

Volume 2, Issue 3

Research Article

Date of Submission: 10 Aug, 2026

Date of Acceptance: 11 Sep, 2026

Date of Publication: 21 Sep, 2026

Knowledge Panels and Person Entity Confidence in the AI Search Era: A Signal, Traffic, and AI-Visibility Audit of 1,005 US Personal Injury Attorneys and Law Firms

Behzad Hussain* 

Independent Researcher, Pakistan

Corresponding Author:

Behzad Hussain, Independent Researcher, Pakistan.

Citation: Hussain, B. (2026). Knowledge Panels and Person Entity Confidence in the AI Search Era: A Signal, Traffic, and AI-Visibility Audit of 1,005 US Personal Injury Attorneys and Law Firms. *J Adv Robot Auton Syst Hum Mach Interact*, 2(3), 01-14.

Abstract

The Google Knowledge Panel has evolved from a rich-snippet convenience into the visible surface of an entity-resolution machinery that also feeds Google AI Overviews and third-party generative answer engines. In the personal injury (PI) legal vertical, where practice is name-driven and commercial intent is dense, the practical question is whether a firm's managing partner has crossed the Knowledge-Graph confidence threshold that Google requires before rendering a Person-template Knowledge Panel. This study audits the Knowledge Panel status of the managing partner or firm owner of 1,005 top-ranking US personal injury law firms across 51 markets, classifies returned panels into a three-class taxonomy (Person Entity Panel, Local Business Panel, and Minimal or Other), and measures corpus and on-page signals associated with each class. We find that 21.3% of firms surface some Knowledge Panel when the managing partner's name is queried, split into 3.5% Person Entity Panels, 16.4% Local Business Panels, and 1.4% Minimal panels. Among the 35 Person Entity Panel holders, 67.6% anchor on a Wikipedia article as the canonical description source and 59.5% exhibit recent national-publisher media coverage versus a matched control rate of 32.7%. Schema-markup adoption does not meaningfully predict Person Entity Panel presence: Attorney/Person schema shows a 0.99x lift, and the only signal with measurable positive lift is Person-nested-in-Organization at 1.54x. Ten firms deploy the full entity-signal schema stack ($SCI \geq 17$) yet still have no Person Entity Panel, evidencing that schema completeness is necessary hygiene rather than sufficient signal. Merging monthly organic-traffic estimates into the panel table reveals that Person-Entity-Panel firms carry a 2.95x traffic multiplier (Mann-Whitney $p = 0.0003$) after controlling for SERP rank, appearance frequency, and schema completeness, while Local-Business-Panel firms carry a 32% traffic deficit (multiplier 0.68x, $p = 0.0023$), consistent with zero-click cannibalization. A qualitative audit of six queries across Google AI Overviews and third-party generative answer engines shows Person-Entity-Panel attorneys are cited by name with unhedged verbs and grounded in Wikipedia and national-tier media, while a schema-heavy no-panel control is cited only to self-attested firm content. We argue that the confidence-threshold construct we term Person Entity Confidence, together with the practitioner discipline we term Brand Context Optimization, is the correct frame for legal-vertical AI Visibility in 2026 and beyond.

Keywords: Knowledge Panel, Knowledge Graph, Person Entity, Entity Recognition, Google, AI Overviews, AI Visibility, Answer Engine Optimization, Generative Engine Optimization, Brand Context Optimization, Wikipedia, Personal Injury Law, Legal SEO, Schema Markup, E-E-A-T, Large Language Models

Introduction

The digital marketplace for personal injury (PI) legal services in the United States is one of the most competitive commercial verticals on the open web. First-page visibility for practice-area queries such as “personal injury lawyer [city]” and “car accident lawyer [city]” translates directly into intake volume and, thus, into firm revenue. An earlier working paper by the present author audited the schema-markup implementation of 1,005 PI law firm websites ranking on Google Page 1 across 51 U.S. markets and documented that only 35.3% deploy the industry-specific Legal Service type and only 20.6% publish Person or Attorney schema for their lawyers. That study frames the on-page discipline. It does not, however, ask whether the discipline has translated into the observable outcome that Google’s own patents identify as the codified endpoint of entity understanding: a Knowledge Panel on the search results page. Two shifts in the search-engine layer have raised the stakes of that endpoint beyond conventional SEO. First, since the launch of Google’s Knowledge Graph Google’s understanding of a query has moved from string matching to entity resolution: the system attempts to resolve “who or what is this query about?” before it returns documents. A Knowledge Panel is the visual surface that Google’s Knowledge Panel patent codifies as the machinery of that resolution, defined explicitly in claim 3 as consuming a larger area of the search results page than any individual organic result. Second, the arrival of generative answer surfaces such as Google AI Overviews and third-party generative answer engines has added a new class of consumer surfaces that retrieve from Knowledge-Graph-adjacent substrates rather than crawling documents at answer time. The consequence is that entity recognition at Google now propagates into visibility across multiple AI answer engines simultaneously; the absence of recognition removes an attorney from that shared retrieval substrate. Absence of a Knowledge Panel is not evidence of absence in the Knowledge Graph. It is evidence that the confidence Google has accumulated for a given person entity has not yet crossed the visibility threshold that would justify rendering a Person template. Google’s own Knowledge Panel Help documentation states this directly: “Knowledge panels are created automatically by Google Search when there is enough information available on the open web” The academically loaded verb in that sentence is enough: the graph carries confidence scores, and only when the score for a person entity exceeds a threshold does the panel render. Throughout this paper we adopt the term Person Entity Confidence to name that threshold construct, and we ask which corpus-level and on-page signals systematically accompany attorneys who have crossed it. Our contribution is fourfold. First, we assemble the first large-N Knowledge Panel audit of the U.S. personal injury legal vertical, covering 1,005 top-ranking firms and their identified managing partners. Second, we propose a three-class Knowledge Panel taxonomy (Person Entity Panel, Local Business Panel, Minimal or Other) and provide a diagnostic decision procedure that practitioners can apply in the wild. Third, we combine external corpus-signal enrichment (Wikipedia lookup, book authorship via Google Books, and national-publisher news coverage) with the companion study’s on-page structured-data signals in a single analytical table, producing what we believe is the first empirical ranking of the factors that systematically accompany Person Entity Confidence for attorneys in this vertical. Fourth, we merge monthly organic-traffic estimates for the same 1,005 firms and report the traffic differential across the panel-class taxonomy, closing the loop from panel presence to business impact [1-6].

