Does AI-generated listing language affect property market outcomes for investors?

Journal of Property Investment & Finance

Clayton Pilat and Chris Leishman
School of Economics, Adelaide University, Adelaide, Australia

Received 16 January 2026
Revised 26 April 2026
29 June 2026
Accepted 29 June 2026

Abstract

Purpose – This study investigates whether residential property listing descriptions continue to function as credible signals in the era of generative artificial intelligence (AI). It examines the linguistic differences between human-authored and AI-generated text and evaluates how these differences relate to three transaction outcomes: sale price, time on market and investor acquisition. The paper is concerned with market pricing dynamics rather than formal valuation in the International Valuation Standards (IVS) or Royal Institution of Chartered Surveyors (RICS) sense, and considers the implications of AI-mediated listing language for the residential investment market and the boundary between owner-occupier and investment-grade stock.

Design/methodology/approach – The research analysed 6,376 residential sales in South Australia between June 2023 and January 2025, drawn from an initial pool of approximately 10,000 listings. The study employed the DeskLib AI detection tool and computational linguistic analysis to construct continuous measures of AI-likeness. A subset of 5,368 cases was linked to a major rental platform to identify properties acquired by investors within twelve months of sale. Ordinary least squares regressions were used to test the relationship between AI-likeness and sale price and time on market, and a logistic regression was used to examine investor acquisition, controlling for standard property attributes.

Findings – AI-generated descriptions are systematically longer, more formulaic and more reliant on emphatic phrases compared to human-authored text. AI-likeness is associated with a small but statistically significant 2% reduction in sale price, properties selling approximately 9% faster, and 21% higher odds of investor acquisition within twelve months of sale. These results suggest that listing language is shifting from a costly signal of quality to a coordination device that improves market efficiency, and that the standardised character of AI-generated text appears to particularly resonate with investor buyers who prioritise transactional clarity over narrative distinctiveness.

Research limitations/implications – The study is limited to South Australian residential transactions between June 2023 and January 2025. AI detection remains probabilistic and vulnerable to obfuscation and subgroup bias, and the time-on-market model has modest explanatory power. Agents who adopt AI writing tools may also be more technologically sophisticated in other dimensions of marketing, which cannot be ruled out as a contributing factor. The investor flag is conservative, and findings may not extend to commercial markets where lease structures and income-based valuation dominate. Theoretically, the research updates signalling theory by showing that as production costs fall, listing language migrates from costly signal toward coordination device.

Practical implications – Generative AI offers efficiency gains, reducing time on market by approximately 9% with only a small effect on sale price. Agencies can use these tools to accelerate workflows but should avoid over-reliance to prevent linguistic homogenisation, supplementing AI outputs with bespoke, locally grounded detail to maintain differentiation. The association between AI-generated listings and higher odds of investor acquisition has implications for portfolio managers and analysts assessing the composition of residential demand, and may signal a quiet shift in the boundary between owner-occupier and investment-grade stock. Platforms may need to innovate via visual and interactive tools as textual variety diminishes.

Originality/value – The paper provides the first quantitative test of AI-likeness in relation to residential housing market outcomes, including investor acquisition as a distinct buyer segment relevant to property investment research. It advances signalling theory by demonstrating how the reduced cost of persuasive text erodes signal credibility while enhancing transaction speed through standardisation, and contributes methodologically by treating probabilistic AI detection scores as continuous research features for analysing market communication rather than as binary classifiers of authorship.

Keywords Market efficiency, Signalling theory, Generative AI, Property market analysis, AI detection, Listing descriptions, Residential property markets

Paper type Academic paper

Journal of Property Investment & Finance
Emerald Publishing Limited
e-ISSN: 1470-2002
p-ISSN: 1463-578X
DOI 10.1108/JPIF-01-2026-0008

1. Introduction

Property valuation [1] has long been central to the functioning of real estate markets, underpinning transactions, financing, taxation and investment decisions (Pagourtzi *et al.*, 2003). Traditional valuation approaches are based on comparable sales, income capitalisation, or cost methods. These provided the foundation for valuation practice, but they remain limited by their reliance on professional judgement, historical comparables and time-intensive processes. Studies show that such valuations often lag market conditions and deviate substantially from transaction prices (Kok *et al.*, 2017). The rapid growth of automated valuation models (AVMs), which leverage large-scale data and algorithmic methods, has sought to overcome these shortcomings by delivering faster, more scalable and more precise estimates (Topraklı, 2025).

AVMs have been particularly notable for their ability to integrate variables that move beyond the most powerful or obvious physical and neighbourhood amenity indicators that dominate traditional valuation methods. Local amenities and neighbourhood attributes have been shown to explain substantial variation in the values of one property compared to another (Kok *et al.*, 2017). The growing utility of AVMs signals a broader reorientation of valuation research, where qualitative cues such as photographs, textual descriptions and marketing techniques such as staging are treated not merely as embellishments but as signals with measurable influence on buyer behaviour (Luchtenberg *et al.*, 2019; Baur *et al.*, 2023; Wheaton and Xu, 2024). The credibility of these signals rests on their costliness and difficulty to imitate, as highlighted by signalling theory (Connelly *et al.*, 2011). Signals are most informative when they are observable and impose higher costs on low-quality sellers than on high-quality ones, making widespread imitation unattractive and limiting deceptive use. When signals become cheap and easy to reproduce, they lose credibility and receivers learn to discount them. With the rise of generative artificial intelligence (AI), this condition is challenged. The ability to cheaply automate persuasive text risks homogenising listing language and eroding its distinctiveness as a signal of quality (Kandipati, 2025; Sandler *et al.*, 2024).

At the same time, advances in AI detection and computational text analysis offer new methodological tools for interrogating these changes. Detectors such as DeskLib, benchmarked under the Robust AI Detection (RAID) framework, provide probabilistic measures of AI-likeness that can be mobilised as research features rather than as conclusive evidence (Dugan *et al.*, 2024; Elkhatat *et al.*, 2023). This methodological innovation creates an opportunity to study how AI-mediated language is diffusing through housing markets and whether its homogenisation affects market outcomes.

Against this backdrop, this study asks whether property listing descriptions continue to function as credible signals in an era of generative AI. Specifically, it investigates how the linguistic features of AI-generated versus human-authored text differ, and whether these differences bear on transaction outcomes such as sale price and time on market. In doing so, it contributes to three areas of knowledge: first, by reconceptualising listing text as a qualitative signal under conditions of technological change; second, by providing the first quantitative test of AI-likeness in relation to housing market outcomes; and third, by advancing methodological integration of computational detection with qualitative theories of signalling. The empirical analysis is confined to residential property, where listings are marketed to a broad buyer pool and comparable evidence is relatively dense. More heterogeneous or complex assets, particularly commercial property, are likely to follow different marketing processes and may not exhibit the same outcomes, a point we return to in the discussion.

A terminological and conceptual clarification is warranted before proceeding. This paper is concerned with market pricing dynamics, specifically how the linguistic features of listing descriptions relate to transaction outcomes including sale price and time on market, rather than with formal property valuation in the professional sense. The distinction is important. As French (2023, p. 300) observes, the overriding requirement of any market valuation is to "price to market": to estimate the price that would be achieved were a property sold on the open market at the date of valuation, a process he characterises as a heuristic exercise requiring the valuer to read the market and weigh comparable evidence through professional judgement. Formal valuation operates under internationally recognised standards: the IVSC (2025) makes explicit that no model, including one employing artificial intelligence or machine learning, can produce an International Valuation Standard (IVS) compliant valuation without the valuer applying professional judgement at each stage of data selection, modelling and reporting. Our empirical analysis concerns the upstream marketing environment, the listing descriptions that form part of what French (2023, p. 304) terms the "signposts" directing valuers toward market value and which the RICS (2023) professional standard on comparable evidence treats as market data supporting the valuation process. Throughout the paper we use the term "signal" in line with Connelly *et al.* (2011), though we acknowledge that the valuation profession typically refers to comparable evidence as "signposts" (French, 2023; RICS, 2023). Where our findings have implications for formal valuation practice, we return to this in the discussion.

2. Related works

2.1 Property valuation and marketing signals

Real estate valuation underpins the efficacy of many aspects of housing markets, including the viability of transactions, financing conditions and arrangements, computation of taxation liabilities and the assessment of investment metrics. Traditional valuation methods, such as comparable sales, income capitalisation, or cost-based approaches, rely heavily on historical evidence and professional judgement (Pagourtzi *et al.*, 2003). While these foundations will undoubtedly remain important, they are increasingly challenged by the growth of data-driven approaches and AVMs. Valuation practice has increasingly shifted away from purely expert-led valuation towards AVMs, driven by their lower cost, faster turnaround and comparable or superior predictive accuracy (Kok *et al.*, 2017; Topraklı, 2025).

