Human Factors, Business Management and Society, Vol. 215, 2026, 70–78
https://doi.org/10.54941/ahfe1007587

Strategic Personas at the Intersection of HCI and Marketing: A Framework Inspired By Virtual Chess Players

Khaldoon Dhou

College of Business Administration, Texas A&M University Central Texas, Killeen, TX 76549, USA

ABSTRACT

Conventional personas have always been attractive tools for HCI and marketing researchers. Although they allow researchers to show differences between personas by providing demographic and psychographic information, they generally do not explain how these personas behave or make decisions. This paper considers a unique perspective that views personas as strategic agents. In other words, the personality of a user emerges while making decisions under conditions of uncertainty. The paper is inspired by previous research on virtual chess players, where he explored players of the same strength, but with different playing styles, and identified that chess styles play an important role in determining game outcomes. Additionally, it develops a new attacker–defender orientation space and identifies three behavioral dimensions that can be detected based on the sequence of decisions or actions performed by a person. The paper also shows how to move these ideas from the domain of chess to marketing strategies and the design of interactive systems. Finally, it discusses the influence of these ideas on customer targeting, brand posture, interaction design, and AI-mediated agents. The findings confirm that strategic personas allow predicting behavior, performing simulations, and real-time adaptation. Most importantly, they augment traditional personas instead of replacing them.

Keywords: Personas, Human–computer interaction, Strategic behavior, Marketing analytics, Virtual agents, Chess

INTRODUCTION

Exploring personas has long been an attractive approach for improving understanding of markets (Jansen et al., 2020; Salminen et al., 2022). That is, it enables marketing and HCI domains to transform large and diverse groups of people into models that can be examined and applied in practical settings (Nielsen, 2019). For example, a B2C marketing campaign might consider modeling `Busy Sophia, a time-pressed working mother', to decide which channels and messages to use for grocery promotions. The campaign employs her demographic profile, shopping preferences, and lifestyle constraints to guide campaign design. However, the effectiveness of these models depends on narrowing the focus, since most personas treat demographic indicators, inferred preferences, and lifestyle characteristics as foundational descriptors of the user segment (Ilkka et al., 2025). Although such personas

Received March 7, 2026; Revised April 13, 2026; Accepted May 1, 2026; Available online July 20, 2026

are intended to describe users, their capacity to anticipate how users behave under changing conditions or uncertainty remains limited (McGinn and Kotamraju, 2008). This limitation is becoming more noticeable due to the transition from static campaigns to computerized, adaptive, and mediated encounters (Jansen et al., 2020). Recommendation engines, conversational agents, dynamic pricing mechanisms, and virtual brand representatives operate under temporal pressure, partial information, and feedback loops that unfold over sequences of interaction (Sidlauskiene et al., 2023). In such environments, user response cannot be reduced to fixed traits or stable preferences. What matters instead are decision patterns, tolerance for risk, sensitivity to tempo, and strategies for managing uncertainty—dimensions that conventional, segmentation-based persona practices tend to overlook, even as adaptive AI systems and advanced segmentation methods highlight the importance of behavior under uncertainty (Shneiderman, 2020; Wedel and Kamakura, 2000).

Parallel concerns have appeared in HCI research on virtual humans and intelligent agents. Early work often emphasized appearance, believability, and social presence. For example, Cassell (2000) described the development of Rea, an embodied conversational agent designed for the real estate domain. Unlike earlier `autonomous synthetic characters' that focused on general social skills and used caricature to encourage users to suspend disbelief, Rea was built on theories of HCI to ensure she shared actual conversational skills grounded in HCI. More recent research efforts have shifted toward action-based models in which behavior, rather than affect alone, carries the burden of personality. For instance, Han et al. (2025) demonstrated that Large Language Models (LLMs) can be leveraged to unify the generation of verbal and nonverbal behaviors (e.g., facial expressions) within a single framework to reflect specific personality traits. Their findings indicate that LLMs can generate personality-aligned behaviors that are statistically consistent with psychological literature. These personalities are also accurately perceived and recognized by human observers across diverse social contexts, such as negotiation. Additionally, prior research has successfully applied action-based modeling to virtual chess players, classifying them by strategic orientations such as attackers and defenders based on observable decision patterns rather than emotional traits (Dhou, 2018, 2019, 2020, 2021, 2023, 2024). These studies showed that agents with equivalent strength but differing action heuristics (e.g., opening aggression, sacrifice willingness) generated distinct interaction dynamics and risk responses (Dhou, 2018, 2021).