Literature Review

The Knowledge Graph and the Knowledge Panel

The intellectual origin of an entity-centric search paradigm is the Semantic Web vision articulated by in which machine-readable metadata renders web content as interpretable to software agents as it is to human readers. Freebase provided the first large-scale collaboratively-constructed graph database of entities and relations, and it was acquired by Google in 2010 and folded into what became the Knowledge Graph Knowledge Vault extended that graph with a probabilistic fusion mechanism that combines extractions from text, tabular data, page structure, and human annotations, where “human annotations” includes declared structured data such as Schema.org markup The Knowledge Panel is the user-facing surface of that graph. Google’s patent US-11836177-B2 defines the panel plainly. Its abstract describes a search system that obtains search results responsive to a query, identifies a factual entity referenced by the query, and presents a knowledge panel of content selected for that entity alongside the results. Claim 3 makes the visual weight of the panel explicit: the Knowledge Panel area consumes a larger area of the search results page than each of the individual search results. For a personal injury attorney or firm, this means the Knowledge Panel is not a decorative appendage. It is patented as the largest single content unit on the results page, and it renders only when Google resolves the query to a factual entity it has in the graph [7-10].

Knowledge Panel Triggering, Templates, and Ranking

Claim 1 of US-11836177-B2 requires the system to identify one or more queries received with a reference to a given factual entity, then obtain resources related to the query, then select content for a knowledge panel that represents the entity. Two implications follow for legal-vertical entities. First, the panel does not appear because the site ranks for the query; it appears because the query has been resolved to a specific factual entity in the Knowledge Graph. Ranking follows entity resolution, not the reverse. Second, panel content is not first-come-first-served: claim 4 states that selecting content for the panel is based on a ranking of content for the given factual entity, and that the ranking is based on user search events related to the content and the entity. Panels are rendered from templates keyed to entity type. Claim 8 of US-11836177-B2 specifies that if the factual entity is a Person, the template includes a placeholder for an image, a description, and at least one fact about the person. Claim 9 specifies that if the entity is a place, the template includes placeholders for a map image, a description of the place, and at least one fact. Claim 7 specifies that content selection comprises identifying types of content items for the identified template. For a personal-injury attorney, this is the mechanical reason the panels observed in this sample split cleanly into Person Entity Panels (image, biographical

facts, Wikipedia-sourced description) and Local Business Panels (address, phone, hours, map). The two templates draw from different placeholders and require different upstream signals to populate. Claim 6 of US-11836177-B2 also codifies the disambiguation mechanism: when a query is associated with multiple distinct meanings, the panel may include content for two or more of those meanings, with a link that allows the user to select a specific meaning. For a personal-injury attorney sharing a common name with a celebrity or another attorney, this is the mechanism that determines which entity receives the panel. The personal-knowledge-panel interface patent extends the surface to combine public entity data with personalized signals for signed-in users [11].

Cross-Source Corroboration and Submitted-Data Accuracy

The confidence with which Google resolves an attorney as a distinct entity depends on cross-source corroboration. The patent on corroborating facts extracted from multiple sources states that facts extracted as attribute-value pairs from multiple sources are corroborated by cross-source agreement on the same subject. For a personal-injury attorney, this is the mechanism that rewards consistency between the firm's own site, the Google Business Profile listing, the state bar registry, LinkedIn, legal directories, and (where present) a Wikipedia article. Each independent source that agrees on a fact (birth year, education, firm affiliation) raises the corroborated confidence for the underlying entity node. A complementary patent family on predicting the accuracy of submitted data describes systems that determine whether a submission of data by a user is accurate. Third-party submissions to a Knowledge Panel (such as "suggest an edit") are not accepted uncritically; the system predicts accuracy and down-weights submissions that conflict with corroborated sources. For a firm with a Local Business Panel this is the machinery that determines whether an incorrect competitor edit or a customer-submitted correction propagates into the panel. Firms whose own site, Google Business Profile, and directory listings are internally consistent are less likely to see inaccurate submissions accepted.

Local Business Panels, Related Entities, and Question-Answering

Local Business Panels are governed by a distinct set of mechanisms. The local business ranking patent describes how local business search combines business data with mapping data to rank results. The sentiment-detection patent claims a method where the system ranks entities using sentiment detected in reviews. The 5local-experts patent identifies expert reviewers of a category of businesses in a geographic area. The detection-of-related-local-entities patent and the entity-review-extraction publication extend the local-entity graph and its review corpus. For a personal-injury attorney whose firm holds a Local Business Panel, this review corpus is a first-class ranking and confidence input, not a decorative badge. The related-entities patent describes a search system that provides one or more knowledge panels in response to a search request and populates information about groups of entities related to the primary entity. This is the mechanism behind the "People also search for" and related-attorney or related-firm clusters that render adjacent to a Knowledge Panel. Query-suggestion mechanisms further populate the "People also ask" and follow-up query surfaces from entity aliases and categories. The question-answering patent grounds answer-strip content in entity references extracted from unstructured data, which is one path by which factual answers about a jurisdiction's statute of limitations or contingency-fee structure surface near or inside a panel [12-21].

AI Answer Engines and Prior Work on Legal-Vertical Entity SEO

The Knowledge Graph substrate that feeds Google's Knowledge Panel also feeds Google's AI Overviews which retrieve from Knowledge-Graph-adjacent stores to synthesize answer text, and which cite external sources as provenance chips inside the AI Overview surface. The same underlying corpus is the retrieval substrate for third-party generative answer engines that ground on the open web through retrieval-augmented generation pipelines. Microsoft's parallel research program on entity retrieval and cards confirms that this cross-engine retrieval logic is not idiosyncratic to Google. Empirical work specifically on Knowledge Panel adoption in the legal vertical is essentially absent from the literature. Related studies on structured-data adoption in the personal-injury legal vertical include Hussain's (2026) 500-firm analysis (April 2026) and 1,005-firm analysis (July 2026) of on-page schema implementation, neither of which measures the Knowledge Panel outcome. Documents the notability criteria that govern whether a lawyer receives a Wikipedia article, which as our findings will show is the dominant corpus-level signal for a Person Entity Panel in this vertical. The present study fills this gap [22-24].