Traditional valuations have well-documented limitations. They are costly and time-consuming, often lag market movements and can deviate substantially from transaction prices. Research by Kok *et al.* (2017) shows that US commercial appraisals deviated by 9–13% on average from sales and cost several thousand dollars per asset, with a three-week lag. By contrast, utilising an AVM achieved a median absolute error of 9%, with over half of estimates within 10% of realised prices. Topraklı (2025) reinforces this transition, arguing that AI-powered models not only improve accuracy and efficiency but also provide greater scalability. However, he cautions that challenges such as data bias, algorithmic opacity and the need for human oversight must be addressed to ensure responsible adoption. Together, these contributions highlight how AVMs disrupt a USD 90 billion appraisal industry while setting new standards for valuation precision and speed. A key strength of AVMs lies in their ability to incorporate a much broader set of inputs than traditional valuation methods, allowing them to capture aspects of location, neighbourhood context and market conditions that are difficult to assess consistently through manual valuation.

Recent AVM research shows that adding numerical representations of property descriptions can materially improve predictive accuracy relative to structured features alone (Baur *et al.*, 2023). More broadly, recent work extends this logic by quantifying qualitative housing attributes that are typically omitted from standard hedonic specifications. Using vision–language models, Wheaton and Xu (2024) derive measures of visual housing quality from listing photographs, capturing attributes such as condition and aesthetic appeal that support greater valuation accuracy and price transparency. The integration of listing-based qualitative signals into AVMs can therefore be understood through the lens of signalling theory, which explains how sellers convey information about otherwise unobservable quality to buyers in markets characterised by information gaps. Signals help reduce uncertainty when direct inspection or full information is not available. Connelly *et al.* (2011) emphasise that signals are effective only when they are observable and costly to fake. In real estate, professional photos, staging, or credible renovation claims can function as costly signals of quality, while generic textual embellishments are easier to imitate and therefore less reliable. Importantly, signals are dynamic: their value erodes as they become easier to reproduce. This raises a forward-looking question with the rise of generative AI. If persuasive descriptions can be cheaply automated, textual signals risk homogenisation, potentially undermining their discriminatory power in valuation contexts.

The literature on valuation reveals a clear evolution: from subjective, lagging and costly traditional valuations, to AVMs that harness structured, spatial and increasingly marketing-driven signals. AVMs deliver superior predictive performance, but their promise is coupled with challenges of bias, transparency and sustainability of signals in an AI-mediated landscape. For real estate research and practice, this evolution underscores both the opportunities and risks of integrating marketing signals into valuation models.

2.2 Qualitative signals in housing markets

Beyond price fundamentals, buyers and sellers rely on a wide range of qualitative cues to make sense of housing markets, reflecting the informational asymmetries that characterise these markets (Connelly *et al.*, 2011). Luchtenberg *et al*. (2019) demonstrate that simple listing signals, including positive wording, photo quality and price tiers, influence a buyer's intention to visit a property. Their experimental results suggest that listing language is particularly influential, exhibiting a stronger effect than images or price cues, which challenges the common assumption that photographs dominate early buyer impressions.

This evidence is further reinforced by research on framing effects in property markets. Levy *et al*. (2020) demonstrate that pessimistic framing produces a larger perceived price decrease than the equivalent optimistic framing produces an increase, consistent with prospect theory and loss aversion. Importantly, they find that market familiarity moderates these effects, with less experienced buyers showing greater susceptibility to listing language. While this paper adopts signalling theory as its primary theoretical lens, given its focus on the credibility and cost of language production rather than its psychological reception, the framing literature underscores that the words agents choose carry measurable behavioural consequences for buyers.

Behavioural economics research shows that these signals are interpreted through the lens of cognitive bias. You (2020) finds that anchoring, loss aversion and over-extrapolation of recent trends shape how buyers and sellers respond to information, producing non-linear price adjustments. Sellers are reluctant to accept nominal losses, while buyers overreact to periods of rapid growth, leading to sticky downward adjustments and overshooting in rising markets. This evidence makes clear that signals cannot be read as neutral, since they interact with systematic psychological tendencies that distort market behaviour. Other forms of signalling affect impressions rather than hard values. Lane *et al*. (2015) examine staging conditions, defined primarily by furniture quality and interior wall colour and find that these features enhance perceived liveability and overall impressions of a property but do not consistently translate into higher willingness to pay. This reveals a gap between perception and pricing in which both agents and sellers may overstate the financial impact of presentation, while buyers ultimately behave more rationally at the point of transaction. A further limitation in the existing literature is the implicit assumption that buyers intend to occupy the properties they purchase. In practice, a substantial share of residential transactions are driven by investor demand. Approximately 20% of Australian taxpayers report an interest in a rental property (Australian Taxation Office, 2024), and residential property investment is a similarly significant asset class in comparable jurisdictions: the private rented sector accommodates around a fifth of households in England (Ministry of Housing, Communities and Local Government, 2024) and Statistics Canada (2023) reports that investors own between 14 and 26% of houses across the provinces examined. Despite this, the existing literature on AVMs and signalling in housing markets has, to our knowledge, not engaged with the investor segment as a distinct buyer type, leaving open the question of whether the same informational dynamics apply.

Signals also arise beyond individual listings, reflecting broader market conditions and shared expectations. Heinig and Nanda (2018) show that sentiment, measured through Google search activity, property market reports and macroeconomic proxies, plays a significant role in shaping short-term demand and yield movements. Rising optimism pushes prices up and yields down, while negative sentiment has the opposite effect. These findings underline the way in which the mood of the market itself becomes a signal, and one that can be tracked in near real time. Lützkendorf and Speer (2005) highlight the role of quality indicators such as energy performance certificates, green building labels and other certifications. These are observable and costly to fake, which makes them credible signals that consistently command price premiums and reduce uncertainty. Unlike stylistic features such as staging or description, they provide institutionalised assurances of building quality and performance.

Together, these studies show that qualitative signals in housing markets range from the fragile to the durable. Words, photos and staging influence perception, sentiment measures capture collective expectations and certifications provide trusted benchmarks. Each helps buyers navigate uncertainty, although their effectiveness and persistence vary. Just as with valuation signals, their role may shift as new technologies change what can be produced easily and what remains costly to imitate.

2.3 Generative AI and language production

The emergence of generative AI introduces a shift in the production of language within housing markets. Where once property descriptions were crafted by agents with creativity and persuasion, the task is now increasingly mediated by large language models. This shift has dual implications: automation of labour and a potential erosion of the distinctiveness of marketing signals.

Kandipati (2025), through a systematic review of 72 studies, documents how natural language processing and generative models are increasingly applied across the real estate sector. Agents are utilising generative AI to automate listing descriptions, deliver personalised recommendations and deploy conversational chatbots that supplement traditional workflows (Kriegbaum *et al.*, 2024). These tools improve efficiency and scalability, but Kandipati also highlights an emerging authenticity challenge: when agents draw on similar generative models, marketing language risks becoming homogenised, weakening its role as a distinctive signal of quality. A parallel concern has emerged in higher education, where the widespread use of generative AI has prompted debate about the credibility of written assessments as signals of student ability, with some institutions reconsidering assessment formats to preserve signal value (Alexander and Belloni, 2024; Perkins, 2023). The implication for property markets is similar: as textual signals become cheaper to produce, their informational content for buyers and valuation models may erode.

Like agents, agencies and platforms are increasingly adopting generative AI tools across a range of real estate functions. Ma and Huang (2023) note that platforms such as Redfin and Zillow have incorporated ChatGPT-based plugins, enabling conversational property search and automated assistance. Beyond text-based tools, Xiong *et al.* (2024) find that virtual reality listings attracted significantly more physical inspections and reduced time on market, suggesting that richer digital marketing tools more broadly accelerate buyer decision-making and reduce information asymmetry between buyers and sellers. Agents report substantial time savings from these tools, with applications extending beyond listing descriptions to customer service, legal queries, valuation support and even staging advice. However, this expansion is not without risk. Generative models trained on historical market data may inherit and reproduce embedded biases, including those linked to discriminatory practices such as redlining, reinforcing inequality if left unchecked. For this reason, Ma and Huang emphasise the importance of transparency and ongoing human oversight, framing generative AI not as a substitute for professional judgement but as a decision-support tool.

At the same time, emerging evidence suggests that the distributional effects of generative AI are context-dependent. In the specific setting of housing price generation, generative models have been shown to reduce observed price discrimination relative to human-generated prices, indicating that automation can, in some cases, mitigate rather than amplify bias (Tanlamai *et al.*, 2024). This nuance highlights that the implications of AI adoption depend not only on the technology itself but also on the domain and task to which it is applied.

Beyond equity considerations, empirical evidence shows that generative tools do not simply replicate human language but systematically reshape it. Sandler *et al.* (2024) demonstrate that ChatGPT-generated dialogues exhibit lower variance across linguistic categories than human conversations, producing more predictable and uniform patterns. While AI outputs scored higher on social behaviours, attentional focus and analytical style, they underperformed on measures of authenticity. Similarly, Arnold *et al.* (2020) show that predictive text nudges authors towards repetitive phrase use, reducing lexical diversity and increasing similarity across writers. Together, these findings point to a structural tendency of AI-mediated text production to converge on a narrower linguistic range, optimised for general acceptability but at the expense of originality. In the context of property marketing, this tendency may work against the creative or distinctive descriptions traditionally produced by human agents.

