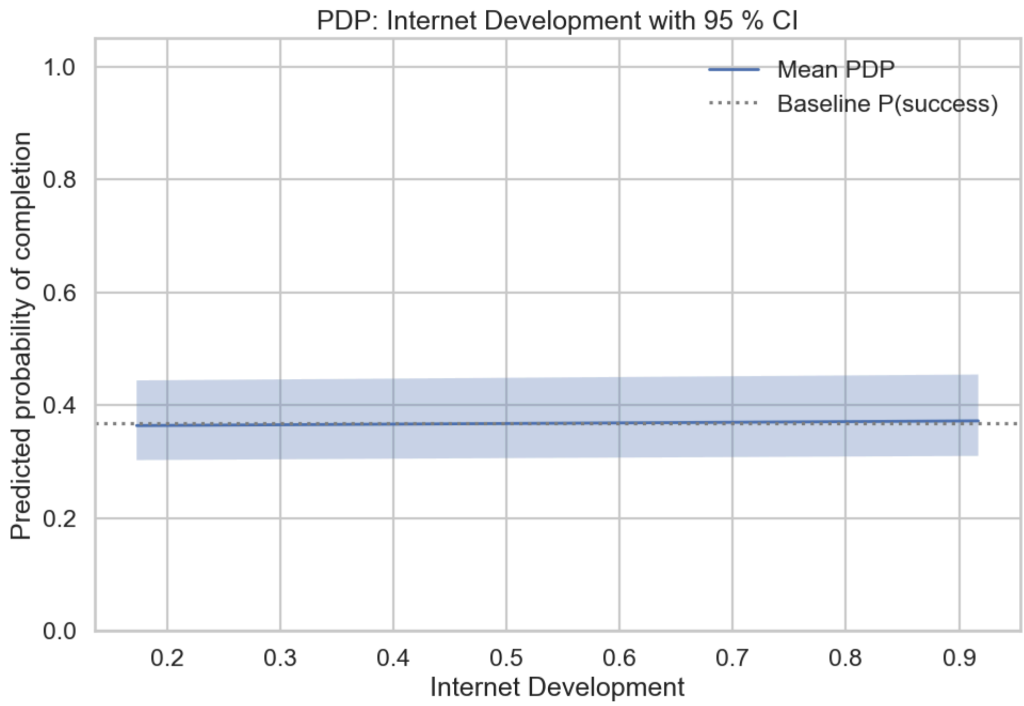

Figure A.8. Partial-dependence of regional Internet Development (Digital-Index sub-score) on predicted completion probability (95 % CI); n = 189, Ministry of Digital Transformation dataset.

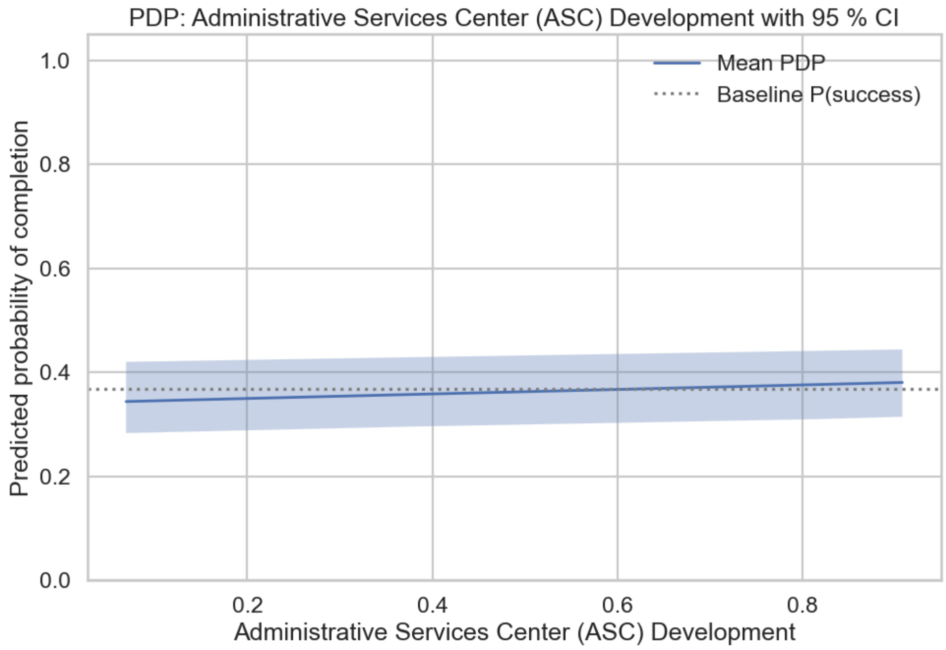

Figure A.9. Partial-dependence of Administrative Services Centre (ASC) development (Digital-Index sub-score) on predicted completion probability (95 % CI); n = 189, Ministry of Digital Transformation dataset.