Knowledge Panels and Person Entity Confidence in the AI Search Era

Section 2 documented the patents and research that describe how Google resolves a query into a factual entity and renders a Knowledge Panel keyed to that entity's type. This section develops the conceptual scaffolding this paper uses to interpret the empirical findings that follow. Four constructs are foundational: the two-panel-class taxonomy, the Person Entity Confidence threshold, the composite property we call AI Visibility, and the practitioner discipline we call Brand Context Optimization.

The Two Panel Classes: Person Entity Panel vs Local Business Panel

When a searcher queries the name of a personal-injury attorney, Google may render one of three distinct surfaces: no panel at all, a Local Business Panel keyed to the firm's Google Business Profile, or a Person Entity Panel keyed to the individual as a person entity in the Knowledge Graph. These two classes are not positions on a maturity ladder that a Local Business Panel firm climbs into a Person Entity Panel firm. They are structurally distinct entity types produced by different templates (claims 8 and 9 of US-11836177-B2, [4] and they carry different implications for AI Visibility. A Local Business Panel carries the firm's address, phone number, hours, click-to-call button, and driving directions. A Person

Entity Panel carries the individual's biographical attributes: date of birth, education, spouse, children, books, awards, or (for attorneys with prior political careers) party affiliation and previous office. The former is a business-listing surface; the latter is an entity-recognition surface.

Person Entity Confidence as a Threshold Construct

Google's own Help documentation describes a Knowledge Panel as being created "when there is enough information available on the open web". The critical construct in that sentence is enough. Google's corroboration mechanism and its submitted-data accuracy predictor both operate on confidence-scored entity nodes, not on binary "in the graph or not" flags. Every attorney in this sample may have some degree of Knowledge Graph representation. The subset that surfaces a Person Entity Panel is the subset whose confidence score has crossed the visibility threshold that Google's rendering machinery requires. Throughout this paper we adopt the term Person Entity Confidence to name that threshold construct. The empirical question we ask in Sections 5.2 through 5.5 is which corpus-level and on-page signals systematically accompany attorneys who have crossed it, not which single signal causes it to be crossed. Absence of a Person Entity Panel is evidence that Person Entity Confidence has not crossed the visibility threshold at the moment of measurement; it is not evidence that the underlying person entity does not exist in the Knowledge Graph. The distinction is load-bearing throughout the paper.

AI Visibility and Knowledge-Graph-Adjacent Retrieval

We use AI Visibility to name the composite property of being retrievable and citable in AI-generated answers surfaced by Google AI Overviews and by third-party generative answer engines built on retrieval-augmented generation over open-web corpora. These systems do not evaluate the corpus at answer time by crawling the open web from scratch; they retrieve from Knowledge-Graph-adjacent substrates and from open-web corpora that overlap heavily with the corpora that feed Google's Knowledge Panel. A person entity with high Person Entity Confidence in Google's Knowledge Graph is therefore structurally retrievable across these surfaces simultaneously; a person entity with low confidence is retrievable only through ad-hoc crawl at answer-generation time, which yields thinner and less consistent citations. The clustered-search-results, search-result-ranking-and-presentation and question-answering patents codify the retrieval mechanism [25,26].

Brand Context Optimization: The Practitioner Discipline

The practitioner discipline that raises Person Entity Confidence over time is the coherent, corroborated construction of brand context across independent sources: the firm's own website (with complete Organization, Legal Service, and Person/Attorney schema), a verified and complete Google Business Profile, the jurisdictional regulatory-body record (state bar in the U.S., SRA and Law Society in the U.K., provincial law societies in Canada), authoritative legal directories (Avvo, Martindale-Hubbell, Super Lawyers, Chambers, Legal 500, Lexpert), consistent national or regional press coverage, and, at the highest tier, an editorially verified anchor such as a Wikipedia article, a book published at a major legal or trade press, or an elected or appointed public office. We use Brand Context Optimization (BCO) to name this discipline. BCO is the compounding investment that raises the corroborated confidence score the Google patent family rewards; individual on-page tactics are components of BCO, not substitutes for it. Section 7 formalizes the discipline as a five-level maturity model [27].

What the Framework Predicts About PI Attorney Panels

The framework above yields four empirical predictions that structure our findings. First, panels observed in the wild should split cleanly into Person Entity Panels and Local Business Panels, with a small residual of Minimal or empty-attribute panels. Second, Person Entity Panels should be strongly associated with corpus-level signals (Wikipedia articles, national-tier media coverage, book authorship, prior political office) and only weakly associated with on-page schema signals. Third, on-page schema signals should nonetheless correlate with Local Business Panel presence, because the local-business schema is the direct feed into the local-entity graph the local-business ranking patent consumes. Fourth, if AI Visibility depends on the same Knowledge-Graph-adjacent retrieval substrate, an attorney with Person Entity Confidence should be cited by name across AI answer engines with unhedged verbs, and an attorney without should be cited only against self-attested firm content or not cited at all. All four predictions are tested in Section 5.

Methodology

Sample Construction

The 1,005-firm sample is inherited directly from the companion 1,005-firm structured-data study. In that study, 255 Google SERP queries were issued across five practice-area keyword templates ("personal injury lawyer", "car accident lawyer", "truck accident lawyer", "medical malpractice lawyer", "wrongful death lawyer") combined with 51 U.S. markets consisting of the most populous city in each of the 50 U.S. states plus Washington, D.C. Legal directories, business directories, social networks, forums, encyclopaedic sources, news outlets, and government or military domains were removed from the top-10 organic results. After root-domain deduplication (each domain kept at its best observed rank across all queries in which it appeared), the sample comprised 1,005 unique law firm and attorney websites. This sample forms the denominator for every rate reported in the present paper [28].

Managing Partner Identification

For each of the 1,005 firms we retrieved the top ten Google organic results and any accompanying Knowledge Graph

payload for the query pattern <domain> managing partner OR founder OR owner. The retrieved results were passed through a language-model-assisted extraction step with a schema-forced JSON return producing {partner name, partner role, confidence, reasoning}. The extraction was instructed to reject firm-brand names, actor and celebrity Knowledge Graph collisions, and multi-firm noise, and to return an empty name string when confidence was low. This procedure identified a managing partner or firm principal for 944 of 1,005 firms (93.9%), with 89.9% of identifications flagged high-confidence. The remaining 61 firms without an identifiable head were retained in the sample for the overall Knowledge Panel presence denominator but excluded from partner-level analyses.

