POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Executive Summary of the Dissertation

Supporting Italian SMEs Facing Succession: An AI-Assisted Acquirer Recommendation Approach

Laurea Magistrale in Management Engineering - Ingegneria Gestionale

Author: **Dorian Silvio de Groot**

Advisor: **Prof. Giancarlo Giudici**

Academic year: **2025-2026**

---

1. Introduction

Italy's productive system is exceptionally dependent on small and medium-sized enterprises (SMEs), which are predominantly family-controlled and owner-dependent. Demographic aging and the decline of intra-family succession are turning business continuity into a systemic problem. A large share of family firms will face a leadership and ownership transition within the coming decade. At the same time, willing and suitable internal successors are increasingly scarce [1]. When intra-family succession is unavailable, delayed, or organizationally unsuitable, continuity comes to depend on transferring the firm to an external owner.

Mergers and acquisitions (M&A) provide a viable external-succession channel, but post-acquisition outcomes are conditional on the quality of the match between target and acquirer [4, 12]. The practical bottleneck is therefore upstream and on the sell side: identifying, prioritizing, and shortlisting plausible acquirers at acceptable cost and time. Recent work shows that AI can support exactly this stage, by predicting deal outcomes and by recommending acquirers from firm characteristics [10, 16], yet no operational framework targets succession-constrained Italian SMEs from the sell side.

As a proof of concept, this dissertation develops and evaluates a transparent, leakage-controlled, and interpretable recommendation prototype that ranks and screens candidate acquirers for Italian SMEs from pre-deal information, paving the way for a more advanced decision-support tool. The conceptual grounding and the research questions that organize this summary are set out next.

2. Theoretical Background

Succession as a systemic problem. Italy's productive base rests on millions of small, family-controlled, owner-dependent firms [7], and demographic aging is turning leadership and ownership transfer into a systemic issue. Recent estimates imply that more than one million firms may face a generational transition within the decade. Meanwhile, exiting family leaders increasingly outnumber entering ones, and only a minority of firms have a formal succession plan [1, 6]. When an internal heir is unavailable, unwilling, or unsuitable, continuity comes to depend on transferring the firm to an external owner.

M&A as a conditional continuity mechanism. Mergers and acquisitions can preserve continuity by keeping a viable firm operating, retaining its workforce and relationships, and embedding it in a larger platform. However, this is highly conditional. Post-acquisition outcomes are heterogeneous and depend above all on how well the acquirer fits the target [2, 4, 12]. The practical question is therefore not only whether an SME can be sold, but whether it can be sold to a *suitable* acquirer. This reframes external succession as a target–acquirer matching problem.

AI for the matching problem, and the gap. A growing literature applies machine learning to M&A, predicting deal outcomes, recommending acquirers from firm characteristics, and broadening the reach of buyer search [8, 10, 16]. Yet these methods are developed for settings unlike the fragmented, family-owned, information-opaque Italian SME landscape. Three gaps follow. First, a *context* gap, as the methods are not built for this environment. Second, a *problem-definition* gap, as most work predicts whether a deal occurs rather than ranking suitable acquirers for a firm that needs an external successor. Third, an *implementation* gap, as little evidence exists on turning such methods into an interpretable framework feasible under realistic data constraints.
Against this background, the main research question is the following:

> *To what extent can an AI-assisted recommendation model support Italian SMEs facing succession-related continuity challenges by identifying and prioritizing external acquirers and generating candidate matches associated with stronger ex post acquisition outcomes?*

It is broken into six sub-questions, answered one by one in Section 4:

RQ1 Does the model rank unseen deals consistently with their realized outcomes, and does it outperform a linear model and simple practitioner heuristics?

RQ2 Can the model shortlist, among previously unseen Italian deals, the target–acquirer matches that achieve top-quartile ex post acquisition performance?

RQ3 Does this ranking performance persist for deals occurring later in time and for acquirers never observed during training?

RQ4 Which broad families of pre-deal target–acquirer characteristics are most informative?