On one hand, generative AI enhances efficiency, scalability and responsiveness in real estate marketing, freeing agents from repetitive labour. But on the other hand, it risks diluting the informational value of language as a market signal. If descriptions become standardised outputs of widely available tools, their capacity to differentiate properties diminishes, potentially undermining their role in valuation models and buyer decision-making. As with earlier signals, from photos to staging, the credibility and distinctiveness of AI-generated text may erode once it becomes ubiquitous. The challenge for real estate markets lies in balancing automation with the preservation of authentic, costly-to-fake signals that continue to convey meaningful information in an increasingly AI-mediated environment.

2.4 AI detection and computational text analysis

The widespread adoption of generative AI has created a demand for reliable methods to detect machine-authored text. Detection tools are typically fine-tuned or task-specific language models designed to identify stylistic regularities associated with AI-generated writing. Much of the policy and education debate has focused on the risks of false positives in high-stakes contexts (Dalalah and Dalalah, 2023). Their analysis shows that detection scores for human- and AI-generated text can overlap substantially, making both type I and type II errors plausible and that literature-review sections are more likely to be flagged as AI-like than abstracts. These findings highlight the contextual sensitivity and inherent uncertainty of detection outputs.

At the same time, recent research positions AI detectors as valuable research instruments rather than enforcement tools (Dugan *et al.*, 2024). This distinction is critical: whereas regulators and educators may require near-perfect reliability, social scientists can leverage probabilistic detection scores to map trends, construct empirical features and generate insight into evolving writing practices even when classification is imperfect.

A central development in this space is the introduction of RAID, a shared benchmark designed by Dugan *et al.* (2024) to evaluate detectors under diverse and adversarial conditions. RAID systematically subjects models to synonym swaps, homoglyph substitutions, alternative decoding strategies and domain shifts (e.g. Reddit posts, reviews). Its contribution is to standardise evaluation and foreground robustness, especially at low false positive rates, which is a critical requirement for real-world deployment. Within this benchmark, researchers used DeskLib (an open-source Hugging Face release), which performed competitively relative to peers such as GPTZero and Binoculars, validating its use as a research instrument. RAID thus provides the framework through which comparative reliability can be understood.

Beyond benchmarking, other strands of research focus on improving specific aspects of detection. Jung *et al.* (2025) highlight the problem of subgroup bias, where short or stylistically distinct human texts are disproportionately misclassified. Their *FairOPT* framework adapts classification thresholds for different subgroups, reducing balanced error rate disparities by more than 12% across detectors. Although DeskLib saw only modest improvements, the study demonstrates how fairness-aware optimisation can sharpen detection without undermining accuracy. Additionally, studies in education test detectors against practical challenges. Weber-Wulff *et al.* (2023) evaluated 14 systems across six document sets, finding Turnitin and Compilatio to lead in raw accuracy, while DeskLib consistently performed in the upper tier at around 74%. Crucially, all tools were vulnerable to obfuscation strategies such as paraphrasing or translation. Elkhatat *et al.* (2023) reached similar conclusions: detectors are inconsistent, unstable under repeated tests and not suitable as decisive evidence in misconduct cases. Yet both studies acknowledge that detectors generate reliable differentiation between broad classes of text, providing useable signals for research.

Taken together, this literature suggests that AI text detection remains imperfect but informative, particularly when used to analyse patterns and distributions rather than to make binary judgements about authorship.

2.5 Research gap and contribution

The literature demonstrates moving from traditional valuations towards AVMs and, more recently, to the integration of textual and qualitative signals. While these advances underscore the growing role of data-driven valuation, several gaps remain.

First, existing research has largely focused on structural variables (e.g. amenities, neighbourhood characteristics, energy labels) or relatively static qualitative cues (e.g. staging, word positivity). Far less attention has been paid to the dynamic role of language itself as a signal, where its informational content evolves over time in response to market conditions and technological change. Listing descriptions are not fixed attributes: they are continuously rewritten, adapted to prevailing norms and increasingly shaped by automated tools. While prior studies show that descriptive language influences buyer behaviour, they pre-date the widespread adoption of AI-assisted writing and therefore do not capture how the scale, speed and uniformity of language generation may be weakening the distinctiveness and credibility of textual signals.

Second, although signalling theory has been applied to housing markets, its application to AI-generated text is underdeveloped. The theory emphasises that signals must be observable and costly to fake, yet generative tools radically reduce the cost of producing persuasive language. On the other hand, it is also currently unproven that generative AI tools produce the same level of persuasion, or that the product of such tools will remain viable or relevant for a meaningful forward period of time. Historically important qualitative signals such as listing text risks losing its discriminatory power precisely because of technological diffusion. Few studies have empirically tested whether language remains an effective signal once it becomes cheap and homogeneous.

Third, while the emergence of AI detection tools has generated extensive policy debate, their use as research instruments remains underexplored. Existing studies primarily evaluate detection systems in adversarial or educational settings, with little attention to how probabilistic measures of AI-generated text can be used to analyse markets, track trends, or construct empirical features. This creates a missed opportunity to treat detection not as a gatekeeping mechanism but as an analytical lens into changing market practices. This paper addresses these gaps in three ways.

First, it reconceptualises listing text as a qualitative signal whose informational value evolves under technological change, moving beyond the question of whether descriptions influence outcomes to examine whether they retain signalling power in an environment shaped by generative AI. Second, it provides empirical evidence linking detection-based measures of AI-likeness to transaction outcomes, offering the first quantitative assessment of whether language homogenisation diminishes market performance in terms of price and time-on-market. Third, it advances the methodological integration of computational and qualitative approaches by framing AI detection as an analogue to qualitative coding, demonstrating how machine learning tools can be used to interrogate meaning and signal erosion rather than replace interpretive analysis.

Together, these contributions show how generative AI is reshaping not only marketing practices but also the informational ecology of housing markets, raising broader questions about the sustainability of qualitative cues in an era of automated language production.

3. Data and methodology

The methodological design of this study is structured to directly address the gaps identified in the preceding literature. Prior research establishes that AVMs and text-based signals increasingly shape real estate analysis (Kok *et al.*, 2017; Topraklı, 2025), while recent work shows that generative AI is transforming the production of marketing language (Kandipati, 2025; Sandler *et al.*, 2024).

However, existing studies do not test whether these changes alter the informational value of listing text in market outcomes. To address this gap, the study integrates structured property attributes with unstructured listing descriptions and examines how variation in linguistic characteristics relates to transaction performance (Baur *et al.*, 2023). Specifically, the empirical framework tests whether measures of AI-likeness in listing language are associated with realised prices and time-on-market, triggering differences in market efficiency, a dimension widely used to assess liquidity despite its inherent complexity (Cajias and Freudenreich, 2024). This approach operationalises listing text as a signal whose credibility may vary with the extent of automated language production.

The methodology combines computational text analysis with AI detection tools to construct probabilistic indicators of generative language use. Consistent with the literature, these tools are not treated as definitive classifiers but as continuous features that capture convergence toward AI-generated writing styles (Dugan *et al.*, 2024; Jung *et al.*, 2025). By embedding these indicators within standard valuation and market performance models, the study tests whether language homogenisation corresponds to weaker signalling effects, thereby providing an empirical assessment of the risks and opportunities associated with AI-mediated marketing.

Building on this framework, the empirical analysis is guided by four research questions.

RQ1. Does the degree of AI-likeness in listing language relate systematically to market outcomes, as measured by transaction prices and time-on-market, after controlling for standard property characteristics?

RQ2. Does greater linguistic convergence toward AI-generated style correspond to weaker signalling effects, consistent with theories of signal dilution?

RQ3. Can probabilistic AI detection scores be used as stable empirical features for analysing changes in market communication practices rather than as binary indicators of authorship?

RQ4. Is there any evidence to suggest that these outcomes differ between properties that appeal to investor buyers rather than purchasers intending to owner occupy?

3.1 Data source and sample

The dataset was drawn from RP Data and covers residential property sales in South Australia between 1 June 2023 and 1 January 2025. This date was also chosen due to the nature of earlier versions of language models, which were easier to detect AI linguistic patterns.

We began with approximately 10,000 listings, but after filtering out records with missing or empty descriptions, failed text extractions, absent sale prices, or missing days on market

(DOM), the final analytic sample comprised 6,376 complete cases. A subsequent data collection phase scraped rental platform listings to construct an investor flag variable, yielding 5,368 useable cases after excluding approximately 248 rows with scrape errors. The full 6,376 case sample is used for the linguistic and pricing analyses, while the investor flag analysis draws on this slightly smaller subset. For each listing, we retained both the structured property attributes (price, bedrooms, bathrooms, car spaces, land and floor size, year built) and the unstructured textual description. A third outcome variable, investor_flag, was constructed by scraping a major rental listings platform to identify whether each sold property appeared as a rental listing within 12 months of sale, consistent with investor purchase behaviour. A positive flag is treated as strong evidence of investor acquisition; a negative flag does not confirm owner-occupier purchase, as properties rented through other platforms or privately would be misclassified.