Knowledge Panel Audit

For each identified head we retrieved the Google Search Knowledge Graph payload for the query pattern <partner name> lawyer with U.S. locale and English language parameters. We captured the returned knowledge Graph object and extracted title, type, description, description Source, description Link, image URL, website, attributes, and social Profiles. Presence of a non-empty title in the returned object was recorded as a Knowledge Panel hit.

Knowledge Panel Classification

Panels were classified into three mutually exclusive classes by the composition of their attributes dictionary. A Person Entity Panel is a panel containing at least one biographical attribute (Born, Died, Education, Spouse, Children, Parents, Siblings, Books, Movies, Awards, Party, Nominations, Previous office, Height, or Net worth) and no business attributes, or two or more biographical attributes regardless of business attributes. A Local Business Panel is a panel containing at least one business attribute (Address, Phone, Hours, Appointments, Menu, Reservations, or Order). Minimal or Other covers cases where a Knowledge Panel object was returned with a title but with no attribute recognized under either classification.

External Corpus-Signal Enrichment

For each of the 214 Knowledge Panel holders and, separately, for a rank-matched control sample of 200 non-panel partners, we ran three additional lookups. First, we searched Wikipedia by the partner's name and inspected the top five article results, recording a match when an article title contained the partner's last name (and, for partners with distinguishing first names, the first name). Article byte size was captured as a proxy for article depth. Second, we searched Google Books by author with an in author: filter and recorded the total-items count and a sample title where present. Third, we retrieved the top ten Google news results for the partner's name and captured source, date, and title for each item. National-publisher hits were counted by substring match against a curated allowlist of major U.S. national outlets (New York Times, Wall Street Journal, Washington Post, CNN, Fox News, Reuters, Associated Press, Bloomberg, USA Today, NPR, BBC, CBS News, ABC News, NBC News, CNBC, Los Angeles Times, Chicago Tribune, Boston Globe, Time, Newsweek, Forbes, Politico, The Guardian, The Atlantic, The Hill, PBS, ProPublica, Law360, Law.com, Bloomberg Law, and ABA Journal).

Schema Markup Cross-Reference

Every Knowledge Panel record was joined on domain to the per-domain schema flags produced by the companion 1,005-firm structured-data study. The joined fields are has_organization, has_localbusiness, has_legalservice, has_attorney_person, has_person_nested, has_at_id, has_sameas, has_review, has_faqp, has_breadcrumblist, at_id_count, sameas_count, estimated_attorneys, firm_size, maturity_label, sci_total (the Schema Completeness Index total), and or_eligible_count (the number of Google Rich Results features for which the site qualifies).

Organic Traffic Integration

Monthly organic-traffic estimates and total organic-keyword counts for each of the 1,005 firm websites were obtained from a third-party keyword-database export and joined on normalized domain (lower-cased, www.-stripped). Traffic values are algorithmic estimates from a single tool; they are treated as relative rather than absolute in every analysis below. The distribution is heavily right-skewed (median 430, mean 2,413, maximum 404,100 monthly visits), so all parametric analyses operate on $\log_{10}(\text{traffic} + 1)$, and descriptive statistics report medians alongside means. Eight of the 1,005 firms lacked a traffic value and were excluded from Section 5.8 analyses ($n = 997$).

AI Visibility Qualitative Audit

Six queries were selected to test whether Person Entity Confidence propagates into third-party AI answer engines. Queries covered two top-tier Person Entity Panel holders ("who is Ben Crump", "Morgan and Morgan founder"), one mid-tier holder ("who is Michele S. Mirman lawyer"), one schema-heavy no-panel control ("who is Fredrick Schreck personal injury lawyer") drawn from the ten schema-complete firms without a Person Entity Panel identified in Section 5.6, one case-driven query ("attorney who represented George Floyd family"), and one generic local commercial query ("best personal injury lawyer in Miami"). We captured Google AI Overviews and a representative third-party generative answer engine in the study's data-capture window and recorded the visible citation chips and named external sources. Analyses of the captured surfaces are in Section 5.7.

Findings

Knowledge Panel Prevalence and Class Distribution

Of the 1,005 firms in the sample, 214 (21.3%) surface a Knowledge Panel when the identified managing partner's name is queried on Google. Table 1 gives the panel-class distribution.

Panel class	Count	% of 1,005
No Panel	791	78.7%
Local Business Panel	165	16.4%
Person Entity Panel	35	3.5%
Minimal / Other	14	1.4%

Table 1: Knowledge Panel Class Distribution Across the 1,005-Firm Sample

Among the 944 firms for which a head was identified at high or medium confidence, the Person Entity Panel rate is 3.7% and the any-panel rate is 22.7%. Panel types (the raw type string in the Knowledge Graph payload) further stratify by geography and specialization: the city-scoped type “Personal injury attorney in {city, state}” is dominant among Local Business Panels (74 of 165), while generic biographical types such as “American attorney”, “Attorney”, “Lawyer”, and “American lawyer” dominate among Person Entity Panels (73 of 35 when types are counted with multiplicity across the small cohort). This is the empirical corollary of the template-selection mechanism codified in claims 8 and 9 of US-11836177-B2 different templates draw from different type strings.




Figure 1: Knowledge Panel Class Distribution across 1,005 US Personal Injury Law Firms

Person Entity Panel Signals: Wikipedia Dominance

Among the 35 firms whose managing partner surfaces a Person Entity Panel, the dominant corpus-level signal is Wikipedia. Combining direct Wikipedia article matches by name with cases where the Knowledge Panel’s description Source is Wikipedia, 67.6% of the Person Entity Panel cohort has a Wikipedia touchpoint. Of these 35 panels, 21 populate the description Link field with a wikipedia.org URL, meaning Google is anchoring the panel description to Wikipedia specifically. A further 4 anchors on books.google.com, corresponding to attorneys with book authorship on record; the remaining 10 populate no description Link at all. Table 2 gives the anchor distribution.