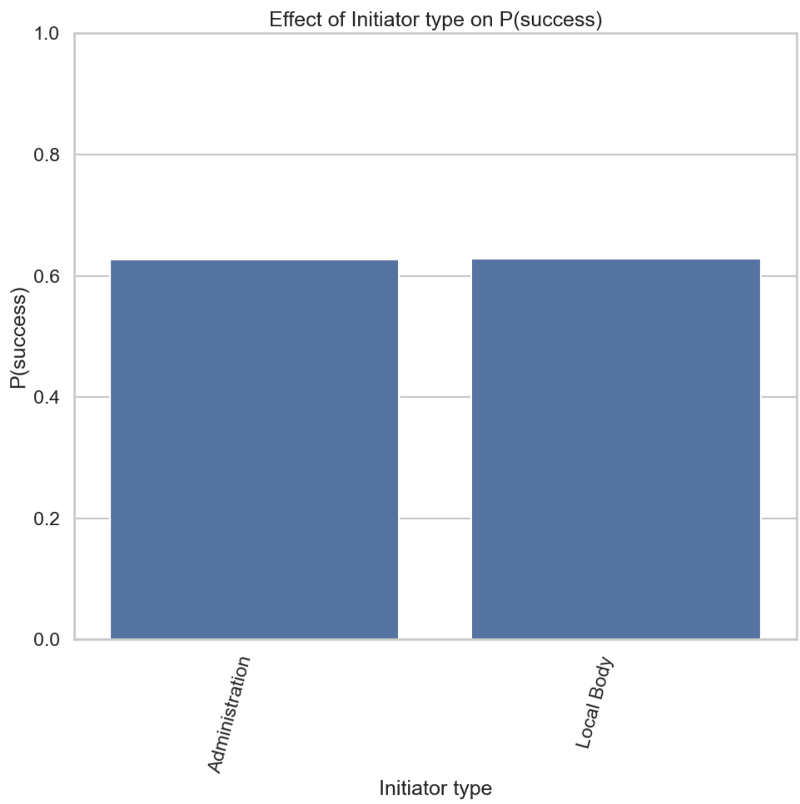

Figure A.10. Average predicted completion probability by initiator type; n = 190, DREAM core register.