RQ5 Which individual features are most informative, and how do they help interpret acquirer suitability in the Italian SME succession context?

RQ6 What are the main practical and methodological limitations of using the model as a decision-support tool for external succession?

3. Methodology

3.1. Data and Performance Targets

The analysis draws on completed controlling transactions (final stake 50%–100%) involving Italian targets from the Orbis M&A database [13], with pre- and post-deal financials for both the target and the acquirer. Following the M&A performance-measurement literature [5, 9, 17], the raw accounting changes are benchmarked against year–industry references from the BACH database [3], which separates transaction-specific performance from broader sector and macroeconomic trends.
Match success is captured by three *benchmark-adjusted* accounting outcomes, a choice well established in the accounting-based M&A success-measurement literature [11, 15]: return on assets (ROA), return on sales (ROS), and revenue turnover (TURN). The three are modeled separately to preserve their distinct economic meaning. Each is expressed as a percentile-rank score from 0 to 100, the direct operationalization of a relative ranking estimand and robust to heavy accounting tails by construction. Two measurement-validity rules then define the estimation sample: post-deal financials are never imputed, and, because the post-deal outcome is read from the acquirer's consolidated accounts, the target must be economically material (at least 10% of combined pre-deal assets). This yields 2,615 deals. The full observed-outcome population (without the materiality constraint) of 5,856 deals serves as a robustness check.

3.2. Recommendation Pipeline

The prototype is organized as four sequential layers.

Input layer. Assembles the two external sources, the Orbis M&A transactions and the BACH year–industry benchmarks, into a single integrated dataset.

Processing layer. Cleans and harmonizes the raw fields, then imputes the missing pre-deal predictors through economically grounded rules. From these it engineers a deliberately broad predictor set. The feature families capture target–acquirer structural relatedness, relative size, and pre-deal financial ratios and levels. They also include the combined entity's pre-deal profile and the target–acquirer fit gaps between them. A final pair of families adds signals captured by the acquirer's track record on strictly earlier deals and text-derived similarity between business descriptions. The result is a leakage-controlled, model-ready dataset.

Model layer. Merges the benchmarks, constructs the three target scores, and fixes the estimation sample and robustness scenario. For each outcome dimension it trains two target-specific CatBoost [14] models: a ranking head (a regressor that orders candidate acquirers) and a screening head (a classifier whose positive class is the top quartile of realized performance used to shortlist matches). Evaluation uses five-fold cross-validation grouped by acquirer, so that all of an acquirer's deals fall in the same fold and every metric is measured on acquirers absent from training. The ranking is benchmarked against a ridge-regression baseline, a mean-predicting dummy, and single-variable practitioner heuristics, and stress-tested out-of-time and on unseen acquirers.

Output layer. Delivers the predicted scores, the ranked and screened shortlist, and the SHAP-based interpretation and feature-block ablation that explain each recommendation.

4. Results and Discussion

Ranking quality and the gain over simpler alternatives (RQ1). The ranking head attains Spearman correlations of 0.665, 0.633, and 0.644 for ROA, ROS, and TURN, separating the top from the bottom predicted decile by 71–76 percentile points (Figure 1). The matches it favors realize outcomes between the 80th and 89th percentile of the sample, while those it discards land near the 12th. This ordering reflects learned structure, not a simple rule. It exceeds a regularized linear ranker on every dimension by +0.058 to +0.084 (paired-bootstrap intervals exclude zero), and no single-variable practitioner heuristic exceeds $|\rho| = 0.15$ for ROA or TURN. Even the intuitive same-industry rule shows essentially no association with realized outcomes, and the strongest heuristic, the acquirer's own pre-deal margins, reaches only 0.26–0.31, less than half the model's 0.63 on the same dimension. What predicts a strong outcome is the financial fit between the two firms, not a shared sector code. The model's gain over both the linear ranker and the simple rules is one of non-linearity. It combines these individually weak signals and learns the target–acquirer fit that no other method can capture.