3.2 AI detection and text processing

To identify whether listing descriptions bore hallmarks of generative AI, we used DeskLib's AI text detector (desklib/ai-text-detector-v1.01), which is a fine-tuned variant of microsoft/deberta-v3-large released on Hugging Face. DeskLib was benchmarked on the RAID framework (Dugan *et al.*, 2024), which systematically evaluates detectors under synonym swaps, homoglyph substitutions, decoding strategies and domain shifts. Its demonstrated robustness makes it suitable for research purposes, aligning with the position that AI detection is best employed as an analytic resource rather than as conclusive evidence (Elkhatat *et al.*, 2023).

The model returns a probability of AI authorship for each description. We applied a 0.5 threshold to create a binary classification (AI vs non-AI) while also retaining the continuous score as an auxiliary measure of "AI-likeness". Following recent methodological guidance (Jung *et al.*, 2025), we interpret these scores probabilistically, recognising both the potential for subgroup bias and the value of such features in mapping linguistic convergence.

Beyond classification, we extracted a suite of linguistic and stylistic features from each description, consistent with prior text analysis in real estate and behavioural economics (Luchtenberg *et al.*, 2019; Arnold *et al.*, 2020). These included:

(1) *Length and complexity metrics:* character count, word count, sentence count, average sentence length, type-token ratio.

(2) *Stylistic markers:* punctuation density, frequency of exclamation marks, ratio of uppercase characters.

(3) *Readability indices:* Flesch Reading Ease and Flesch-Kincaid grade scores.

(4) *Persuasive vocabulary:* binary indicators for common marketing adjectives (e.g. *stunning*, *spacious*, *modern*).

We further applied n-gram chi-square tests to identify words disproportionately associated with AI vs non-AI text, Latent Dirichlet Allocation (LDA) to extract latent topics and term frequency-inverse document frequency (TF-IDF) with t-SNE visualisation to illustrate clustering patterns. These approaches mirror the use of computational text analysis in housing and marketing studies to differentiate authentic from imitative signals (Connelly *et al.*, 2011; Sandler *et al.*, 2024).

3.3 Empirical strategy

The empirical analysis proceeded in two stages. First, we compared AI and non-AI texts directly using Mann–Whitney *U* tests on numeric linguistic features, reporting Cliff's delta as a non-parametric effect size. This follows the literature's emphasis on identifying systematic linguistic convergence (Sandler *et al.*, 2024; Arnold *et al.*, 2020). Vocabulary distributions and topic models were also examined to assess semantic divergence across AI and non-AI groups.

Second, we investigated whether the presence of AI-like text was associated with different property market outcomes compared to property listing descriptions that were more obviously generated by human agents. We decided to focus the analysis on three dependent variables: log-transformed DOM, investor_flag and log sale price. The DOM and sale price outcomes were estimated using ordinary least squares (OLS) regressions of the form:

$$\mathrm{Yi} = \alpha + \beta\,\mathrm{AIi} + \gamma \mathrm{Xi} + \varepsilon \mathrm{i}$$

where Y denotes the outcome of interest and X includes property characteristics. The investor_ flag outcome was estimated using logistic regression with the same set of controls, with results reported as odds ratios.

Explanatory variables included bedrooms, bathrooms, car spaces, land size, floor size and year built. Property type was excluded due to sparsity across categories. Fine-grained geographic fixed effects (suburb or postcode) are not feasible given the large number of locations relative to the sample size, with many areas represented by very few observations. This model specification is consistent with prior AVM studies that emphasise parsimonious inclusion of structural attributes (Kok *et al.*, 2017; Topraklı, 2025). Coefficients were interpreted as semi-elasticities and conventional OLS standard errors were reported. Diagnostics included checks for multicollinearity and residual distributions. In line with the broader literature (Ma and Huang, 2023; Heinig and Nanda, 2018), the goal was not perfect prediction, but to assess whether AI-like linguistic signals bear systematic relationships with market outcomes.

4. Results

The empirical analysis proceeded in three stages: descriptive statistics to characterise the sample and establish baseline comparability, linguistic analysis to examine systematic differences between AI and non-AI listing descriptions and regression models to evaluate whether AI-generated language influenced property market outcomes. This structure mirrors the methodological design, moving from general distributional patterns through to specific hypothesis tests.

4.1 Descriptive statistics

The analytic dataset comprised 6,376 property listings in South Australia between June 2023 and January 2025. This was derived from an initial pool of approximately 10,000 listings, with exclusions for missing text, incomplete transaction records, or absent outcomes such as price or DOM. The retained listings therefore represent a relatively clean and comprehensive sample of the local residential sales market.

Using the desklib model from Hugging Face, the fine-tuned AI detection model classified approximately 52% of listing descriptions as likely to have been AI-generated. This proportion suggests that AI adoption in marketing text is already non-trivial within the study window, offering sufficient variation to allow for comparative analysis. Importantly, the two groups (AI vs non-AI) did not differ systematically in their structural property attributes. Average lot size, floor area, bedroom count and age of dwelling were broadly comparable across groups, indicating that differences observed in outcomes are unlikely to be driven by underlying property characteristics alone. At the descriptive level, AI-listed properties showed a slightly higher investor rate (10.6 versus 8.5% for non-AI listings), sold faster (median 29 days versus 31 days) and at a marginally lower price (median $715,000 versus $750,000), though the practical significance of these differences is assessed through the regression models reported in Section 4.3.

Nevertheless, the descriptive comparison revealed a clear temporal pattern. The share of listings flagged as AI-generated increased steadily month-on-month across the study period, suggesting a diffusion process rather than one-off experimentation. Early in the sample window, only a small fraction of descriptions exhibited AI-like characteristics, but adoption accelerated from late 2023 onwards, with some months showing double-digit growth relative to the preceding quarter (Figure 1). This upward trajectory highlights how generative tools are becoming normalised within marketing practice, underscoring the need to examine not only cross-sectional differences but also the implications of a shifting baseline over time.

4.2 Linguistic features

The most immediate distinction between AI- and non-AI-generated text lay in length and structure. AI listings were consistently more verbose, with significantly higher average word counts (401 vs 375, $p < 0.001$, $\delta = 0.113$, negligible effect), sentence counts (19.0 vs 16.4, $p < 0.001$, $\delta = 0.183$, small effect) and character counts (2,592 vs 2,380, $p < 0.001$, $\delta = 0.137$, negligible effect). These findings indicate that AI outputs are longer overall, though the magnitude of difference varies across measures. By contrast, human-authored descriptions often exhibited greater variability, ranging from terse bullet-like entries to highly detailed narratives. Sentence length followed the reverse pattern: non-AI texts contained significantly longer sentences on average (25.7 vs 22.4, $p < 0.001$, $\delta = -0.146$, negligible effect), underscoring a tendency for humans to combine ideas into fewer but more information-dense units.

Figure 2 further illustrates this pattern by showing a tighter distribution of average sentence length for AI-generated listings, indicating greater consistency in sentence construction. By contrast, non-AI texts exhibit a wider dispersion and heavier upper tail, reflecting greater stylistic variation in how information is packaged into sentences. Stylistically, AI-generated text also displayed higher uniformity in punctuation use, with punctuation counts significantly exceeding those of non-AI texts (45.6 vs 37.2, $p < 0.001$, $\delta = 0.242$, small effect). This suggests that AI descriptions follow a more standardised punctuation rhythm, while human-authored listings vary more widely in their use of commas, dashes, or colons.

A more striking divergence was observed in the use of emphatic phrases. Nearly half of AI-generated listings (≈46%) contained at least one emphatic phrase such as *"DON'T MISS"* or *"PRIME LOCATION"*, compared to just 23% of non-AI texts, a gap of 23% points. The phrase *"DON'T MISS"* appeared in nearly one-third of AI listings (≈29%) alone, signalling a strong reliance on formulaic persuasion. By contrast, non-AI texts employed emphatic language less frequently and with greater variation, reflecting more individualised authorial choices. Together, these results suggest that AI listings are not less emphatic, but emphatic in a highly standardised and repetitive way.

Readability indicators further differentiated the groups, though with smaller effect sizes. On the Flesch Reading Ease scale, AI-generated descriptions averaged slightly lower (40.7 vs 41.4, $p < 0.001$, $\delta = -0.102$, negligible effect), clustering around mid-range values that indicate accessible but somewhat formalised text. On the Flesch–Kincaid Grade scale, no significant difference was observed ($p = 0.122$, $\delta = -0.024$, negligible effect). Non-AI texts exhibited wider dispersion across both indices, with some listings highly accessible and others more complex, reflecting the idiosyncratic nature of human writing. This aligns with signalling theory: the cost of producing varied and personalised language is higher for humans, whereas AI tends to converge toward a standardised middle. Table 1 summarises the Mann–Whitney *U* test results across all linguistic features.

Lexical analysis provided further evidence of divergence between the two groups. Chi-square tests on unigram and bigram frequencies showed that AI-authored listings disproportionately used generic adjectives (*"beautiful home," "ideal location"*), whereas non-AI texts included more context-specific references (*"walk to Glenelg beach," "minutes from Rundle Mall"*) (Figure 3). This supports the hypothesis that generative systems, trained on broad corpora, tend to produce language optimised for general persuasiveness but less anchored in place-specific detail.