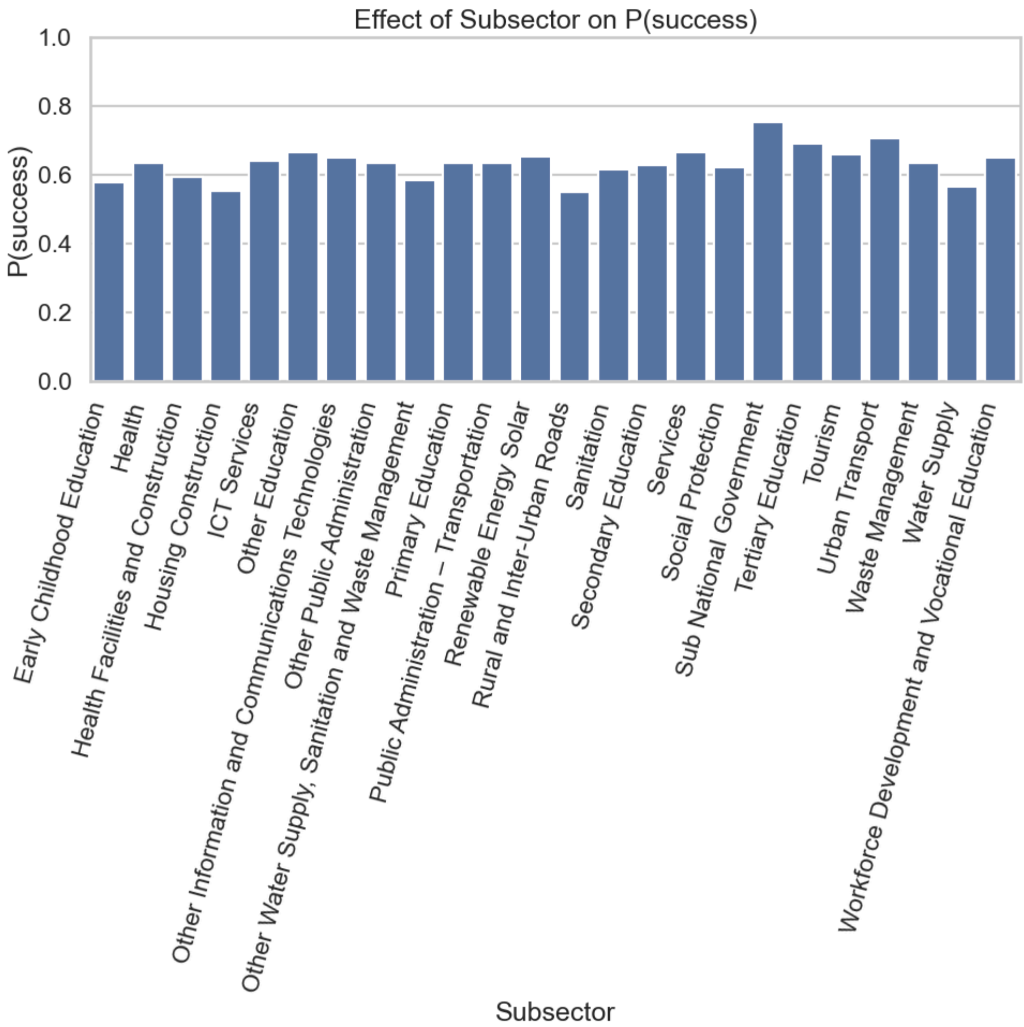

Figure A.11. Average predicted completion probability by subsector; n = 190, DREAM core register.

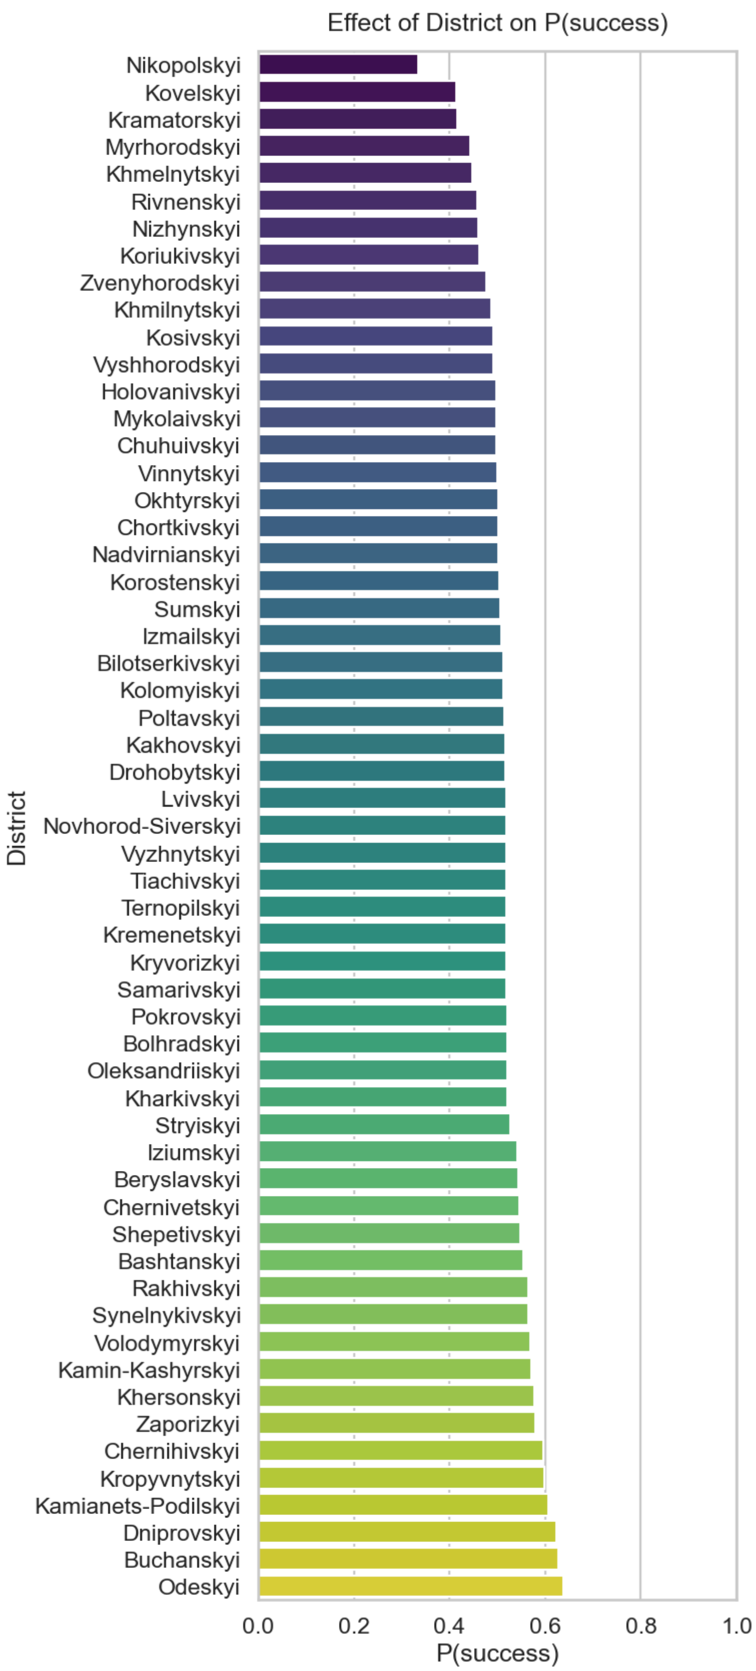

Figure A.12. Average predicted completion probability by district; n = 190, DREAM core register.