The present work extends this line of thinking into the domain of marketing personas by drawing on prior research that modeled virtual chess players as strategic agents as opposed to expressive characters. Virtual humans were constructed to emulate real players not through emotional scripting but through action-oriented parameters: opening selection, sacrifice propensity, tempo control, and responses to identical game states under uncertainty. Across multiple studies, agents with comparable strength but different strategic orientations (e.g., early attackers, solid defenders) produced systematically different interaction dynamics, error profiles, and game trajectories. What distinguished these agents was not outcome alone, but how they interacted

with opponents with different strategies and playing styles. These findings suggest a reframing of persona development. Instead of treating personas as static profiles anchored in demographics or psychographics, they can be reconceived as context-dependent strategic orientations. Using this framing, an "attacking"' persona is defined by a preference for early commitment, visibility, and initiative, even at the cost of stability. On the other hand, a "defensive" persona absorbs pressure, avoids unnecessary exposure, and optimizes for endurance. Crucially, such orientations are not fixed identities. In chess, the same player may shift strategy depending on position, opponent, or tournament context. Similarly, consumers and brands may oscillate between strategic postures across journey stages, market conditions, or competitive landscapes.

The present paper advances a conceptual framework that formalizes this shift. Building on action-based personality modeling from virtual humans and strategic behavior under uncertainty, it proposes marketing personas as dynamic agents whose defining properties lie in decision tempo, risk tolerance, and adaptive response patterns. The goal is not to displace existing segmentation approaches, but to augment them with a behavioral layer that supports prediction, simulation, and real-time adaptation—particularly in systems where artificial agents must engage humans continuously and strategically.

ORIGINS OF THE FRAMEWORK: VIRTUAL CHESS PLAYERS

The framework proposed in this paper does not originate from marketing theory but from a different problem space: how to model strategic individuality in artificial agents without reducing it to affect, narrative, or outcome-based metrics. The virtual chess player studies that motivate this work were initially concerned with a narrow HCI question—how virtual humans might emulate the playing styles of real chess players in a way that is perceptible, consistent, and analytically tractable (Dhou, 2018, 2024). What emerged instead was a more generalizable insight about behavior under uncertainty.

In these studies, virtual players were not differentiated through emotional states, dialogue, or expressive cues. Personality was encoded structurally, through action selection constraints and preferences. Agents shared equivalent ratings—thus comparable computational strength—but differed along parameters such as opening aggressiveness, willingness to sacrifice material, tolerance for positional imbalance, and control of tempo. These parameters were not cosmetic. They shaped the order, timing, and risk profile of decisions throughout play (Dhou, 2018; Dhou, 2019).

A consistent pattern appeared across experiments: agents exposed to identical board states diverged systematically in how they responded. Attacking-oriented agents tended to commit earlier, accepting local instability in exchange for initiative and pressure. Defensive-oriented agents delayed commitment, absorbed threats, and optimized for structural resilience. Importantly, these orientations were legible not only in isolated moves but in the unfolding trajectory of games—game length, error distribution, and opponent behavior shifted accordingly (Dhou, 2020; Dhou, 2021).

What proved most consequential was not victory or defeat but engagement with uncertainty. Attacking agents treated uncertainty as an opportunity to impose tempo, forcing opponents into reactive positions. Defensive agents treated the same uncertainty as something to be managed, constrained, and deferred. These differences persisted even when overall performance converged, suggesting that strategic orientation operates independently of strength. Two agents could "play equally well" while playing in fundamentally different ways.

Several features of this research are especially relevant for persona theory. First, personality was inferred from patterns of action, not from stated preferences or post hoc interpretation. Second, behavior was evaluated temporally, across sequences of decisions, rather than at isolated moments. Third, orientation was context-sensitive but not unstable; agents adapted without losing identity. Finally, strategic posture influenced not only the agent's own outcomes but the behavior of others interacting with it. These properties map cleanly onto contemporary marketing and HCI contexts. Brands, consumers, and digital agents increasingly interact in environments characterized by rapid feedback, competitive pressure, and partial information. As in chess, the question is often not what choice is made, but when, with what degree of commitment, and under what tolerance for risk. The virtual chess player work demonstrates that such patterns can be modeled explicitly, without recourse to demographic or psychographic proxies.