Table 1: Ranking and screening performance on the base sample ($n = 2{,}615$).

Target	Rank ρ	Spread	ROC-AUC	P@10%
ROA	0.665	76.0	0.85	0.87
ROS	0.633	71.5	0.84	0.83
TURN	0.644	70.6	0.81	0.72

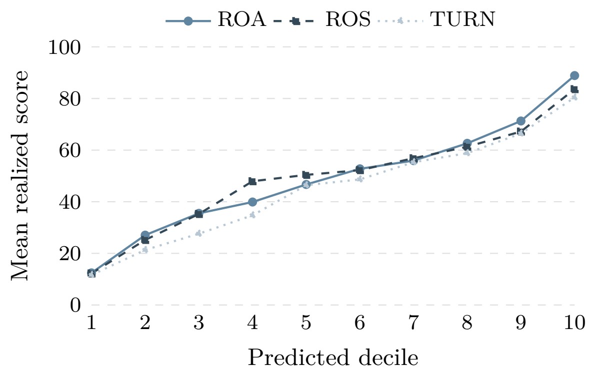

Figure 1: Mean realized benchmark-adjusted outcome by model-predicted decile (out-of-fold).

Identifying top-quartile matches (RQ2). Table 1 reports headline performance on the base sample. The dedicated screening classifier separates top-quartile matches from the rest with ROC-AUC between 0.81 and 0.85 and top-decile precision between 0.72 and 0.87, against a base rate near 0.25: a roughly three to three-and-a-half-fold concentration of genuinely strong matches. The cumulative-gains curve in Figure 2 makes the operational payoff explicit. Pursuing only the top 10% of the ranked list already captures about a third of all top-quartile matches (34% for ROA, 33% for ROS, 29% for TURN), and the top 20% captures roughly half (52%, 50%, 47%). The trade-off is the familiar one: extending the shortlist captures more of the strong matches but lowers precision. The advantage over random search nonetheless stays large across the depths an owner can realistically screen, so the shortlist can be matched to the evaluation capacity available. This is the result that matters most for the succession use case, where what is needed is a short list of candidates worth pursuing rather than a precise score for every buyer.

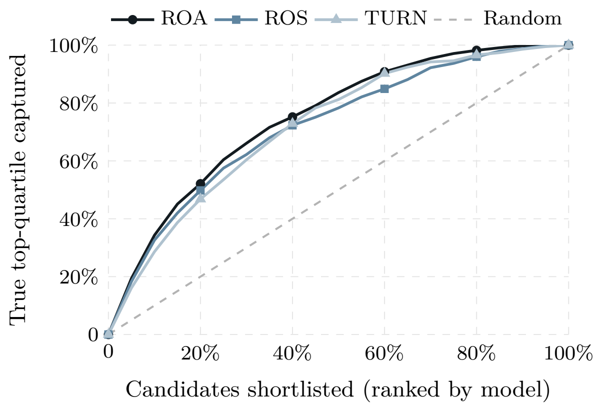

Figure 2: Cumulative gains of the model-ranked shortlist: share of all true top-quartile matches captured as the shortlist is extended.

Generalization across time and to unseen acquirers (RQ3). Acquirer-grouped cross-validation already evaluates every metric on acquirers held out of the training fold. However, two further tests probe stronger forms of generalization (Table 2). Trained only on deals completed up to 2018 and tested on later transactions, the model retains strong rank correlations, for ROS even exceeding the in-sample benchmark. This temporal stability also reinforces the value of the benchmark-adjusted pipeline, in which the BACH year–industry information helps separate transaction-relevant signal from broader sectoral and macroeconomic trends. Restricting that later test set to acquirers that never appear in the pre-2018 training period at all (genuinely new buyers rather than merely held-out ones), leaves performance essentially unchanged. The recommendations therefore rest on a transferable mapping from pre-deal characteristics to outcomes, not on memorizing specific recurring acquirers.