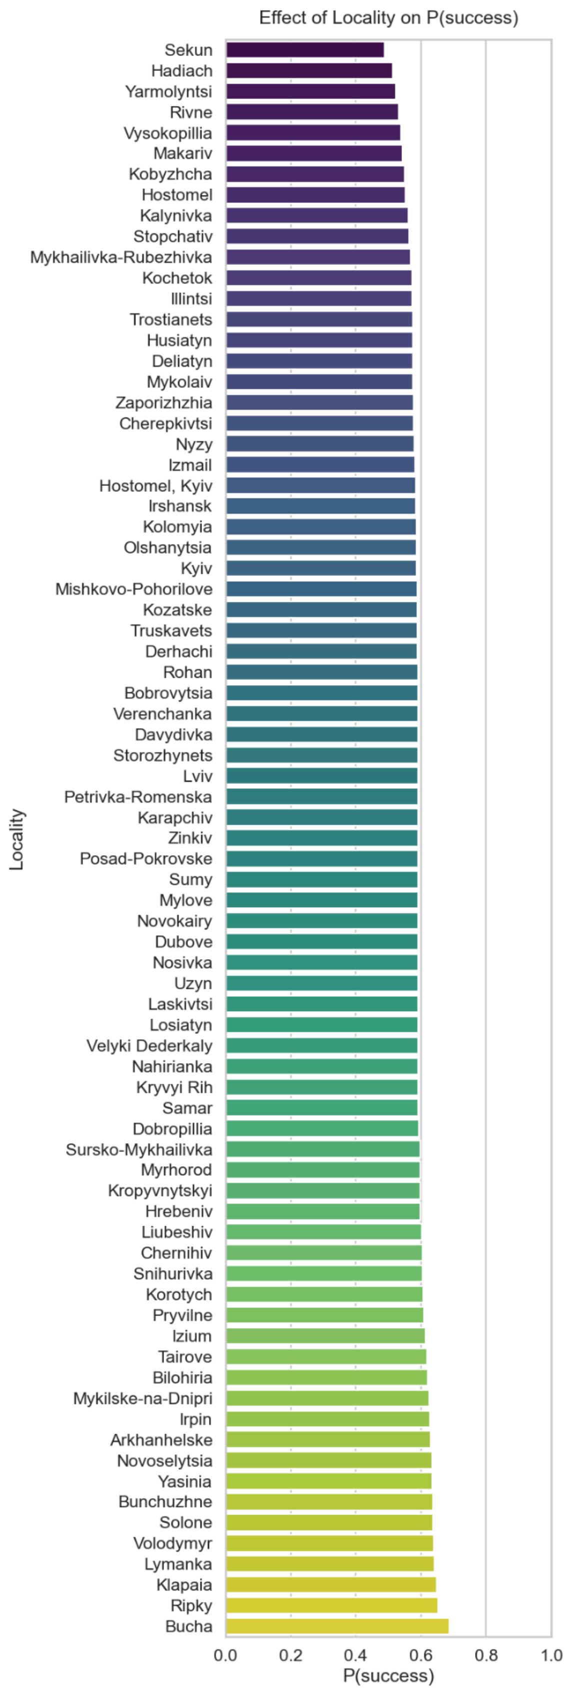

Figure A.13. Average predicted completion probability by locality; n = 190, DREAM core register.

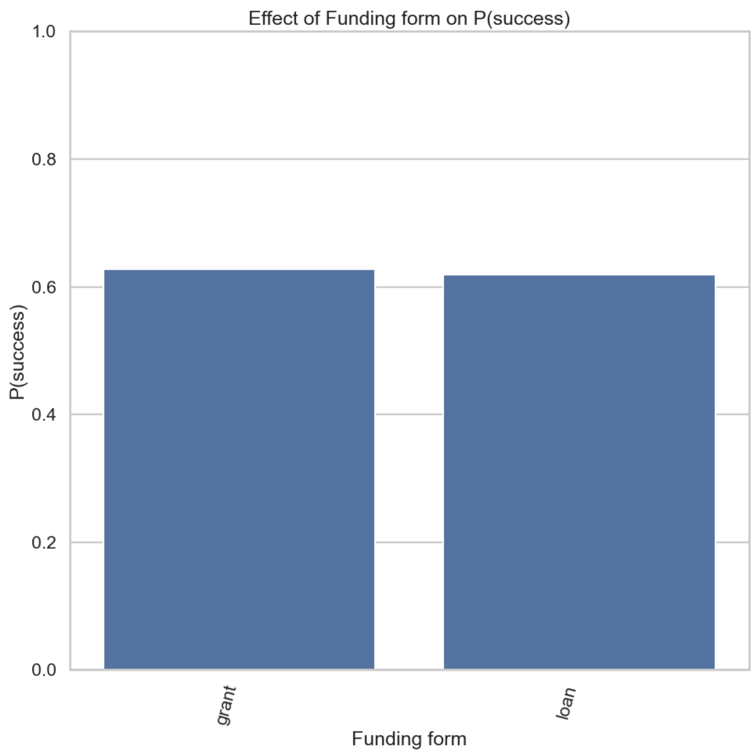

Figure A.14. Average predicted completion probability by funding form; n = 190, DREAM core register.