The attacker–defender distinction that anchors the present framework should therefore be understood as a structural abstraction, not a metaphor. It captures a spectrum of strategic engagement that emerged empirically from action-based modeling. In chess, these orientations manifest through opening choices, sacrifices, and tempo. In marketing contexts, they may manifest through campaign timing, pricing aggressiveness, feature rollout cadence, or response to competitor moves. The underlying logic is the same: strategic agents reveal themselves through how they navigate uncertainty over time.

STRATEGIC PERSONA FRAMEWORK

The strategic persona framework reframes personas as agents rather than profiles. Its core assumption is that personas are best understood through the structure of their decisions under uncertainty, not through static descriptors. Drawing on action-based modeling from virtual chess players, the framework treats persona identity as a pattern that emerges across sequences of choices—when commitment occurs, how risk is traded for initiative, and how tempo is regulated as conditions evolve (Dhou, 2018, 2021).

At the center of the framework lies a strategic orientation space rather than a typological catalog. The attacker–defender spectrum serves as an initial axis, not as a taxonomy. Attacking-oriented personas are characterized by early commitment, preference for initiative, and a willingness to tolerate imbalance. Defensive-oriented personas defer commitment, absorb pressure, and prioritize structural resilience. These orientations are not labels applied ex ante; they are inferred from behavior across repeated interactions, much

as strategic orientation was inferred from move sequences in virtual chess agents (Dhou, 2019, 2021).

Three behavioral dimensions operationalize this space: decision tempo, risk tolerance, and uncertainty handling. Decision tempo captures how quickly a persona commits when multiple viable options exist. High-tempo personas act early, often shaping the interaction environment for others. Low-tempo personas delay, seeking additional information or stability before acting. Tempo is not speed alone; it reflects a strategic stance toward uncertainty—whether it is exploited or contained.

Risk tolerance describes how willingly a persona accepts short-term instability in pursuit of longer-term advantage. In chess, this appears as material sacrifice or structural imbalance; in marketing contexts, it may surface as aggressive pricing, rapid feature release, or provocative messaging. Risk tolerance is revealed through patterns of escalation and retreat, not through stated preferences.

Uncertainty handling refers to how ambiguity is managed across time. Some personas externalize uncertainty, forcing counterparts into reactive positions; others internalize it, buffering volatility and prolonging optionality. This dimension proved decisive in virtual agent interactions, where agents of similar strength diverged markedly in how they shaped opponent behavior under identical conditions (Dhou, 2023).

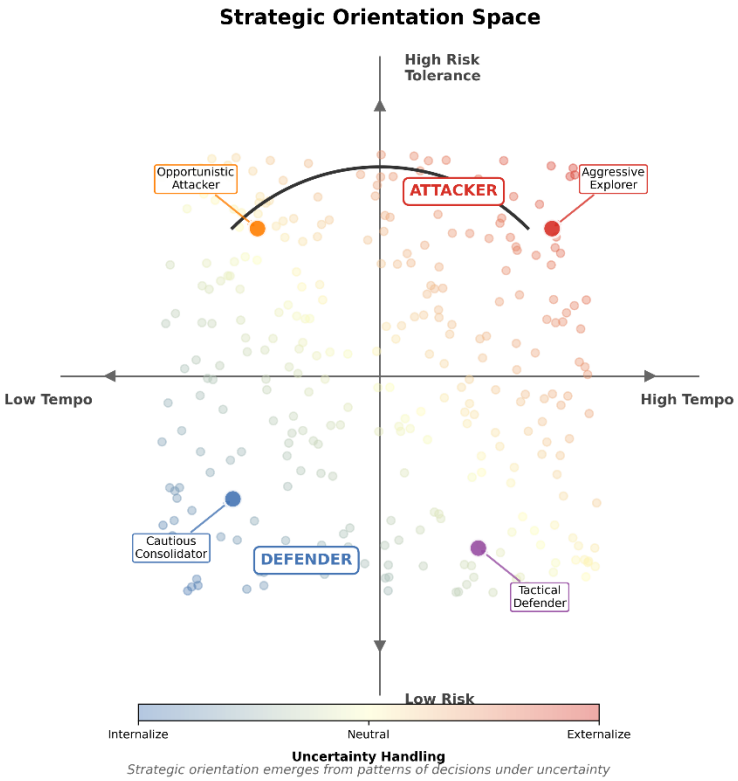

Figure 1: Three-dimensional mapping of persona decision-making under uncertainty.