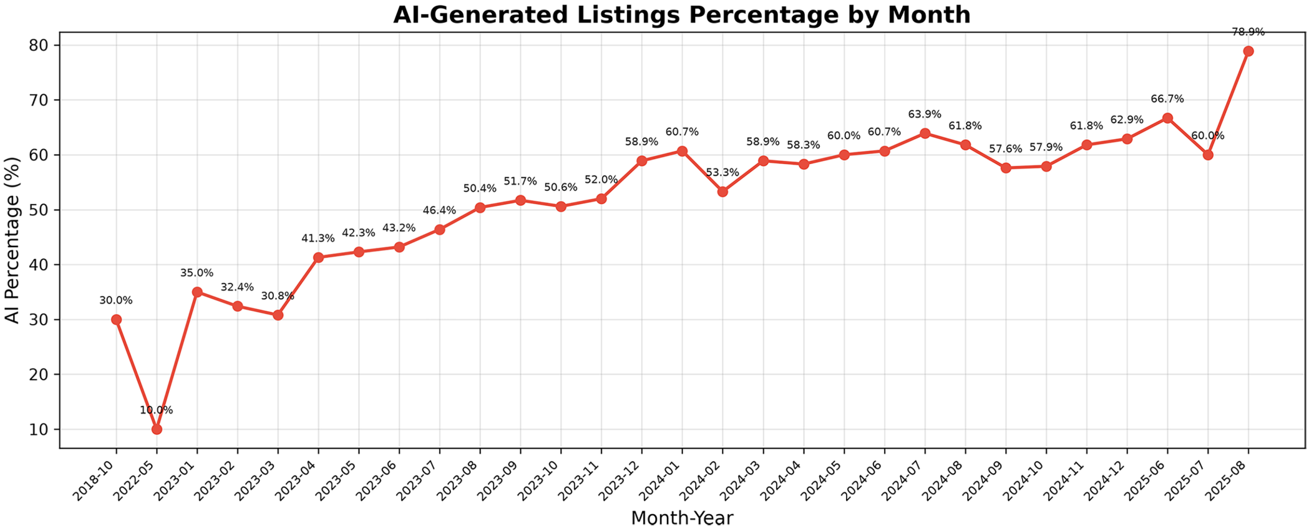

Figure 1. AI-detected listings month on month

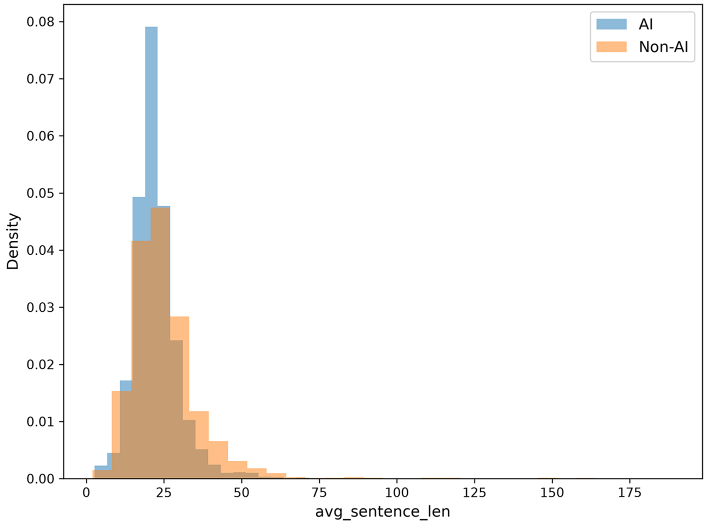

Figure 2. Distribution of avg_sentence_len by label

Table 1. Summarising Mann–Whitney *U*-test results for linguistic features with Cliff's delta values

Feature	*U*-statistic	*p*	Cliff's delta	AI mean	Non-AI mean
Punctuation count	4948693.0	$5.85 \times 10^{-56***}$	0.242	45.570	37.199
Sentence count	4714464.0	$7.19 \times 10^{-33***}$	0.183	18.987	16.356
Uppercase character ratio	3373993.5	$2.41 \times 10^{-23***}$	−0.153	0.033	0.043
Type-token ratio	3382952.0	$1.04 \times 10^{-22***}$	−0.151	0.624	0.640
Average sentence length	3402999.5	$2.53 \times 10^{-21***}$	−0.146	22.400	25.704
Character count	4529449.5	$4.98 \times 10^{-19***}$	0.137	2592.177	2379.091
Word count	4434613.5	$1.78 \times 10^{-13***}$	0.113	400.749	374.637
Flesch Reading Ease	3576121.5	$2.86 \times 10^{-11***}$	−0.102	40.674	41.357
Flesch-Kincaid grade	3888791.5	0.122	−0.024	13.151	13.876
AI detection score (DeskLib)	2549323.0	0.539	−0.011	0.523	0.533
Exclamation count	3964664.5	0.734	−0.005	0.812	0.935

Note(s): Asterisks denote statistical significance: $*p < 0.05$, $**p < 0.01$, $***p < 0.001$. No symbol indicates $p \geq 0.05$

One potential concern in interpreting lexical frequency differences is that AI-generated listings are, on average, longer, which mechanically increases the probability that any given word or phrase appears. However, this length effect alone cannot explain the observed patterns. If longer text were the primary driver, AI listings would be expected to contain higher frequencies of both generic and context-specific terms. Instead, the results show a selective overrepresentation of generic adjectives and aspirational phrases in AI-authored text, alongside an underrepresentation of place-specific references that are common in human-authored listings. This suggests that the differences reflect stylistic convergence rather than simple verbosity. Nonetheless, future work could further isolate these effects by normalising lexical frequencies by document length or by modelling relative term shares rather than absolute counts.

Topic modelling with LDA identified distinct thematic structures. AI-generated text clustered more heavily around aspirational lifestyle themes, grouping words related to

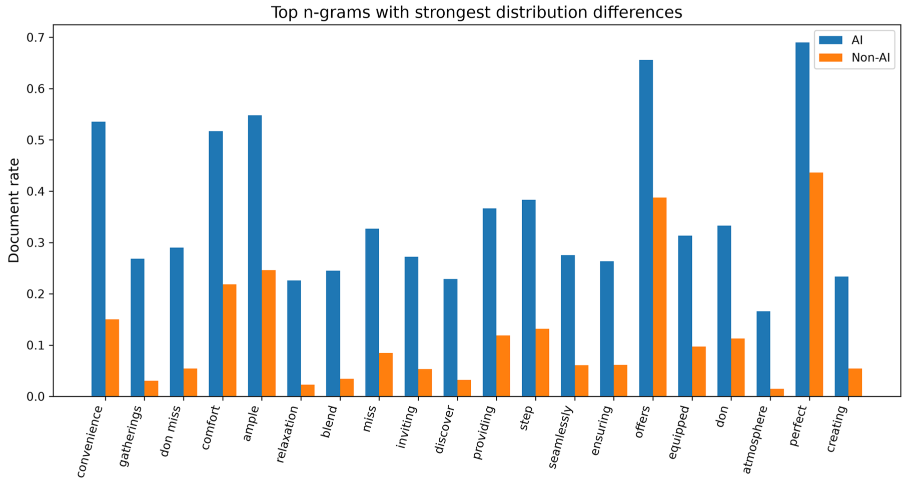

Figure 3. N-gram differences between AI vs. Non-AI

comfort, design and modernity. Non-AI text produced topics with stronger emphasis on location, neighbourhood amenities and agent-specific narratives. This divergence underscores how AI homogenisation may reduce the discriminatory power of listing text: if all descriptions emphasise generic lifestyle appeals, their ability to credibly signal value diminishes. These are also larger giveaways into how the author adds knowledge to the listing that the AI model might not be trained on and be able to provide within the listing text.

Visualisation using TF-IDF embeddings and t-SNE projections illustrated these findings graphically (Figure 4). While AI and non-AI listings occupied overlapping space in the

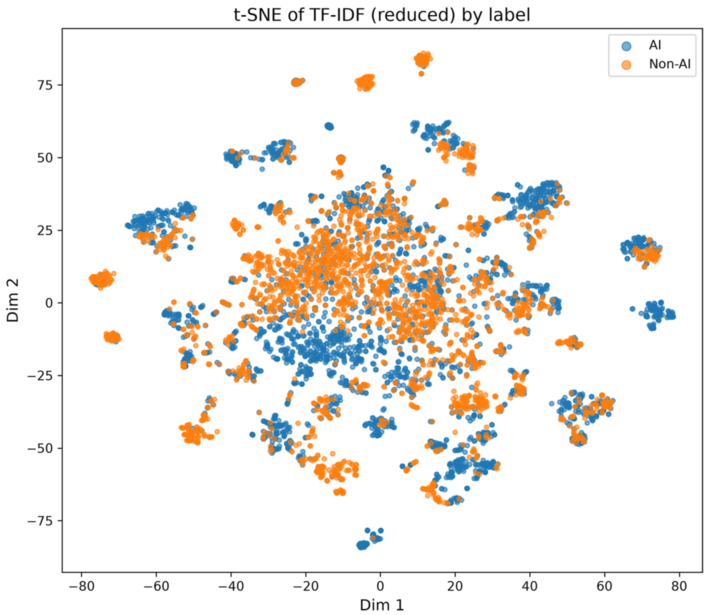

Figure 4. t-SNE scatter plot of document embeddings by AI vs non-AI label

semantic map, there was visible clustering that aligned with detector classifications. The degree of separation was sufficient to suggest that stylistic convergence is occurring, but not so extreme as to make AI descriptions indistinguishable from human text in practice.