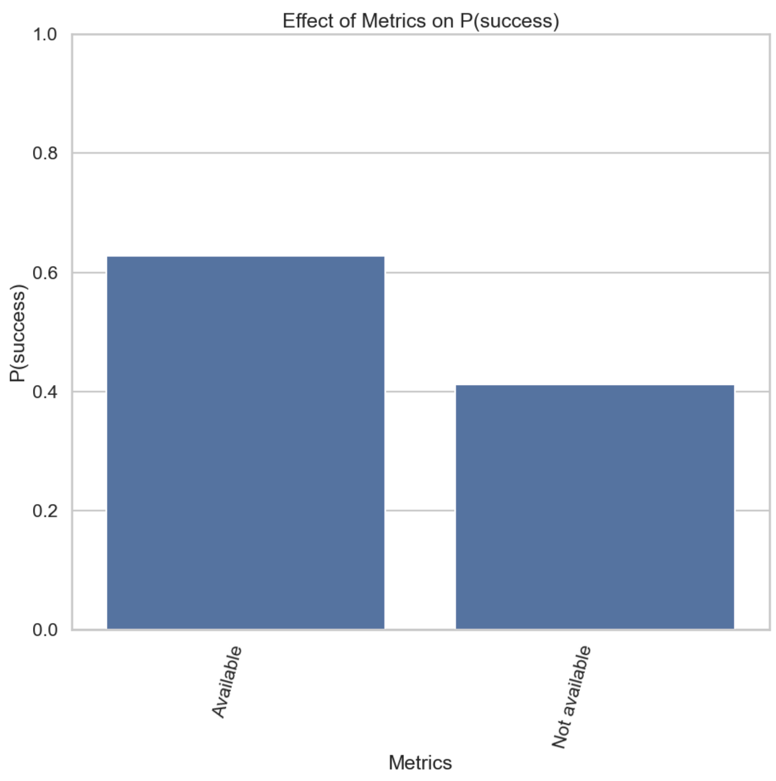

Figure A.15. Average predicted completion probability by availability of project-level performance metrics; n = 190, DREAM Completeness dataset.

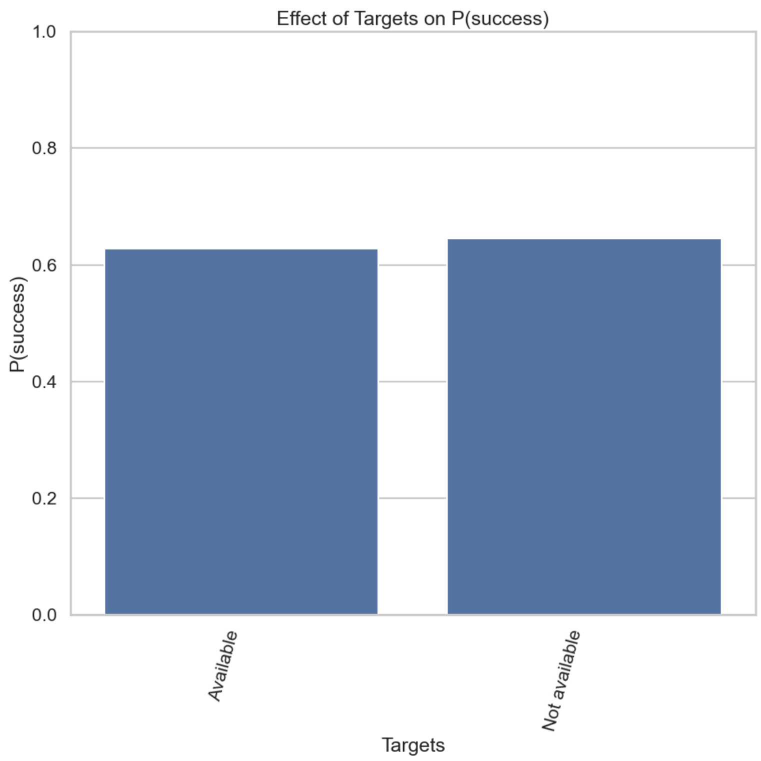

Figure A.16. Average predicted completion probability by availability of target indicators; n = 190, DREAM Completeness dataset.

