

Epistara

NSS 2025 Authorship Attribution

Research Publication

*A Forensic Stylometric Analysis of the
National Security Strategy of the United States (Nov 2025)*

*Research Area: Forensic Stylometry & Institutional
Baselines*

Prepared by Epistara December 13, 2025

Epistara Personae Project Prepared Dec 13, 2025**URL:** <https://www.epistara.ai/research-publications>**Target document:** *National Security Strategy of the United States of America (November 2025)***Candidate set with individual training corpora:** Stephen Miller, Michael Anton, Elbridge Colby, Mauricio Claver-Carone, J.D. Vance, Marco Rubio. A generic institutional voice was explicitly added with its own generic training corpora.**Approach:** Supervised stylometric attribution with explicit institutional class and open-set rejection

Executive Summary

This report presents **two qualitatively different approaches** to the question "who contributed what?" to the recently published *National Security Strategy of the United States of America (November 2025)*. The two approaches analyzed are:

- 1) **Heuristic/Interpretive Attribution** (earlier heuristic analysis): A human-style analysis based on *content themes, rhetorical motifs, and role plausibility*. This method captures **policy influence** (who likely drove ideas) even when prose is staff-written.

Forensic Stylometric Attribution (current analysis) A text-internal method detecting **authorship signal** in the *surface form* of writing (character n-grams) after controlling for institutional baseline. This captures **writing hand/stylistic residue**, not policy influence. It uses Python analysis tools generated with assistance from GPT-5.2 and executed by the Epistara human author on a local Mac Studio M3 Ultra.

- 2) Comparison with Heuristic Attribution

Before constructing training corpora and running a supervised stylometric model, an **interpretive (heuristic) attribution** was also produced based on content alignment, rhetorical motifs, and role plausibility. That approach answers a *different question* from the forensic ML attribution run.

Side-by-Side Comparison

Author / Category	Heuristic Estimate	Forensic ML Estimate
Institutional / generic staff voice	~15-20%	51.64% + 26.82%
Stephen Miller	~17-21%	~0%
Michael Anton	~27-32%	4.36%
Elbridge Colby	~22-26%	4.04%
Marco Rubio	~9-12%	8.87%
Mauricio Claver-Carone	~4-6%	4.27%
J.D. Vance	~5-8%	~0%

Important: The heuristic column reflects *estimated influence or agenda-setting*, not detected authorship. The forensic column reflects *detectable stylistic authorship signal* in the final NSS prose.

What the Forensic Run Found

Contribution Estimates (Length-Weighted)

Label	% of NSS (chars)
Institutional (with training data)	51.64%
INSTITUTIONAL_UNKNOWN (untrained)	26.82%
Rubio (with personal training data)	8.87%
Anton (with personal training data)	4.36%
Claver_Carone (with personal training data)	4.27%
Colby (with personal training data)	4.04%
Miller (with personal training data)	~0%

High-Level Interpretation

- Roughly **52%** of the NSS looks like **generic interagency policy writing** (explicit institutional class)
- Another **27%** could not be confidently assigned to any single candidate (open-set rejection)
- The remaining **~22%** shows localized stylistic alignment with listed individuals, led by Rubio

Why There Is Little-to-No Miller Signal in Stylistic Analysis

1) Stylometry Detects "How It's Written," Not "Who Drove the Ideas"

This is the central point. A senior policy actor can exert very strong influence via outlining priorities, directing staff, approving or rejecting language, and insisting on concepts without leaving **detectable stylometric residue** in the final prose. Stylometry is much better at finding **hands** than **hands behind the hands**.

2) NSS Drafting Erases High-Variance Rhetorical Styles

In committee-authored, heavily edited documents: **distinctive voices** (especially rhetorical/polemical) tend to be smoothed out first, **institutional prose** survives, and **analytic/essayistic prose** sometimes survives in localized pockets. Miller's public-facing voice (often transcript/speech cadence, high rhetorical intensity) is precisely the kind of signal that editorial processes tend to flatten.

3) Miller's available corpus is structurally heterogeneous for stylometric comparison

Miller-type material typically includes press briefings, Q&A interactions, short punchy segments, and staff-mediated releases. These differ structurally from an NSS and raise within-author variance, reducing classifier confidence and increasing rejection into *institutional* or *INSTITUTIONAL_UNKNOWN*.

4) The Model Is Not "Confusing Topic with Author" (Good Sign)

After adding the institutional class, the model did **not** simply assign migration/border windows to "Miller" because of topical overlap. That is what you want in a defensible run: **style, not topic**.

5) Miller Influence May Be Absorbed into "Institutional" or "Unknown"

A practical reading: If Miller affected the document, that influence appears to have come through **institutional drafting channels** rather than through directly authored text. This is common in high-level strategy documents.