Anchor source	Count	% of Person Entity Panels
Wikipedia	21	60.0%
Google Books (books.google.com)	4	11.4%
No description-link populated	10	28.6%

Table 2: Canonical Description-Source Anchor Domains for Person Entity Panels (n = 35)

The prevalence of Wikipedia in this cohort is consistent with Google’s own documentation, which describes Knowledge Panel information as coming from “various sources across the web”. Wikipedia is the largest and most-editorially-verified open-web source about notable persons, and its info box structure surfaces exactly the biographical fields (Born, Education, Spouse) that Google’s Person-template placeholder claim 8 of US-11836177-B2, is engineered to consume.

National-Publisher Media Coverage

For a rank-matched control comparison we sampled 200 non-panel partners with a similar best-rank distribution and issued the same /news query for each. Person Entity Panel holders are 1.82x more likely to have at least one top-10 news item from a major national publisher than the no-panel control (59.5% versus 32.7%), while Local Business Panel holders are statistically indistinguishable from the no-panel control on this signal (33.3%). Table 3 gives the group-level breakdown.

Group	n	Mean items	% any national hit
No Panel (control)	199	8.03	32.7%
Local Business Panel	165	8.36	33.3%
Person Entity Panel	37	9.68	59.5%

Table 3: News-Coverage Metrics by Knowledge Panel Class. Top-10 Google News Results Per Partner

Raw top-10 volume differs less across the three groups because the /news endpoint caps at 10 items per query and both the Person-Entity-Panel and control cohorts' approach that cap; the differentiating signal is publisher tier, not raw count. The Local-Business-Panel cohort's indistinguishability from the no-panel control on national-media rate confirms the framework prediction of Section 3.5: business-class panels are driven by local-listing and review signals not by media notoriety.




Figure 2: National-Publisher News-Coverage Rate by Knowledge Panel Class

Biographical-Attribute Fingerprints and Alternative-Corpus Paths

The Wikipedia-derived biographical-attribute fingerprint is present in 88.6% of Person Entity Panels: at least one of Born, Died, Education, Spouse, Children, Parents, Siblings, or Height appears. This exceeds the direct Wikipedia touchpoint rate of 67.6% because Google's Person template can populate biographical attributes from a Wikipedia info box even when the panel's description text is drawn from a different source; the attributes propagate independently of the description link for the 12 Person Entity Panel holders in the sample without a Wikipedia touchpoint, three alternative-corpus paths account for recognition. Book authorship, evidenced by the Books attribute, is present in 14.3% of Person Entity Panels (5 of 35); Google Books serves as description Link in a comparable share. Political office, evidenced by Party, Previous office, Candidate for, or Office, is present in 17.1% (6 of 35). Awards or nominations attributes are present in a further 2.9%. A small number of single-attribute Person panels surface with only Education, corresponding to attorneys with dense regional media saturation whose entity signal is strong enough for a minimal Person template but not for the full Wikipedia-fed panel.




Figure 3: Signal Prevalence Across the Person Entity Panel Cohort, Ranked

Schema Markup Adoption Does Not Predict Person Entity Confidence

Cross-tabulating each schema-signal flag from the companion structured-data study (Hussain, 2026, 1,005-firm audit) against Person Entity Panel presence yields the lift table in Table 4. Lift is defined as the Person Entity Panel rate among firms with the signal divided by the rate among firms without.

Schema signal	P-KP rate WITH	P-KP rate W/O	Lift
has_attorney_person	3.47%	3.49%	0.99x
has_person_nested	5.13%	3.34%	1.54x
has_at_id	3.05%	4.78%	0.64x
has_sameas	2.83%	4.59%	0.62x
has_organization	3.76%	3.41%	1.10x
has_localbusiness	5.45%	3.24%	1.68x
has_review	3.01%	3.65%	0.82x
has_faqqpage	3.23%	3.54%	0.91x

Table 4: Schema-Signal Lift on Person Entity Panel Probability (n = 1,005)

Only one schema flag shows a positive lift on Person Entity Panel probability: Person nested inside Organization, at 1.54x. This is consistent with the grouped-related-entities mechanism described in the related-entities patent (Google LLC, 2016), which rewards a clean nested graph for downstream cluster traversal. The apparent 1.68x lift on has_localbusiness is driven by its association with the Local Business Panel class, not the Person Entity Panel class. The lack of positive lift on has_attorney_person alone, and the mildly negative lift on has_at_id and has_sameas, are not evidence that schema is inert. They are evidence that the on-page declaration must be corroborated by corpus-level signals (Section 5.2 through 5.4) to elevate the underlying Person Entity Confidence past the panel-triggering threshold.




Figure 4: Schema Signal Lift on Person Entity Panel Presence (n = 1,005)

Schema-Complete Firms Without Person Entity Panels: An Existence Proof

The most direct evidence that schema completeness alone does not produce a Person Entity Panel is the existence proof: firms that have implemented the full entity-signal schema stack and still have no Person Entity Panel for their managing partner. Table 5 lists the ten firms in the sample with the highest Schema Completeness Index (SCI) scores, filtered to those with an identified managing partner at high or medium confidence and no Knowledge Panel of any class.

Firm domain	Managing Partner	SCI	Wikipedia
schreckpersonalinjurylawyer.com	Fredrick Schreck	19	no
richardsonlawfirmnpc.com	Charles L. Richardson	17	no
cochranlaw.com	Terry L. Cochran	17	no
davis-adams.com	Chad Adams	17	no
walterslawwww.com	Ron D. Walters Jr.	17	yes
yelvertonlitigators.com	Matthew Yelverton	17	no
barkanmeizlish.com	Jason Cox	17	no
utahadvocates.com	Matthew Driggs	17	no
doctorspracticinglaw.com	Pat Cullan	17	no
gideonasen.com	Benjamin Gideon	17	no

Table 5: The Ten Schema-Complete Firms without a Person Entity Panel for the Managing Partner

All ten firms have @id and same as implemented; five have Person nested inside Organization, the one schema pattern with positive lift in Section 5.5. Only one (walterslawwww.com) has a Wikipedia article for the managing partner, and that firm still has no Person Entity Panel. The reading is not that schema is ineffective. The reading is that the on-page payload sits unconsumed when the corpus-level corroborated signal required by the confidence-threshold construct of Section 3.2 is absent. Section 6.1 develops this distinction as the central discussion point of the paper.