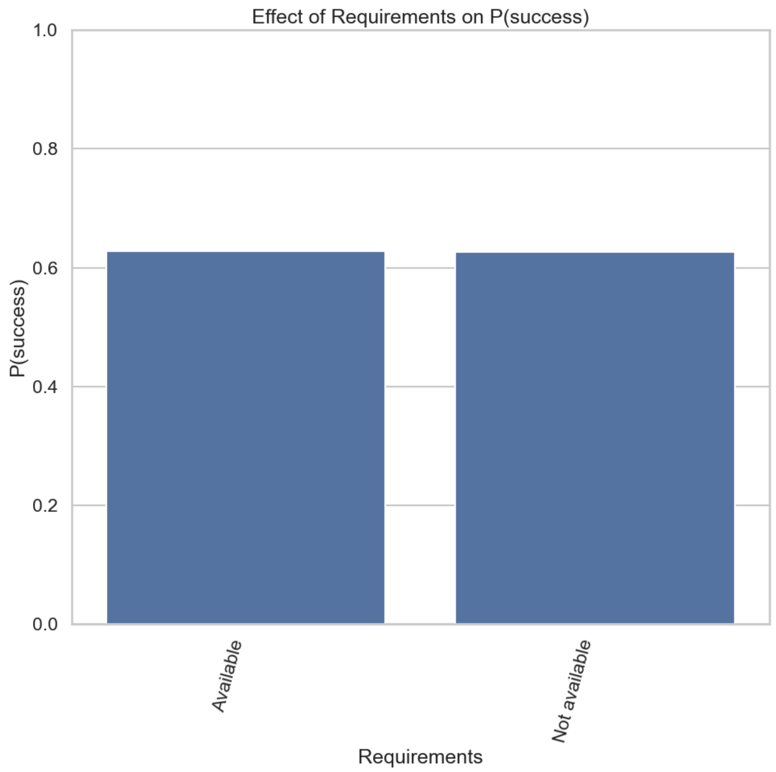

Figure A.17. Average predicted completion probability by availability of project-execution requirements; n = 190, DREAM Completeness dataset.

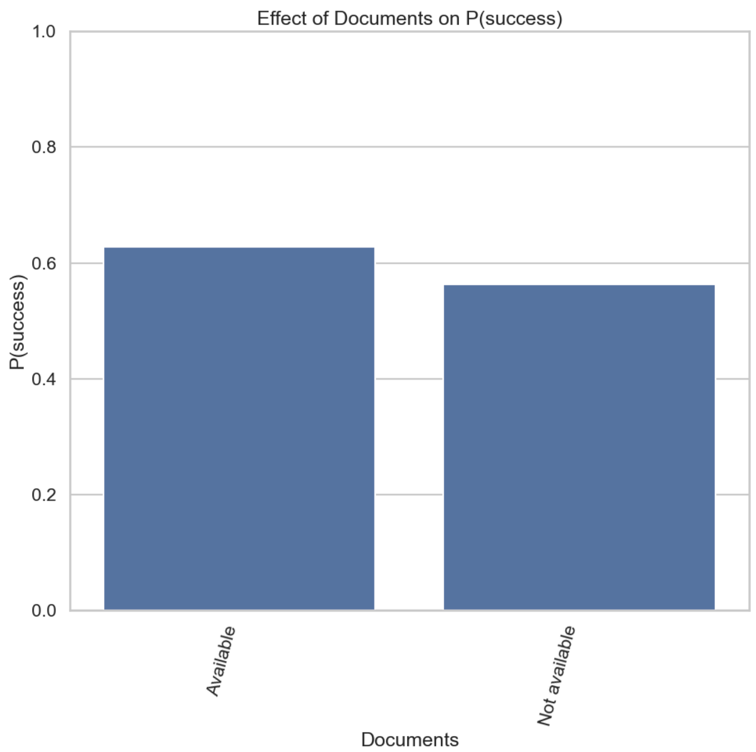

Figure A.18. Average predicted completion probability by availability of strategic or planning documents; n = 190, DREAM Completeness dataset.