Together, these analyses show that AI-generated listings are systematically distinguishable from human-authored ones across multiple dimensions of text. They are longer, more structurally uniform and more likely to rely on emphatic phrases, while non-AI texts remain more variable and context-specific. Vocabulary and topic modelling further reinforce this distinction, with AI tending toward generic lifestyle appeals and humans embedding more place-based or idiosyncratic details. These patterns suggest a clear divergence in how language is produced across groups, setting the stage for examining whether such differences carry measurable implications for market outcomes.

4.3 Regression analysis

The coefficient for AI-generated listings on log sale price was negative and statistically significant but substantively very small ($\beta = -0.019$, 95% confidence interval (CI): -0.038 to -0.001, $p = 0.043$), suggesting a difference of approximately 2% in sale price associated with AI-authored descriptions after controlling for property characteristics (Table 2). Given the effect size, this finding should be treated cautiously. Among the significant control variables, bathrooms and floor size were positively associated with price, while year built carried a small negative coefficient, consistent with the premium often attached to established dwellings in the South Australian market.

The effect of AI-generated descriptions on time to market was more consistent. The coefficient for AI-generated listings on log DOM was negative and statistically significant ($\beta = -0.095$, 95% CI: -0.141 to -0.048, $p < 0.001$), indicating that properties with AI-authored text sold approximately 9% faster than comparable homes with agent-written descriptions (Table 3). One plausible explanation is that the greater structural uniformity and generic phrasing of AI-generated descriptions reduce cognitive processing costs for buyers, thereby facilitating faster comparison and decision-making.

It is worth noting that the R^2 of approximately 0.012 for the days-on-market regression indicates low explanatory power. This is consistent with prior research showing that time on market is inherently difficult to model and remains sensitive to a wide range of interacting and often unobserved factors (Cajias and Freudenreich, 2024). Factors such as pricing strategy, seller urgency, micro-neighbourhood demand and seasonal timing likely play a much larger role. Even so, the consistently negative and statistically significant association with AI-authored text suggests that, on average, these descriptions are associated with a modest improvement in market efficiency as measured by time on market.

The logistic regression for investor acquisition found that AI-listed properties had 21% higher odds of appearing on a rental platform within 12 months of sale (OR = 1.21, 95% CI: 1.01–1.46, $p = 0.044$) (Table 4). While statistically significant, this result sits at the margins of

Table 2. OLS regression results for log sale price

Variable	Coef.	Std. Err.	t	p	95% CI
AI listing	*−0.0194*	*0.0096*	*−2.027*	*0.043**	*[−0.0381, −0.0006]*
Bedrooms	−0.0131	0.0084	−1.555	0.120	[−0.0296, 0.0034]
Bathrooms	0.1383	0.0106	13.081	<0.001***	[0.1176, 0.1590]
Car spaces	0.0076	0.0044	1.733	0.083	[−0.0010, 0.0162]
Land size (m^2)	−0.0000	0.0000	−0.390	0.696	[−0.0001, 0.0000]
Floor size (m^2)	0.0037	0.0001	34.976	<0.001***	[0.0035, 0.0039]
Year built	−0.0060	0.0002	−39.102	<0.001***	[−0.0063, −0.0057]
Constant	24.6302	0.3011	81.798	<0.001***	[24.0399, 25.2205]

Note(s): Asterisks denote statistical significance: *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$. No symbol indicates $p \geq 0.05$

Table 3. OLS regression results for log DOM

Variable	Coef.	Std. Err.	t	p	95% CI
AI listing	*−0.0949*	*0.0237*	*−3.996*	*<0.001******	*[−0.1415, −0.0484]*
Bedrooms	0.0539	0.0209	2.575	0.010*	[0.0129, 0.0949]
Bathrooms	0.0437	0.0263	1.663	0.096	[−0.0078, 0.0951]
Car spaces	0.0162	0.0109	1.485	0.138	[−0.0052, 0.0375]
Land size (m^2)	−0.0001	0.0001	−2.023	0.043*	[−0.0003, −0.0000]
Floor size (m^2)	−0.0001	0.0003	−0.466	0.641	[−0.0006, 0.0004]
Year built	0.0013	0.0004	3.398	<0.001***	[0.0005, 0.0020]
Constant	0.7748	0.7479	1.036	0.300	[−0.6914, 2.2410]

Note(s): Asterisks denote statistical significance: *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$. No symbol indicates $p \geq 0.05$

Table 4. Logistic regression results for investor acquisition within twelve months of sale

Variable	Coef.	Std. Err.	z	p	Odds ratio	95% CI (OR)
AI listing	*0.1909*	*0.0949*	*2.012*	*0.044**	*1.210*	*[1.005, 1.458]*
Bedrooms	0.0134	0.0902	0.149	0.882	1.014	[0.849, 1.209]
Bathrooms	−0.1911	0.1111	−1.720	0.085	0.826	[0.664, 1.027]
Car spaces	−0.1775	0.0524	−3.389	<0.001***	0.837	[0.756, 0.928]
Land size (m^2)	0.0004	0.0003	1.325	0.185	1.000	[1.000, 1.001]
Floor size (m^2)	−0.0061	0.0015	−4.177	<0.001***	0.994	[0.991, 0.997]
Year built	0.0009	0.0016	0.560	0.575	1.001	[0.998, 1.004]
Constant	−2.7711	3.1043	−0.893	0.372	0.063	[0.000, 27.473]

Note(s): $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

conventional thresholds and should be interpreted with care given the conservative nature of the investor flag variable and the multiple testing context. Nonetheless, the direction of the finding is consistent with the possibility that the standardised, efficiency-oriented character of AI-generated listings may be particularly well suited to investor buyers who prioritise transactional clarity over individualised property narratives.

Taken together, the three models point to a consistent pattern. AI-listed properties sell modestly faster, attract slightly more investor buyers and sell at a marginally lower price, though the sale price and investor effects are small and should not be overstated. The DOM finding is the most robust across both datasets and model specifications. The evidence points to a modest efficiency advantage associated with AI-generated listing language rather than any systematic effect on property value.

5. Discussion

This study advances our understanding of how language functions as a signal in housing markets under conditions of technological change. Prior work has positioned listing text as a persuasive cue capable of shaping buyer perceptions and intentions (Luchtenberg *et al.*, 2019). Our findings suggest a shift: while AI-authored descriptions were systematically longer, more formulaic and more reliant on emphatic phrases, these features were associated with only a marginal and cautiously interpreted reduction in sale price and no meaningful premium or discount in practical terms. In signalling terms, the reduced cost and widespread availability of generative AI appear to weaken the credibility of listing text as a costly signal of quality.

Instead, the primary effect observed was on market efficiency. Properties with AI-generated descriptions sold notably faster, implying that the uniformity and clarity of machine-authored text may reduce frictions in buyer search and decision-making. Rather than functioning as a differentiator of quality, listing language appears to be evolving into a coordination device that streamlines transactions. This reframes the role of text in housing markets: its signalling power may be diminishing, while its instrumental value in reducing transaction costs is increasing. Theoretically, these findings contribute in three ways.

First, the findings update signalling theory for the AI era by showing that when the cost of producing persuasive signals falls, their credibility and distinctiveness diminish, even as new forms of efficiency may emerge. Second, the results extend real estate pricing research by demonstrating that while price effects of textual variation are marginal and should not be overstated, listing language continues to influence market performance through time-on-market. Third, the study contributes methodologically by treating AI-likeness as a probabilistic research feature rather than a binary classification. Framing detection tools as analytical instruments rather than enforcement mechanisms enables scholars to trace linguistic convergence and link it meaningfully to market outcomes.

A further finding with particular relevance for investment markets is the association between AI-generated listings and investor purchase. Properties listed with AI-authored descriptions had 21% higher odds of appearing on a rental platform within 12 months of sale than those with manually written descriptions. This may suggest that the standardised, transactionally efficient character of AI language may resonate with investor buyers who prioritise clarity over narrative distinctiveness. It has a certain intuitive appeal given that property investors are likely to focus on key metrics such as size, bedrooms, property condition, likely rental value and ease of finding a suitable tenant. By contrast, purchasers intending to owner-occupy are more likely to be swayed by aesthetic considerations and liveability – signalling or signposting more easily conveyed in a manually written description. This has implications beyond the individual transaction. If AI-mediated listings systematically attract a higher share of investor buyers, the composition of demand in residential markets may shift in ways that are not immediately visible in price or time-on-market data alone. For investment portfolio managers and analysts, this raises questions about whether listing language is a useful signal of likely buyer type and whether the diffusion of AI tools is quietly reshaping the boundary between owner-occupier and investment-grade residential stock.