Table 2: Generalization: Spearman ρ under acquirer-grouped cross-validation, out-of-time, and unseen-acquirer validation.

Validation	ROA	ROS	TURN
Acquirer-grouped CV	0.665	0.633	0.644
Out-of-time ($\leq$ 2018)	0.596	0.641	0.616
+ unseen acquirer	0.590	0.652	0.607

What drives the recommendations (RQ4 and RQ5). Feature-importance and ablation evidence give the prototype an interpretability uncommon in tree-based pipelines, and point to a single, economically coherent mechanism: the quality of the match. At the family level (Table 3), financial ratios and levels dominate (58%–64% of importance), the target–acquirer gap measures form the second-largest family (24%–33%, rising from ROA through ROS to TURN), and text-derived features add almost nothing. At the individual level, three ingredients recur across all targets. The first is the *starting position* from which performance is measured: the target's own pre-deal profitability, margin, and turnover, together with the combined pre-deal baseline. The target's pre-deal ROA is the single most important predictor for ROA. The second is *target–acquirer fit*: the margin and turnover gaps, the turnover gap alone carrying about 22% of importance for TURN. The third is the acquirer's *track record* (5%–12%), led by the score of its most recent prior deal. Removing the financial block that contains these signals costs only about 0.07–0.11 in rank correlation, and the findings are stable across tighter sample definitions and on the full observed-outcome population. The same mechanism explains the generalization above: a match-based signal defined for every candidate pair does not depend on recognizing a specific acquirer.

Table 3: Feature-family importance: share of total mean absolute SHAP by outcome (base sample).

Feature family	ROA	ROS	TURN
Financial ratios & levels	64.0	60.1	57.6
of which combined baseline	12.2	5.9	7.7
of which acquirer history	5.2	11.6	9.2
Target–acquirer gaps	24.4	29.2	33.2
Log-levels	5.1	5.1	8.1
Text-derived	5.9	5.1	—
Interactions	0.5	0.5	1.1

The prototype as a proof-of-concept decision-support tool. Taken together, the evidence supports three claims: observable pre-deal information identifies top-quartile matches with precision far above the base rate, ranks them in line with realized outcomes and beyond what a linear model or practitioner heuristics achieve. It does so through an interpretable, match-based mechanism that generalizes prospectively and to unseen buyers. As a proof of concept, the prototype is best understood as a first cut at decision support that widens and orders the candidate set, intended to complement rather than replace valuation, due diligence, and negotiation. The potential gain is concrete: increasing the hit rate of a shortlist from roughly one-in-four to between three-in-four and nearly nine-in-ten would be a material reduction in search cost. The results therefore pave the way for a more advanced tool, built on the same approach. Such a tool could serve three groups: a transparent shortlist for an owner weighing succession, a first-pass filter for advisors, and, inverted, a way for investors to surface targets that fit their own profile. The trained framework is exposed through a publicly accessible web prototype at `acquirerfit-demo.com`, which makes the decision-support role directly testable and shows that the route from this proof of concept toward such a tool is open.

Limitations (RQ6). The framework is bounded in four ways. First, the data: Orbis covers the smallest SMEs least well, the extraction holds only a single pre-deal year and the first post-deal year, and about a quarter of the pre-deal financial fields are imputed. Second, the track-record signal is available only for serial acquirers, about 28% of deals. Third, the scores are relative percentile positions rather than absolute forecasts, and the validated scope covers only economically material targets. Fourth, success is judged on first-year accounting outcomes, so non-accounting dimensions of continuity, such as employment, supplier and customer relationships, and brand survival, lie outside the frame.

Toward a more advanced tool. Each limitation marks a development direction. The first-appearance acquirer gap could be narrowed by transferring a track-record-like signal across similar acquirers or from external capability proxies. Longer pre-deal histories and explicit relational and territorial data, such as industrial-district and supplier–customer linkages, would widen the observable space. The notion of success could be broadened to longer post-deal windows and non-accounting outcomes such as employment and brand survival. On the tool side, connecting the prototype to a live company database would turn it from a manual evaluator into a systematic acquirer-discovery engine. Expert feedback loops and explainability dashboards would make it fit for advisory use.