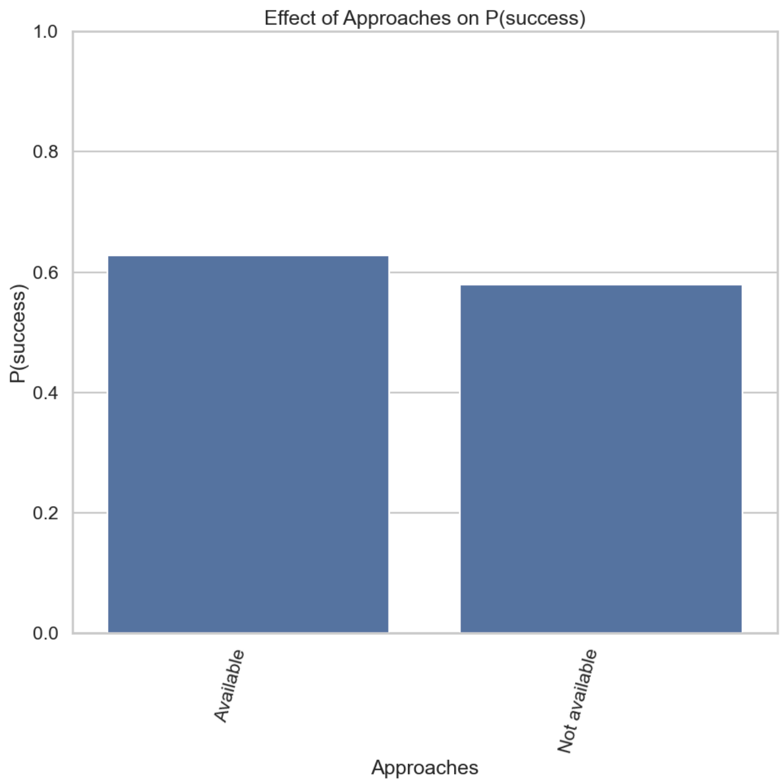

Figure A.19. Average predicted completion probability by availability of methodological approaches; n = 190, DREAM Completeness dataset.

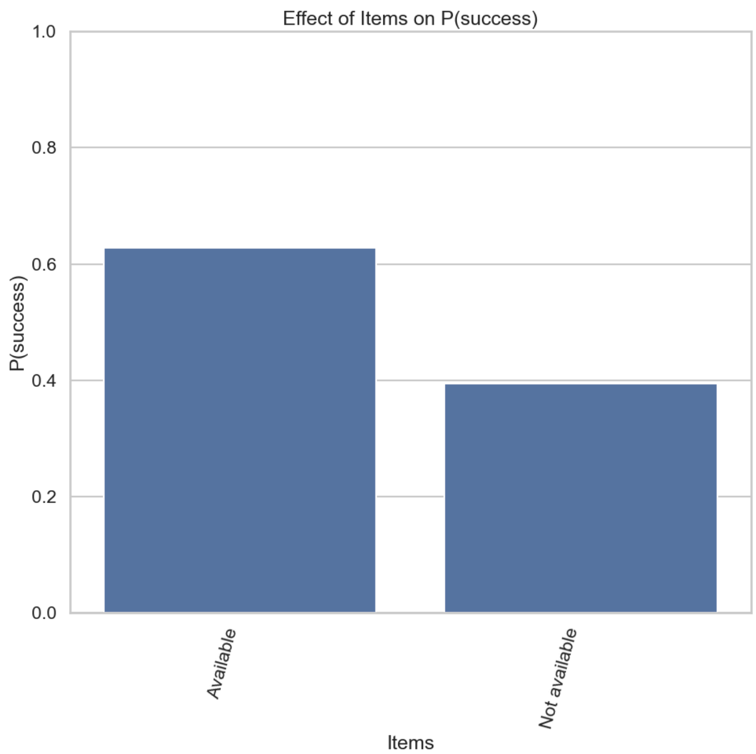

Figure A.20. Average predicted completion probability by availability of an itemised bill-of-materials; n = 190, DREAM Completeness dataset.

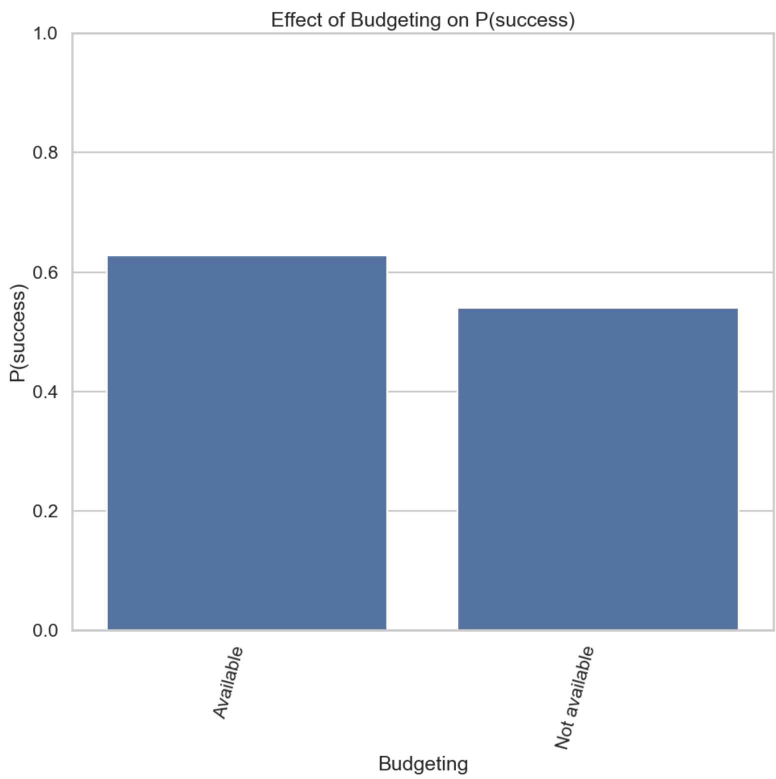

Figure A.21. Average predicted completion probability by availability of a detailed budgeting schedule; n = 190, DREAM Completeness dataset.