These dimensions interact rather than operate independently. An attacking persona typically combines high tempo with elevated risk tolerance, but may modulate uncertainty handling depending on context. A defender may operate at low tempo with conservative risk exposure, yet still seize initiative selectively when structural conditions favor it. The framework therefore resists fixed categorization. Strategic orientation is treated as context-dependent but bounded (i.e., adaptive without being incoherent). Figure 1 shows the Strategic Orientation Space, where personas are positioned according to decision tempo and risk tolerance. The horizontal axis represents tempo (low to high), and the vertical axis represents risk tolerance (low to high). Color encodes uncertainty handling, from internalizing (blue) through neutral (yellow) to externalizing (red). The diagonal arc marks the attacker–defender spectrum, with example personas placed along it to illustrate how orientations vary across the space.

Crucially, the framework distinguishes orientation from outcome. In both chess and marketing contexts, success metrics can converge even as strategies diverge. Two personas may achieve similar conversion rates, market share, or engagement while exhibiting fundamentally different strategic signatures. Evaluating personas solely through outcomes collapses these distinctions. Evaluating them through decision patterns preserves them.

The framework also accommodates orientation shifts. In virtual chess studies, agents adjusted posture in response to opponent style and game phase without abandoning their underlying strategic identity. The same logic applies here. A brand persona may adopt an attacking posture during market entry, then transition toward defense as maturity and saturation set in. A consumer persona may oscillate between exploration and consolidation across the customer journey. These shifts are meaningful only when anchored in a stable decision logic.

Formally, the strategic persona framework treats personas as parameterized agents embedded in interaction environments. Traditional demographic or psychographic attributes remain relevant, but they function as constraints rather than definitions. What differentiates personas at the strategic level is how they act when constrained—how they trade speed for certainty, exposure for control, initiative for endurance.

This reframing enables personas to support simulation rather than description. Because strategic orientation is expressed through decision rules, personas can be instantiated in agent-based models, adaptive interfaces, and AI-driven marketing systems. More importantly, they can be tested—not for narrative plausibility, but for behavioral consistency across scenarios. This property, inherited directly from the virtual chess player lineage, marks a departure from conventional persona practice.

IMPLICATIONS FOR MARKETING, HCI, AND AI SYSTEMS

Reconceptualizing personas as strategic agents alters how they function across the marketing pipeline. In targeting and segmentation, the shift redirects attention from who customers are to how they engage. Strategic orientation provides leverage where demographic similarity masks behavioral

divergence. Two customers may share age, income, and usage patterns yet differ sharply in tempo and risk tolerance. An attacking-oriented customer responds to early commitment—limited releases, time-bound offers, visible escalation—while a defensive-oriented customer resists such pressure and rewards stability and reassurance. Treating these orientations as first-order variables reframes targeting as anticipation of play style rather than classification of traits (Wedel and Kamakura, 2000).

Brand strategy benefits from a similar reframing. Competitive posture is often discussed in metaphorical terms—aggressive launches, defensive moats—but rarely operationalized at the level of decision patterns. The strategic persona framework renders posture legible as a sequence of choices over time: when a brand escalates, when it consolidates, and how it manages exposure under uncertainty. This perspective accommodates strategic drift without incoherence. A brand can shift from attack to defense across lifecycle stages without losing identity, provided the underlying decision logic remains consistent. Such coherence mirrors findings from virtual chess agents, where orientation adapted to context without dissolving into randomness (Dhou, 2021, 2023).

In creative design and messaging, the implications are more subtle. Strategic orientation constrains not only what is said but when and how forcefully. High-tempo personas favor messages that invite immediate response and signal initiative; low-tempo personas privilege pacing, accumulation of evidence, and delayed commitment. These differences are not reducible to tone alone. They shape cadence, sequencing, and the acceptable degree of ambiguity in communication. Treating personas as strategic agents thus informs orchestration across touchpoints rather than isolated message optimization.

For HCI, the framework aligns naturally with interaction design in adaptive systems. Interfaces increasingly respond to user behavior in real time, yet the logic guiding adaptation is often shallow—reactive rather than strategic. Embedding strategic personas allows systems to reason about users' decision patterns across sessions, not just within them. A system interacting with an attacking-oriented user may surface options earlier, tolerate exploratory error, and escalate affordances quickly. With a defensive-oriented user, the same system may prioritize reversibility, preview, and constraint. This mirrors the design logic of virtual humans whose personalities are inferred from action histories rather than declared states (Cassell, 2000; Riedl and Young, 2010).