Commercial property marketing may require a more specialised and targeted process than residential property marketing. Residential marketing often focuses on lifestyle, location, comparable sales and broad public exposure. Commercial marketing, by contrast, requires a deeper analysis of income-producing and legal characteristics, including lease terms, tenant covenant strength, income yield, expectations for rental growth, regulatory compliance, market depth and liquidity. Tenant covenant strength is relevant to property pricing and equivalent yields (Hutchison *et al.*, 2011). The definition of market value in the RICS and International Valuation Standards requires an estimate of the price achievable in a hypothetical arm's-length transaction following proper marketing (RICS, 2023; IVSC, 2025). The appropriate form and extent of marketing depend on the asset and market conditions. In commercial markets, segmentation by investor size and preferences for particular property characteristics may make a tailored campaign directed towards relevant investor groups more appropriate than broad public advertising (Cvijanović *et al.*, 2022). Consequently, the effects of GenAI-generated listing language observed in residential property should not be assumed to apply unchanged to commercial property, where purchasers are likely to place greater weight on verifiable income, lease and risk information.

A related implication concerns the differential value of human-authored text for distinctive or unusual properties. Our finding that non-AI descriptions embed more place-specific detail and idiosyncratic content mirrors emerging evidence from the valuation literature that AI tools struggle with properties where standard comparable evidence is limited, such as those with unusual aspect, significant views, or atypical configurations (RICS, 2023). Just as valuers must exercise greater judgement when automated models cannot adequately capture a property's unique characteristics, agents marketing distinctive dwellings may find that human-authored narratives retain a comparative advantage precisely because they can convey what generic AI language cannot. This suggests a potential segmentation in optimal listing strategy: AI-generated text may be well suited to standard stock where transactional efficiency is the priority, while human-authored descriptions may better serve properties where differentiation drives value.

For practitioners, the results suggest that generative AI tools may deliver modest but consistent efficiency gains without adversely affecting price outcomes. Agencies may benefit from adopting these tools to accelerate marketing workflows and reduce time on market, particularly in high-volume or resource-constrained contexts. However, the risk of over-reliance is evident. If listings increasingly converge on generic lifestyle appeals, differentiation between properties may weaken, undermining the role of marketing in attracting specific buyer segments. To preserve distinctiveness, agents may need to supplement AI-generated content with personalised, property-specific, or locally grounded details.

For platforms that utilise listing search tools, the diffusion of AI-authored text raises important questions about content homogeneity and user experience. If buyers increasingly encounter standardised language across listings, platforms may need to innovate in other dimensions of presentation, such as interactive features, visual search tools, or enhanced imagery, to preserve informational variety. Regulators may also need to consider the fairness implications of widespread AI-mediated marketing. Historical biases embedded in training data could be reproduced or amplified through homogenised outputs, reinforcing inequality in subtle but systematic ways. In this context, transparency around AI use and the maintenance of human oversight are likely to play an important role as generative tools become more deeply embedded in everyday marketing practice.

While the study provides novel insights, several limitations should be noted. The dataset is restricted to South Australian transactions during a period when earlier generations of language models were more readily identifiable by existing detection tools. As generative models evolve and their stylistic signatures converge further with human writing, classification accuracy may decline, potentially affecting replication in future periods.

More importantly, the validity of AI detection itself requires careful qualification. Detectors such as DeskLib, used here, have been evaluated through the RAID benchmark (Dugan *et al.*, 2024), which tests robustness against synonym swaps, homoglyph substitutions and domain shifts. These results support its use as a research instrument, but not as conclusive evidence of authorship. Studies in educational contexts (Weber-Wulff *et al.*, 2023; Elkhatat *et al.*, 2023) further demonstrate that detectors remain vulnerable to paraphrasing, translation and repeated trials, and can produce inconsistent classifications. Jung *et al.* (2025) also show that subgroup biases distort error rates for shorter or stylistically distinct texts. Taken together, this literature underscores that AI detection should be interpreted probabilistically, as a noisy indicator of linguistic convergence rather than a ground-truth classifier.

A further limitation concerns the potential confounding effect of agent technological sophistication. Agents who adopt AI writing tools may also be more technologically advanced in other aspects of their marketing, including photography quality, platform selection and listing timing. If so, the observed reduction in time on market may partly reflect these complementary practices rather than the linguistic features of the listing description alone. Without agent-level data on broader marketing behaviour, this cannot be ruled out empirically and the findings should be interpreted with this caveat in mind.

The regression models explain a relatively modest share of variation in DOM, a pattern that is well documented in the time-on-market and valuation literature. Prior research emphasises that marketing duration is highly sensitive to idiosyncratic and often unobserved factors, including market opacity, heterogeneity of assets, pricing strategy and local demand conditions, which limits the attainable explanatory power of even richly specified models (Tajani *et al.*, 2018). Within this context, the results indicate that listing language is one factor among many shaping market outcomes, and its effects should be interpreted as incremental rather than determinative.

Finally, this study is restricted to residential property transactions in South Australia. The dynamics identified here may not translate directly to commercial property markets, where information asymmetries, buyer sophistication and marketing conventions differ substantially. Lease structures, covenant strength and income-based valuation methods mean that listing language plays a different and arguably less central role in commercial transactions than in residential ones. Future research examining AI-mediated language in commercial property marketing would help determine whether the efficiency and homogenisation effects observed here are specific to the residential context or represent a broader market phenomenon.

Future research could extend this analysis both geographically and temporally to assess whether the observed efficiency effects persist across different market settings and under more advanced generations of language models. Incorporating richer outcome measures, such as buyer engagement indicators including click-through rates, enquiries, or open-home attendance, would help clarify whether homogenised language improves informational clarity or instead reduces distinctiveness from a demand-side perspective. Complementary qualitative interviews with agents and buyers could further illuminate how authenticity and credibility are perceived in AI-mediated listings, capturing dimensions that quantitative detection methods cannot fully observe. A further priority lies in methodological validation. Detection tools should be benchmarked not only through adversarial stress tests but also against labelled datasets of property listings to assess domain-specific classification accuracy. As models continue to evolve, newer detection frameworks and fairness-aware approaches, such as FairOPT (Jung *et al.*, 2025), should be evaluated to ensure robustness and subgroup validity.

6. Conclusion

This study examined how generative AI is reshaping the informational role of property listing descriptions in housing markets. Building on signalling theory and prior work on AVMs, it asked whether AI-mediated language retains its credibility as a qualitative signal and how its adoption relates to market outcomes.

The findings reveal a clear divergence between AI- and human-authored text. Machine-generated descriptions were longer, more formulaic and more reliant on emphatic phrasing, while human-authored listings embedded greater variability and context-specific detail. Yet these stylistic differences were associated with only a marginal reduction in sale price, with no meaningful practical effect on transaction values. AI-generated listings were associated with faster sales and with modestly higher odds of investor acquisition, suggesting that the standardised character of AI language may particularly resonate with investment buyers. In signalling terms, listing language appears to be shifting from a costly indicator of quality to a coordination device that reduces search frictions.

Theoretically, this contributes to an updated understanding of signalling in markets where persuasive cues become cheap to produce. For real estate transaction research, the results show that while textual signals do exhibit detectable price effects in this setting, they are marginal and should not be overstated. Methodologically, the use of AI detection as a probabilistic research instrument demonstrates how computational tools can extend qualitative inquiry, offering insight into linguistic convergence without claiming definitive attribution.

An alternative interpretation, however, warrants consideration. It is possible that agents selectively deploy generative AI for more standard or routine properties, where descriptive differentiation offers limited returns, while reserving bespoke, human-authored narratives for atypical or highly distinctive dwellings. Under this interpretation, shorter time-on-market for AI-authored listings may partly reflect underlying property homogeneity and clearer market comparability, rather than a causal effect of language alone. Conversely, non-AI listings may correspond to more idiosyncratic properties, where extended marketing periods reflect greater uncertainty but also the potential for price premium if differentiation is successful. While the marginal nature of the observed price effects suggests that such selection is not driving higher realised values in this sample, future research incorporating measures of property uniqueness or listing strategy would help disentangle these mechanisms more fully.

For practice, the results highlight both opportunity and caution. These findings are drawn from residential transactions and should not be assumed to apply directly to commercial property markets, where more heterogeneous investor groups, income-based valuation, lease structures and formal due diligence may alter both the role of listing language and the effects of GenAI-mediated marketing. Agencies may realise efficiency gains from adopting generative AI, but risk eroding differentiation if all listings converge on generic lifestyle appeals. Platforms and regulators must also anticipate the consequences of linguistic homogenisation, from reduced informational variety to the amplification of embedded biases. Future research should expand geographically and temporally, test newer detection frameworks and incorporate richer measures of buyer behaviour. A mixed-methods approach that combines computational detection with qualitative perspectives would deepen understanding of how authenticity and credibility are negotiated in AI-mediated markets.

In sum, generative AI is not fundamentally altering the pricing of housing, but it is reshaping the ecology of signals through which buyers and sellers navigate transactions. Its impact lies less in changing what properties are worth and more in how quickly and uniformly markets process the information surrounding them.

Note

1. Throughout this paper we use the term “valuation,” consistent with UK, European and Australian usage. In some jurisdictions, notably the United States, the equivalent term is “appraisal.” We retain “appraisal” only where referring specifically to US practice or to the titles of cited works.