AI Visibility Propagates Through Person Entity Confidence

The qualitative audit of six queries across Google AI Overviews and a representative third-party generative answer engine confirms the framework prediction of Section 3.5. For the query “who is Ben Crump”, both surfaces return a confident biographical answer citing Wikipedia and multiple national-tier outlets (PBS, NPR, CNN, WBUR, The Hill). For the query “attorney who represented George Floyd family”, both surfaces name Ben Crump unhedged and support the claim with recent 2026 case coverage across nine independent outlets, demonstrating the compounding effect that a high-Person-Entity-Confidence attorney gathers over time. For the query “who is Michele S. Mirman lawyer”, the third-party engine cites Wikipedia by name in the visible in-answer citation tag, and Google AI Overview cites the firm site

plus additional national sources, both producing a structured biographical answer. The critical control is the query “who is Fredrick Schreck personal injury lawyer”, a schema-complete no-Person-Entity-Panel attorney from Table 5. The third-party engine’s answer opens with the hedged verb “appears to be” and cites only LinkedIn. Google AI Overview does synthesize an answer for this attorney, but every citation chip in the answer is a page of the attorney’s own firm website; no external editorial anchor is present. The full on-page schema implementation on the firm’s site does not surface in either AI answer surface’s citation set, consistent with the framework claim that AI answer engines retrieve from corpus-level Knowledge-Graph-adjacent substrates, not from on-page structured data. The generic commercial query “best personal injury lawyer in Miami” produced no Google AI Overview, and the third-party engine collapsed to directory-tier sources (Justia, Super Lawyers, Avvo), confirming that generic commercial queries do not surface Person Entity Panels or activate the entity-retrieval pathway.

Business Impact: Organic Traffic Correlated with Panel Class

Merging the monthly organic-traffic estimates described in Section 4.7 into the panel table yields the group-level descriptive statistics in Table 6. Given the heavy right-skew of the traffic distribution, medians are the honest headline figures.

KP class	n	Median	Mean	90th %ile
No Panel	785	440	1,657	3,375
Local Business Panel	164	290	1,175	2,690
Person Entity Panel	34	1,600	19,393	47,090

Table 6: Monthly Organic Traffic Distribution by Knowledge Panel Class

Person-Entity-Panel firms carry a median 1,600 monthly organic visits versus 440 for no-panel firms, a 3.64x lift on median firm-site traffic. Local-Business-Panel firms carry a median 290 monthly visits, a 34% deficit against the no-panel baseline. Both differences survive control for SERP rank: within every rank bucket, Person-Entity-Panel firms carry the highest median traffic and Local-Business-Panel firms carry the lowest.




Figure 5: Rank-Controlled Traffic Comparison by Knowledge Panel Class

Mann-Whitney U tests on raw traffic values are highly significant across all three comparisons: Person Entity Panel versus No Panel yields $p = 0.0003$, Local Business Panel versus No Panel yields $p = 0.0023$, and Person Entity Panel versus Local Business Panel yields $p = 8.7e-06$. A multivariable ordinary-least-squares regression on log10-transformed traffic ($n = 997$, $R^2 = 0.097$) confirms that both class coefficients survive control for rank, appearance frequency, and schema completeness. Table 7 summarizes.

Predictor	Coefficient	95% CI	Traffic multiplier
has_personal_kp	+0.470	[+0.050, +0.968]	x2.95
has_business_kp	-0.169	[-0.298, -0.058]	x0.68
log_appearances	+0.716	[+0.394, +1.034]	x5.20
sci_total	+0.030	[+0.016, +0.042]	x1.07 / SCI point
has_wiki_signal	+0.104	[-0.557, +0.637]	x1.27 (n.s.)
best_rank	-0.019	[-0.038, +0.001]	x0.96 / position

Table 7: OLS Regression of Log10(Organic Traffic + 1) On Panel Class and Confound Variables Bootstrap 95% CIs

Holding rank, appearance frequency, and schema completeness constant, a Person Entity Panel is associated with 2.95x more organic traffic and a Local Business Panel is associated with 32% less organic traffic. The Person-KP multiplier reflects the compounding of AI-Visibility-driven referral traffic and direct brand-query navigation for attorneys who have crossed the Person Entity Confidence threshold. The Local-KP deficit reflects the well-documented zero-click cannibalization pattern in which Google's Local Business Panel completes the searcher's task inside the SERP itself, leaving fewer clicks for the firm's website. The model R^2 of 0.097 indicates that the six predictors account for about 10% of traffic variance; the remainder is firm-level idiosyncratic factors (brand equity, paid media spend, PR events, content velocity) not captured in this dataset.




Figure 6: OLS Coefficient Plot with 95% Bootstrap CIs On Log-Traffic

Discussion

Corpus Signals vs On-Page Signals: The Fundamental Distinction

The findings of Section 5 map cleanly onto the framework of Section 3. Person Entity Confidence is driven by corpus-level corroborated signals (Wikipedia, national-tier media, biographical-attribute fingerprints from indexed sources) and only weakly associated with on-page structured-data signals. The one schema pattern with measurable positive lift, Person nested inside Organization, is the pattern that supports the grouped-related-entities graph traversal codified in Google LLC (2016). The existence proof of Section 5.6, in which ten schema-complete firms have no Person Entity Panel and only one of them has a Wikipedia article about the managing partner, establishes that on-page schema is necessary hygiene that conditions the panel-triggering payload without, on its own, driving the confidence-threshold crossing that surfaces a panel.

The Two-Class Panel Landscape as Structural, Not Ordinal

The three-class taxonomy of Section 3.1 and the empirical distribution of Section 5.1 together establish that Person Entity Panels and Local Business Panels are not ordinal positions on a maturity ladder. They are two structurally distinct outcomes produced by two templates (claims 8 and 9 of US-11836177-B2, Google LLC, 2013) that consume different

upstream signals. A firm can hold a mature Local Business Panel indefinitely without ever crossing into a Person Entity Panel, and the two classes carry sharply different implications for AI Visibility and, as Section 5.8 shows, for organic-traffic economics. The practitioner implication is that a firm targeting Person Entity Confidence for its managing partner should not treat further Local-Business-Panel optimization as progress toward that goal.