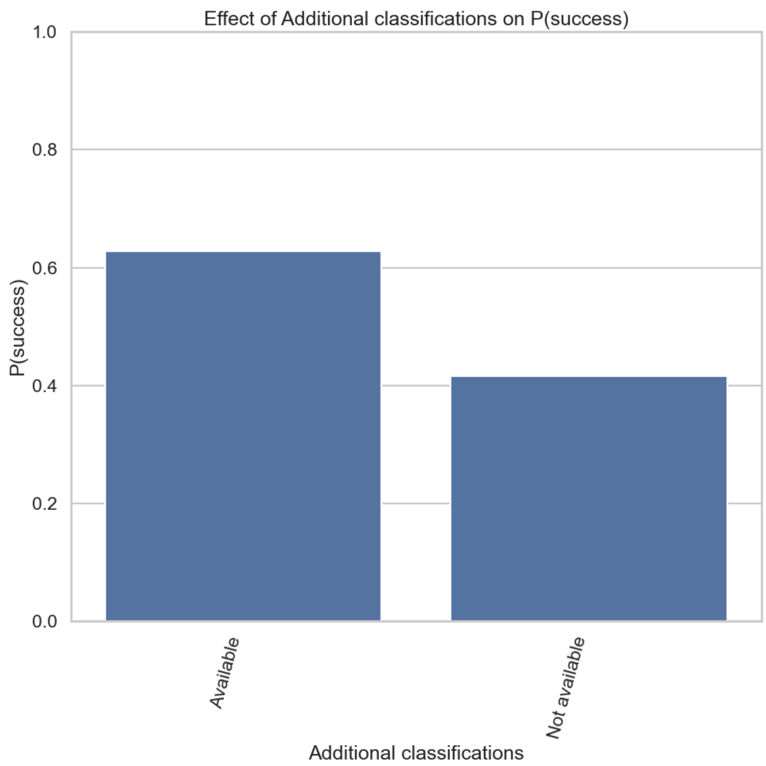

Figure A.22. Average predicted completion probability by presence of supplementary project classifications; n = 190, DREAM Completeness dataset.

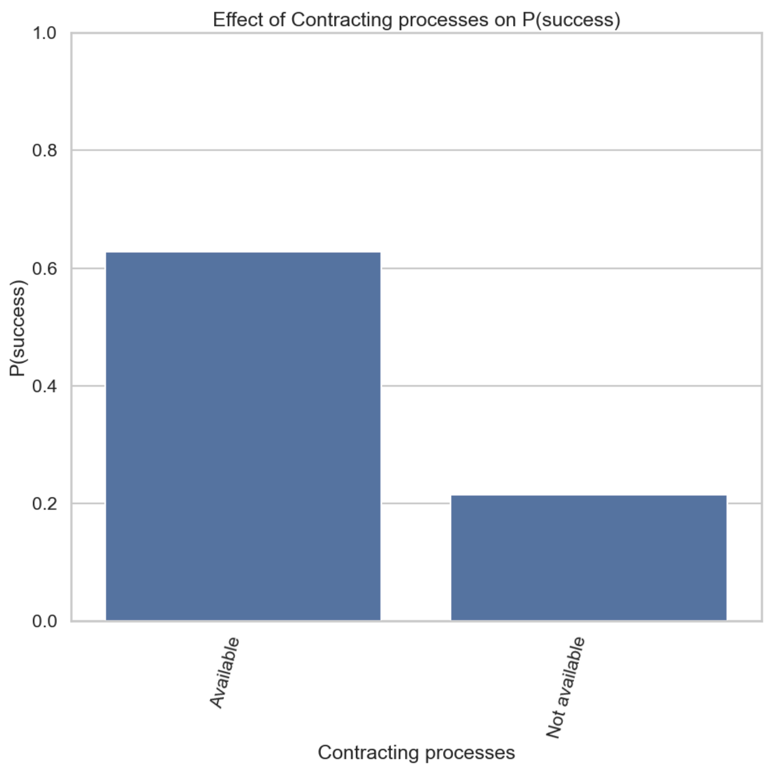

Figure A.23. Average predicted completion probability by publication of contracting-process details; n = 190, DREAM Completeness dataset.