References

Alexander, W.R.J. and Belloni, R. (2024), “Artificial intelligence and the sustainability of the signaling and human capital roles of higher education”, *Sustainability*, Vol. 16 No. 20, p. 8802, doi: 10.3390/su16208802.

Arnold, K.C., Chauncey, K. and Gajos, K.Z. (2020), “Predictive text encourages predictable writing”, *Proceedings of the 25th International Conference on Intelligent User Interfaces*, pp. 128-138.

Australian Taxation Office (2024), “Taxation statistics 2022-23: individuals statistics”, Commonwealth of Australia, Canberra, available at: https://www.ato.gov.au/about-ato/research-and-statistics/in-detail/taxation-statistics/taxation-statistics-2022-23/statistics/individuals-statistics

Baur, K., Rosenfelder, M. and Lutz, B. (2023), “Automated real estate valuation with machine learning models using property descriptions”, *Expert Systems with Applications*, Vol. 213, 119147, doi: 10.1016/j.eswa.2022.119147.

Cajias, M. and Freudenreich, A. (2024), “What are tenants demanding the most? A machine learning approach for the prediction of time on market”, *Journal of Property Investment and Finance*, Vol. 42 No. 2, pp. 151-165, doi: 10.1108/jpif-09-2023-0083.

Connelly, B.L., Certo, S.T., Ireland, R.D. and Reutzel, C.R. (2011), “Signaling theory: a review and assessment”, *Journal of Management*, Vol. 37 No. 1, pp. 39-67, doi: 10.1177/0149206310388419.

Cvijanović, D., Milcheva, S. and van de Minne, A. (2022), “Preferences of institutional investors in commercial real estate”, *The Journal of Real Estate Finance and Economics*, Vol. 65 No. 2, pp. 321-359, doi: 10.1007/s11146-020-09816-y.

Dalalah, D. and Dalalah, O.M.A. (2023), “The false positives and false negatives of generative AI detection tools in education and academic research: the case of ChatGPT”, *International Journal of Management in Education*, Vol. 21 No. 2, 100822, doi: 10.1016/j.ijme.2023.100822.

Dugan, L., Hwang, A., Trhlik, F., Ludan, J.M., Zhu, A., Xu, H., Ippolito, D. and Callison-Burch, C. (2024), “RAID: a shared benchmark for robust evaluation of machine-generated text detectors”, arXiv:2405.07940.

Elkhatat, A.M., Elsaid, K. and Almeer, S. (2023), “Evaluating the efficacy of AI content detection tools in differentiating between human and AI-generated text”, *International Journal for Educational Integrity*, Vol. 19 No. 1, pp. 1-16, doi: 10.1007/s40979-023-00140-5.

French, N. (2023), “Pricing to market: an investigation into the use of comparable evidence in property valuation”, *Journal of Property Investment and Finance*, Vol. 41 No. 3, pp. 300-318, doi: 10.1108/jpif-02-2023-0007.

Heinig, S. and Nanda, A. (2018), “Measuring sentiment in real estate: a comparison study”, *Journal of Property Investment and Finance*, Vol. 36 No. 3, pp. 248-258, doi: 10.1108/jpif-05-2017-0034.

Hutchison, N.E., Adair, A.S. and Findlay, N. (2011), “The impact of covenant strength on property pricing”, *Journal of Property Research*, Vol. 28 No. 2, pp. 167-188, doi: 10.1080/09599916.2010.538479.

IVSC Technology in Valuation Working Group (2025), “Navigating the rise of artificial intelligence in valuation: opportunities, risks, and standards”, IVSC Perspectives Paper. London: International Valuation Standards Council.

Jung, M., Panizo, C.F., Dugan, L., Chen, P.Y. and Liang, P.P. (2025), “Group-adaptive threshold optimisation for robust AI-generated text detection”, arXiv: 2502.04528.

Kandipati, A. (2025), “A systematic literature review on artificial intelligence based techniques for nurturing operations in the real estate sector”, *Journal of Real Estate Literature*, Vol. 33 No. 2, pp. 61-86, doi: 10.1080/09277544.2025.2466370.

Kok, N., Koponen, E.L. and Martínez-Barbosa, C.A. (2017), “Big data in real estate? From manual appraisal to automated valuation”, *Journal of Portfolio Management*, Vol. 43 No. 6, pp. 202-211, doi: 10.3905/jpm.2017.43.6.202.

Kriegbaum, A., Ebert, C. and Chatterjee, R.K. (2024), “Goal-chatbots selling condos: generative artificial intelligence in real estate”, *Journal of Marketing Development and Competitiveness*, Vol. 18 No. 4, pp. 69-92.

Lane, M.A., Seiler, M.J. and Seiler, V.L. (2015), “The impact of staging conditions on residential real estate demand”, *Journal of Housing Research*, Vol. 24 No. 1, pp. 21-36, doi: 10.1080/10835547.2015.12092095.

Levy, D., Frethey-Bentham, C. and Cheung, W. (2020), “Asymmetric framing effects and market familiarity: experimental evidence from the real estate market”, *Journal of Property Research*, Vol. 37 No. 1, pp. 1-20, doi: 10.1080/09599916.2020.1713858.

Luchtenberg, K.F., Seiler, M.J. and Sun, H. (2019), “Listing agent signals: does a picture paint a thousand words?”, *The Journal of Real Estate Finance and Economics*, Vol. 59 No. 4, pp. 617-648, doi: 10.1007/s11146-018-9674-z.

Lützkendorf, T. and Speer, T.M. (2005), “Alleviating asymmetric information in property markets: building performance and product quality as signals for consumers”, *Building Research and Information*, Vol. 33 No. 2, pp. 182-195, doi: 10.1080/0961321042000323815.

Ma, J. and Huang, K. (2023), “ChatGPT in real estate”, in Huang, K., Wang, Y., Zhu, F., Chen, X. and Xing, C. (Eds), *Beyond AI: Future of Business and Finance*, Springer Nature Switzerland, Cham, pp. 219-242.

Ministry of Housing, Communities and Local Government (2024), “English private Landlord Survey 2024: main report”, HMSO, London, available at: https://www.gov.uk/government/statistics/english-private-landlord-survey-2024-main-report

Pagourtzi, E., Assimakopoulos, V., Hatzichristos, T. and French, N. (2003), “Real estate appraisal: a review of valuation methods”, *Journal of Property Investment and Finance*, Vol. 21 No. 4, pp. 383-401, doi: 10.1108/14635780310483656.

Perkins, M. (2023), "Academic integrity considerations of AI large language models in the post-pandemic era: ChatGPT and beyond", *Journal of University Teaching and Learning Practice*, Vol. 20 No. 2, doi: 10.53761/1.20.02.07.

RICS (2023), "Comparable evidence in real estate valuation", in *RICS Professional Standard*, 1st ed., Royal Institution of Chartered Surveyors, London.

Sandler, M., Choung, H., Ross, A. and David, P. (2024), "A linguistic comparison between human and ChatGPT-generated conversations", *International Conference on Pattern Recognition and Artificial Intelligence*, Springer Nature Singapore, Singapore, pp. 366-380.

Statistics Canada (2023), "Residential real estate Investors and investment Properties in 2020", Canadian Housing Statistics Program, Catalogue No. 46-28-0001, Statistics Canada, Ottawa, available at: https://www150.statcan.gc.ca/n1/pub/46-28-0001/2023001/article/00001-eng.htm

Tajani, F., Morano, P. and Ntalianis, K. (2018), "Automated valuation models for real estate portfolios: a method for the value updates of property assets", *Journal of Property Investment and Finance*, Vol. 36 No. 4, pp. 324-347.

Tanlamai, J., Khern-am-nuai, W. and Cohen, M.C. (2024), "Generative AI and price discrimination in the housing market", working paper, Desautels Faculty of Management, McGill University, available at: https://ssrn.com/abstract=4764418

Topraklı, A.Y. (2025), "AI-driven valuation: a new era for real estate appraisal", *Journal of European Real Estate Research*, Vol. 18 No. 1, pp. 105-120, doi: 10.1108/jerer-04-2024-0031.

Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P. and Waddington, L. (2023), "Testing of detection tools for AI-generated text", *International Journal for Educational Integrity*, Vol. 19 No. 1, pp. 1-39, doi: 10.1007/s40979-023-00146-z.

Wheaton, W.C. and Xu, C. (2024), "Using AI to improve price transparency in real estate valuation", *MIT Center for Real Estate Research Paper No. 24/16*, Massachusetts Institute of Technology, Cambridge, MA, available at: https://ssrn.com/abstract=5056315

Xiong, C., Cheung, K.S., Levy, D. and Allen, M. (2024), "The effect of virtual reality on the marketing of residential property", *Housing Studies*, Vol. 39 No. 3, pp. 671-694, doi: 10.1080/02673037.2022.2074971.

You, T. (2020), "Behavioural biases and nonlinear adjustment: evidence from the housing market", *Applied Economics*, Vol. 52 No. 46, pp. 5046-5059, doi: 10.1080/00036846.2020.1752902.

Corresponding author
Clayton Pilat can be contacted at: claytonpilat@gmail.com