The implications are especially pronounced for AI-mediated agents. Chatbots, recommenders, and virtual influencers increasingly operate as autonomous participants in interaction loops rather than passive tools. In such contexts, personas guide not only representation but behavior. A recommender system aligned with an attacking persona may introduce novel options aggressively and probe boundaries; a defensive counterpart may refine within known preferences and minimize surprise. Because strategic orientation is parameterizable, these personas can be instantiated, tested, and tuned within agent-based simulations—an inheritance from the virtual chess player lineage (Dhou, 2018, 2020).

More broadly, the framework suggests a different evaluation criterion for persona effectiveness. Rather than asking whether a persona is vivid or relatable, the relevant question becomes whether it produces consistent decision behavior across scenarios. Consistency here does not imply rigidity. It implies bounded adaptability—a capacity to respond to context while preserving strategic identity. This property proved decisive in virtual agent studies, where agents that were either overly rigid or overly reactive produced brittle interactions (Dhou, 2019). The same risk applies in marketing and HCI systems that adapt without a stable strategic core.

These implications point toward a convergence between persona theory and agent-based modeling. Personas cease to be static artifacts and become lightweight behavioral models. This does not require abandoning qualitative insight or ethnographic grounding; it requires situating those insights within a framework that treats decision-making under uncertainty as the primary unit of analysis. For domains increasingly shaped by real-time interaction between humans and artificial agents, this shift is less an enhancement than a necessity.

REFERENCES

Cassell, J. (2000). Embodied conversational interface agents. *Communications of the ACM*, *43*(4), 70–78.

Dhou, K. (2018, June). Towards a better understanding of chess players' personalities: a study using virtual chess players. In *International Conference on Human-Computer Interaction* (pp. 435–446). Cham: Springer International Publishing.

Dhou, K. (2019, July). An innovative employment of virtual humans to explore the chess personalities of Garry Kasparov and other class-A players. In *International Conference on Human-Computer Interaction* (pp. 306–319). Cham: Springer International Publishing.

Dhou, K. (2020, July). A novel investigation of attack strategies via the involvement of virtual humans: a user study of Josh Waitzkin, a virtual chess grandmaster. In *International Conference on Human-Computer Interaction* (pp. 658–668). Cham: Springer International Publishing.

Dhou, K. (2021). An exploration of chess personalities in grandmasters and class-a players using virtual humans. *International Journal of Entertainment Technology and Management*, *1*(2), 126–145.

Dhou, K. (2023, July). An exploration of feared versus fearless attack attitudes using the chess personalities of virtual chess players. In *International conference on human-computer interaction* (pp. 176–185). Cham: Springer Nature Switzerland.

Dhou, K. (2024, June). Exploring Virtual Chess Personalities: Analyzing Class-C Players vs. Nimzowitsch and Fine Virtual Grandmasters. In *International Conference on Human-Computer Interaction* (pp. 311–321). Cham: Springer Nature Switzerland.

Han, B., Kwon, D., Lin, S., Shrestha, K., & Gratch, J. (2025, September). Can LLMs Generate Behaviors for Embodied Virtual Agents Based on Personality Traits?. In *Proceedings of the 25th ACM International Conference on Intelligent Virtual Agents* (pp. 1–10).

Ilkka, K., Joni, S., Soon-Gyo, J., Joao, S., Kholoud, A., Essi, H., ... & Jansen, B. J. (2025). Demographics do not matter?: Exploring the impact of gender and ethnicity on users' identification with AI-generated personas. *International Journal of Human-Computer Studies*, *202*, 103548.

Jansen, B. J., Salminen, J. O., & Jung, S. G. (2020). Data-driven personas for enhanced user understanding: Combining empathy with rationality for better insights to analytics. *Data and Information Management*, *4*(1), 1–17.

McGinn, J., & Kotamraju, N. (2008, April). Data-driven persona development. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 1521–1524).

Nielsen, L. (2019). *Personas: User focused design* (2nd ed.). Springer.

Riedl, M. O., & Young, R. M. (2010). Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, *39*, 217–268.

Salminen, J., Wenyun Guan, K., Jung, S. G., & Jansen, B. (2022, April). Use cases for design personas: A systematic review and new frontiers. In *Proceedings of the 2022 CHI Conference on human factors in computing systems* (pp. 1–21).

Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human–Computer Interaction*, *36*(6), 495–504.

Sidlauskiene, J., Joye, Y., & Auruskeviciene, V. (2023). AI-based chatbots in conversational commerce and their effects on product and price perceptions: J. Sidlauskiene et al. *Electronic Markets*, *33*(1), 24.

Wedel, M., & Kamakura, W. A. (2000). *Market segmentation: Conceptual and methodological foundations*. Springer Science & Business Media.