Wikipedia as the Universal Retrieval Substrate

The 60% direct anchor rate to Wikipedia in Person Entity Panel description links, combined with the 88.6% Wikipedia-derived biographical-attribute fingerprint and the AI-Visibility audit's explicit citation of Wikipedia by name for the Mirman query, points to Wikipedia as the single most load-bearing corpus source for entity recognition of personal-injury attorneys in this vertical. The corollary is that Wikipedia-related editorial work, subject to the notability requirements codified in Wikimedia Foundation (2024), sits at the highest-leverage tier of Brand Context Optimization. This is not a mechanism that a firm can manufacture directly. It is the outcome of substantive underlying practice (landmark cases, book authorship, elected office, sustained media coverage) that meets Wikipedia's independent-source requirements.

6.4 Zero-Click Cannibalization in the Local Business Panel The 32% traffic deficit observed for Local-Business-Panel firms in Section 5.8 is, to our knowledge, the first quantification of the zero-click cannibalization mechanism in the U.S. personal-injury legal vertical. The panel completes the user's immediate task inside the SERP itself: address renders in the panel, phone renders as a click-to-call button, directions render as a map link, hours render inline, and Google reviews render as a rating with review count. Each of these surface actions is a click that does not go to the firm's website. The Local-Business-Panel firm captures the intake channel (a phone call from a searcher) at the cost of the referral traffic (a page view followed by contact-form submission). Which trade-off a firm should prefer is a business-model question that this paper does not resolve. But firms should be aware that both trade-offs are being made, and that the panel is not a strictly additive benefit.

Practical Recommendations: A Brand Context Optimization Maturity Model

The recommendations that follow formalize the Brand Context Optimization (BCO) discipline introduced in Section 3.4 as a five-level maturity model. Each level maps to expected AI Visibility outcomes. Levels 1 and 2 are firm-controlled work. Level 3 is a partly firm-controlled, partly editorial-community-controlled outcome. Levels 4 and 5 are earned through substantive practice over years, not manufactured through short-term SEO tactics. The paper does not claim the model is a shortcut; it claims the model is the empirical shape of what Person Entity Confidence requires.

Level 1: Baseline Hygiene

Complete Legal Service schema on the firm entity; Attorney (or Person with job Title) schema on every named partner and lawyer; @id and same AS on every entity node; verified and completed Google Business Profile; consistent Name-Address-Phone data across the firm site, GBP, and all directory listings. Companion patents feeding this level: US-11836177-B2 template placeholders claims 8-10); US-8682913-B1 corroboration [12]; US-8312010-B1 local business ranking . Expected AI Visibility outcome: Local Business Panel becomes achievable; firm becomes retrievable in local commercial queries.

Level 2: Corpus Diversification

Verified profiles on the state bar member directory (U.S.), the Solicitors Regulation Authority "Find a solicitor" register (U.K.), or the provincial law society directory (Canada); populated profiles on the reputable jurisdictional legal directories (Avvo, Martindale-Hubbell, Super Lawyers in the U.S.; Chambers and Legal 500 in the U.K.; Lexpert and Best Lawyers Canada in Canada); Justia profile with case documents where applicable; LinkedIn, X, and Facebook profiles cross-linked to the firm site and to each other. Companion patents feeding this level: US-2023113420-A1 submitted-data accuracy prediction US-9411857-B1 grouping related entities .Expected AI Visibility outcome: entity signal strengthens across the retrieval substrate; firm and named partners appear in "People also search for" and related-entity clusters.

Level 3: Editorial Anchor

At least one high-quality, editorially-reviewed third-party profile of the individual. Candidates in ascending order of impact: bar-association feature article; Law360 or Bloomberg Law profile; book authorship listed in Google Books; Wikipedia article. Wikipedia is the highest-impact editorial anchor at this level (60% of Person Entity Panels in the sample use it as description Source). The Wikipedia route requires meeting the notability criteria in Wikimedia Foundation (2024): significant coverage in independent reliable sources, book authorship at a major press, elected or appointed public office, or landmark litigation. Expected AI Visibility outcome: Person Entity Panel becomes achievable; AI answer engines begin citing the individual by name with unhedged verbs.

Level 4: Recurring Media Pattern

Sustained pattern of national or major-regional media coverage: named case coverage, expert commentary quotes, thought-leadership placements. This is what the Section 5.7 AI-Visibility audit demonstrated for Ben Crump: nine independent outlets covering a single 2026 case is the recurring corpus fingerprint that keeps Person Entity Confidence alive over time. Companion patents feeding this level: US-9317559-B1 sentiment detection US-9477759-B2 question answering .Expected AI Visibility outcome: named citation across every major AI answer surface for both name queries and case-driven queries; presence in Google AI Overviews for personal-name queries [29-31].

Level 5: Signature Contribution

Book published at a major legal or trade press; elected or appointed public office; philanthropic gift that names an institution; landmark case that becomes a citable precedent; or founding of a professional organization. This is the level of the 3.5% cohort in the sample. Expected AI Visibility outcome: durable citation across all AI surfaces, including personal-name queries, case queries, thematic queries in the practice area, and academic queries about the subject of the signature contribution.




Figure 7: Brand Context Optimization (BCO) Maturity Model; Each Level Maps to Expected AI Visibility Outcomes

Limitations and Future Research

Five limitations should temper interpretation of the findings above and structure future work. First, the programmatic Knowledge Graph capture used in this study does not expose the website field on Person Entity Panels or the social Profiles field on any panels in our sample; direct on-page SERP inspection would be required for a complete field-level audit of what Google renders visually versus what programmatic capture surfaces. Second, news-volume metrics computed from top-ten Google news results for a partner's name are inflated by unrelated persons whenever the name is common; publisher-tier metrics rather than raw volume are therefore the more robust signal in Section 5.3.

Third, the six-query AI Visibility audit is qualitative rather than large-N and should be extended to a matched-sample study of AI Overview citation rates across the full 214 Knowledge Panel holders. Fourth, the traffic values in Section 5.8 are algorithmic third-party estimates rather than publisher-verified analytics; absolute values are approximate and relative comparisons within the same-source sample are the load-bearing claim. A cross-check against a publisher-verified analytics source with a brand-versus-non-brand traffic decomposition would tighten the magnitude estimates. Fifth, the findings are specific to the U.S. personal-injury vertical in mid-2026; replication in medical-malpractice defence, real-estate, and financial-services verticals in the U.S., and in the U.K. and Canadian PI verticals under their respective regulatory-body corpora, is the natural extension.