5. Conclusions

The evidence answers the main research question in the affirmative: pre-deal information carries enough signal to both rank and shortlist target–acquirer matches for Italian SMEs on an evidence-based footing. This performance generalizes across time and to acquirers never seen in training. It rests on an interpretable, match-based mechanism: the target's own pre-deal starting position, the fit between the two firms' profiles, and the acquirer's demonstrated track record. As a proof of concept, the framework reframes external-succession support as a structured, transparent matching problem and provides a foundation on which a more advanced acquirer-recommendation tool, of the kind outlined above, can be built. It shows that AI-assisted acquirer recommendation is a viable direction for SME succession support, and that the route from this proof of concept toward a more advanced tool is open.

Bibliography

[1] AUB Observatory. XVII Osservatorio AUB: NextGen choice in generational transitions: Ownership family structure, changing competencies, and emerging trends. Technical report, AIDAF–EY Chair in Family Business Strategy, Università Bocconi, 2026. Edited by Fabio Quarato and Carlo Salvato.

[2] Selena Aureli. Performance of unlisted Italian companies acquired by multinationals from emerging markets: The case of Indian acquisitions. *Journal of Organizational Change Management*, 28(5):895–924, 2015.

[3] Banque de France. BACH database (bank for the accounts of companies harmonized), 2026.

[4] Luyi Chen and Sofia Poggi. Breaking the stereotype of Chinese acquisitions in Italy: An empirical study on Italian target firm performance and the influential factors affecting it. Master's thesis, Uppsala University, 2020.

[5] Margaret Cording, Petra Christmann, and Carmen Weigelt. Measuring theoretically complex constructs: The case of acquisition performance. *Strategic Organization*, 8(1):11–41, 2010.

[6] Andrea Costa. Leadership change in Italian family businesses, 2026.

[7] EIB. The digitalisation of small and medium-sized enterprises in Italy: Models for financing digital projects – summary report. *European Investment Bank*, 2025.

[8] Zhishuo Han. Does AI adoption in M&A teams improve deal performance? SSRN Working Paper, 2024.

[9] Jeffrey S. Harrison, Michael A. Hitt, Robert E. Hoskisson, and R. Duane Ireland. Synergies and post-acquisition performance: Differences versus similarities in resource allocations. *Journal of Management*, 17(1):173–190, 1991.

[10] Wilhelm Johnson Swegmark and Didrik Tvedt. AI assisted matching in mergers and acquisitions: A data-driven approach to identifying potential acquirers. Master's thesis, Chalmers University of Technology, Gothenburg, Sweden, 2026.

[11] John B. Jr. Kusewitt. An exploratory study of strategic acquisition factors relating to performance. *Strategic Management Journal*, 6(2):151–169, 1985.

[12] Tingting Liu, Zhongjin Lu, Tao Shu, and Fengrong Wei. Unique bidder-target relatedness and synergies creation in mergers and acquisitions. *Journal of Corporate Finance*, 73:102196, 2022.

[13] Orbis M&A. Orbis database: Global private company data & intelligence, 2026.

[14] Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. CatBoost: Unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems 31*, pages 6638–6648, 2018.

[15] Ioannis C. Thanos and Vassilis M. Papadakis. The use of accounting-based measures in measuring M&A performance: A review of five decades of research. In Sydney Finkelstein and Cary L. Cooper, editors, *Advances in Mergers and Acquisitions*, volume 10, pages 103–120. Emerald Group Publishing Limited, 2012.

[16] Zixuan Wang. *Strategic M&A Decisions Powered by Predictive Insights: Media and AI*. PhD thesis, The University of Hong Kong, 2025.

[17] Maurizio Zollo and Degenhard Meier. What is M&A performance? *Academy of Management Perspectives*, 22(3):55–77, 2008.