Why the Earlier View "Rated Miller Higher"

Earlier Heuristic Conclusion (What It Was Really Measuring)

The earlier narrative attribution leaned on **content alignment** (border/migration sovereignty emphasis), **rhetorical motifs** (civilizational framing, absolutist sovereignty language), and **process plausibility** (who would care and push). This tends to elevate Miller because he is strongly associated with immigration/border framing, sovereignty language, and demographic/cultural-security framing. But that does **not** imply he wrote large shares of the final text.

Forensic Model Conclusion (What It Is Measuring)

The forensic model measures stable stylometric fingerprinting (char n-grams), trained on corpora for each person, with an explicit institutional sink class and an open-set rejection rule. This approach can *down-rank* a person if their corpus is structurally different from the target (transcripts vs strategy document), their personal style was edited out, or their "signature" is more rhetorical than syntactic.

So: Who Should We Believe?

You should "believe" each method **for the question it can actually answer.**

If the Question Is...	Use This Method
Who likely drove the migration/sovereignty emphases?	Heuristic/role analysis
Whose writing hand is detectable in the final NSS prose?	Forensic stylometric result

A good final analytic stance is to treat them as complementary:

- **Stylometry:** conservative evidence about *authorship signal* in the final text.
- **Heuristic role+content analysis:** evidence about *policy influence and agenda-setting*, which can occur without authorship.

Best practice: Treat the forensic result as the more trustworthy answer to "authorship," and treat the heuristic result as a separate, useful answer to "influence."

Concretely:

- **Believe the forensic run** when you want to say: "This looks like it was written (stylistically) by X rather than Y."
- **Believe the heuristic analysis** when you want to say: "This reflects priorities closely associated with X, regardless of who wrote the prose."

If you collapse these two claims into one, you will overstate certainty.

"Under a stylometric attribution model trained on candidate corpora and an explicit institutional baseline, the 2025 NSS is predominantly institutional in voice. No segments exceeded the confidence threshold for stylistic signal from Stephen Miller. This does not imply lack of policy influence; rather, it suggests that any Miller influence was mediated through staff drafting and interagency editorial processes rather than reflected as a detectable authorial hand in the final text."

Appendix A: Methodological Details

A.1 Heuristic / Interpretive Method

The heuristic analysis was conducted **before** building training corpora or running machine learning. It relied on:

- **Content alignment:** matching themes (e.g., migration, sovereignty, alliances) to figures publicly associated with those issues
- **Rhetorical motifs:** civilizational language, absolutist framing, economic nationalism, or balance-of-power logic
- **Role plausibility:** who plausibly had the authority or portfolio to shape particular sections of an NSS
- **Comparative reading:** informal comparison to well-known writings, speeches, and policy positions

What this method measures well:

- Policy influence
- Agenda-setting
- Ideational fingerprints

What it cannot measure:

- Whether a person actually drafted or directly authored the final prose
- Whether stylistic features survived staff drafting and editorial smoothing

This method tends to **overestimate highly influential but rhetorically distinctive figures**, especially when their ideas are implemented by staff writers.

A.2 Forensic Stylometric ML Method

The forensic method uses a **supervised machine-learning stylometric pipeline**, with the following properties:

- **Training data:** curated per-author corpora (testimony, essays, speeches) plus an explicit *institutional* corpus (prior NSS, NDS, RAND reports)
- **Features:** character n-grams (3-5), which capture punctuation, affixes, spacing, and function-word morphology rather than topic
- **Model:** HashingVectorizer + SGDClassifier (logistic loss)
- **Segmentation:** 300-1200 character windows with overlap, applied consistently to training texts and the NSS
- **Open-set handling:** explicit *institutional* class and *INSTITUTIONAL_UNKNOWN* rejection for low-confidence windows

What this method measures well:

- Detectable authorship signal ("writing hand")
- Stylistic residue that survives editing
- Relative similarity in surface form

What it does not measure:

- Policy influence
- Who approved or insisted on particular ideas
- Who shaped the document indirectly through staff or process

This method tends to **underestimate senior figures whose influence is mediated through staff drafting**, and to allocate much of the text to institutional voice when editing is heavy. Variable-length windows were used to balance signal stability against boundary

effects in heavily edited policy prose, while maintaining consistent overlap across training and target texts.

A.3 How to Reconcile the Two

The two methods are **complementary, not competing**:

- Use the **heuristic method** to reason about *who likely drove ideas and priorities*.
- Use the **forensic ML method** to make conservative claims about *detectable authorship in the final prose*.

Where the two diverge (most notably for Stephen Miller), the divergence itself is analytically meaningful: it indicates **strong influence without surviving stylistic authorship**.