The paper's central framing precludes one common misreading: the absence of a Knowledge Panel for a given attorney is not evidence that no Person entity exists for that individual in Google's underlying Knowledge Graph. It is evidence that Person Entity Confidence for that individual has not crossed the visibility threshold at the moment of measurement. Future work with longitudinal capture (the same 1,005 firms measured across successive 12-month windows) would surface the rate at which attorneys move across the confidence threshold in either direction and the corpus events that accompany such moves.

Conclusion

The Google Knowledge Panel has become the visible surface of an entity-resolution and confidence-scoring machinery that now feeds Google's own AI Overviews as well as third-party generative answer engines. In the personal-injury legal vertical, the machinery separates the 1,005 top-ranking firms in our sample into three structurally distinct outcomes:

78.7% surface no panel, 16.4% surface a Local Business Panel keyed to the firm's Google Business Profile at the partner's name, and 3.5% surface a Person Entity Panel keyed to the managing partner as a recognized person entity in the Knowledge Graph. The Person-Entity-Panel cohort is not distinguished by superior on-page schema markup implementation; the schema-completeness signal is uncorrelated with panel presence at the individual-flag level and lifts only 1.54x for one specific pattern (Person nested inside Organization). The Person-Entity-Panel cohort is instead distinguished by corpus-level corroborated signals: a Wikipedia touchpoint at 67.6%, top-tier national-media coverage at 59.5% versus 32.7% for a rank-matched control, and, in the roughly one-third of the cohort without a Wikipedia article, alternative corpus paths through book authorship and elected office. Person Entity Confidence, the threshold construct we introduce, is the correct frame for reasoning about entity visibility in this vertical.

The business impact of the panel-class distinction is direct and material. Person-Entity-Panel firms carry 2.95 times the organic traffic of no-panel peers at the same SERP rank and schema completeness (Mann-Whitney $p = 0.0003$), while Local-Business-Panel firms carry 32% less traffic at the same rank tier, consistent with the zero-click cannibalization mechanism in which the Local Business Panel completes the searcher's task inside the SERP itself. Brand Context Optimization, the five-level maturity discipline formalized in Section 7, is the practitioner path that raises the corroborated corpus signals the Google patent record rewards. In the AI Search era, in which entity recognition at Google propagates simultaneously into Google AI Overviews and third-party generative answer engines, that discipline is no longer optional for firms whose managing partners aspire to be cited by name rather than by URL.

References

1. Clio. (2023). Legal Trends Report 2023. Themis Solutions Inc.
2. Hussain, B. (2026). Schema Markup Adoption in Top-Ranking Personal Injury Law Firm Websites: A Structured Data Audit of 1,005 Google Page-1 Sites Across 50 US States.
3. Singhal, A. (2012). Introducing the knowledge graph: things, not strings. *Official google blog*, 5(16), 3.
4. Google LLC. (2023a). Providing knowledge panels with search results. U.S. Patent No. US-11836177-B2.
5. Google LLC. (2024). Generative search assistant experience. AI Overviews product documentation. Google LLC.
6. Google Developers. (2024). Introduction to structured data markup in Google Search. Google LLC.
7. Berners-Lee, T., Hendler, J., & Lassila, O. (2001). Web semantic. *Scientific American*, 284(5), 34-43.
8. Bollacker, K., Evans, C., Paritosh, P., Sturge, T., & Taylor, J. (2008, June). Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data* (pp. 1247-1250).
9. Dong, X., Gabrilovich, E., Heitz, G., Horn, W., Lao, N., Murphy, K., ... & Zhang, W. (2014, August). Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 601-610).Google Developers.
10. Guha, R. V., Brickley, D., & Macbeth, S. (2016). Schema. org: evolution of structured data on the web. *Communications of the ACM*, 59(2), 44-51.
11. Google LLC. (2016). Personal knowledge panel interface. U.S. Patent No. US-9311362-B1.
12. Google LLC. (2014). Corroborating facts extracted from multiple sources. U.S. Patent No. US-8682913-B1.
13. Google LLC. (2023b). Predicting accuracy of submitted data. U.S. Patent Publication No. US-2023113420-A1.
14. Google LLC. (2012). Local business ranking using mapping information. U.S. Patent No. US-8312010-B1.
15. Google LLC. (2016). Sentiment detection as a ranking signal for reviewable entities. U.S. Patent No. US-9317559-B1.
16. Google LLC. (2017). Identifying local experts for local search. U.S. Patent No. US-9792330-B1.
17. Google LLC. (2018). Detection of related local entities. U.S. Patent No. US-9858291-B1.
18. Google LLC. (2015). Entity review extraction. U.S. Patent Publication No. US-2015112981-A1.
19. Google LLC. (2016). Grouping related entities. U.S. Patent No. US-9411857-B1.
20. Google LLC. (2016). Query suggestion templates. U.S. Patent Publication No. US-2014358940-A1.
21. Google LLC. (2016). Question answering using entity references in unstructured data. U.S. Patent No. US-9477759-B2.
22. Bing Team. (2017). Bring rich knowledge of people, places, things and local businesses to your apps. Bing Blogs, Microsoft Corporation.
23. Microsoft Corporation. (2024). Bing Entity Search API overview. Microsoft Learn.
24. Wikimedia Foundation. (2024). Wikipedia: Notability (people): Academics, lawyers, and legal professionals. Wikimedia Foundation.
25. Google LLC. (2015). Clustered search results. U.S. Patent Publication No. US-2015269231-A1.
26. Google LLC. (2016). Search result ranking and presentation. U.S. Patent No. US-9390174-B2.
27. Google LLC. (2019). Query suggestions based on entity collections of one or more past queries. U.S. Patent No. US-10360225-B1.
28. Hussain, B. (2026). Schema Markup Adoption in Personal Injury Law Firm Websites: A Systematic Analysis of Structured Data Implementation Across North American Legal Services. Available at SSRN 6551638.
29. Google LLC. (2016). Annotating query suggestions with descriptions. U.S. Patent Publication No. US-2016217181-A1.
30. Google LLC. (2016). Associating a search query with an entity. U.S. Patent Publication No. US-2016224621-A1.
31. Ilyas, I. F., Rekatsinas, T., Konda, V., Pound, J., Qi, X., & Soliman, M. (2022, June). Saga: A platform for continuous construction and serving of knowledge at scale. In *Proceedings of the 2022 international conference on management of data* (pp. 2259-2272).