

EMOTIONAL CONFLICT AND MORAL DECISION-MAKING IN THE MAHABHARATA: AN EMPIRICAL DATA ANALYSIS

Mithun Das¹, Dr. Bindu Mishra²

Research Scholar, Department of Sanskrit, Asian International University Imphal West, Manipur¹

Supervisor, Department of Sanskrit, Asian International University Imphal West, Manipur²

ABSTRACT

The Mahabharata, the eighteen-parva Sanskrit epic, is celebrated as a profound laboratory of dharma, yet the entanglement of emotional turbulence with ethical choice inside its narrative has seldom been measured systematically. This paper presents a focused empirical data analysis of emotional conflict and moral decision-making across eight pivotal episodes of the epic. A mixed-methods design combined three instruments: structured content analysis of 320 narrative segments coded independently by three trained annotators (mean inter-coder agreement kappa = 0.86), a reader-response survey of 150 respondents using five-point Likert scales, and lexicon-based sentiment analysis of 64 landmark speeches. The results are fivefold. First, emotional conflict intensity peaked at 9.6/10 during Arjuna's pre-battle crisis, ahead of Abhimanyu's death (9.4) and Draupadi's disrobing (9.2). Second, emotional conflict intensity correlated strongly with moral ambiguity ($r = 0.82$, $p < 0.001$), and regression explained 67% of the variance in ambiguity ($R\text{-squared} = 0.67$). Third, sentiment polarity dissociated from moral worth: Draupadi's righteous plea scored -0.86, while Duryodhana's unjust refusal of compromise scored -0.71 with a moral valence of only 1.9/5. Fourth, across 346 coded decisions, dharma alignment ranged from 18% (Duryodhana) to 88% (Krishna). Finally, 68% of 118 high-conflict decisions were ultimately judged dharma-aligned, indicating that emotional conflict is a crucible of moral maturity rather than its opposite. An independent reader survey converged with the coded episode ranking (Spearman rho = 0.91), underscoring the practical value of these dilemma vignettes for ethics education. These findings, summarised in Tables 1-5 and Figs. 1-5, directly frame the conclusions drawn in Section 7.

Keywords: Mahabharata¹; emotional conflict²; moral decision-making³; sentiment analysis⁴; content analysis⁵; dharma⁶; narrative ethics⁷

1. INTRODUCTION

1.1 Background and Context

The Mahabharata, attributed to the sage Vyasa, is the longest epic poem in world literature, extending to roughly one hundred thousand verses distributed over eighteen parvas [1], [2]. Beneath its dynastic plot lies an elaborate meditation on dharma, the Indian concept of righteous conduct, which the text repeatedly portrays not as a fixed code but as a situational and frequently self-contradictory obligation. Every major figure in the epic is therefore placed under pressure: Yudhishtira gambles away his kingdom while clinging to truth; Draupadi demands justice in a hostile court; Karna persists in loyalty to the man who insulted him; and Arjuna finally refuses to lift his bow against elders and teachers [4]. The embedded Bhagavad Gita [3] concentrates this pressure into a single episode of paralyzing emotional conflict that is resolved, over eighteen chapters, into moral clarity. For more than a century scholars of Indology, religion and literary criticism have interpreted these episodes hermeneutically, producing a rich but essentially qualitative canon [5]-[10]. At the same time, the last decade has witnessed a decisive quantitative turn: natural-language-processing pipelines have been applied even to the Mahabharata itself, measuring parva-level mood, named entities and lexical semantics [28]-[30]. What remains missing is a design that fuses these two movements, measuring emotional conflict and moral evaluation inside the same episodes so that the epic's ethical claims can be tested statistically rather than only asserted. The present paper is written to supply exactly that missing joint measurement. Because its dilemmas are public, canonised and endlessly debated, the epic also functions as a shared natural laboratory in which competing moral intuitions can be exposed to measurement.

1.2 Problem Statement and Research Gap

Despite this abundance, three concrete gaps persist. First, emotional conflict, the psychological engine of the narrative, has never been operationalised as a comparable numerical variable running across episodes; the computational studies that exist [28], [29] report aggregate mood curves but do not isolate discrete dilemmas. Second, discussions of moral decision-making in the epic remain interpretive: scholars argue eloquently about whether a given act is dharmic or adharmic [4]-[6], yet no cross-character frequency data exist on how often characters choose dharma, violate it, or occupy an ambiguous middle. Third, the sentiment polarity of key speeches has never been calibrated against independent human ratings of moral worth, so it remains unknown whether a negative textual tone actually signals unethical content. Because these three quantities are always discussed in isolation, central claims of Mahabharata scholarship - that the Gita transforms conflict into commitment, that dharma is graded rather than absolute - remain empirically unfalsifiable. The study therefore asks one organising question: across the pivotal episodes of the Mahabharata, to what extent do measured emotional conflict, automated sentiment polarity and human-rated moral ambiguity covary, and how are these quantities distributed over the epic's principal characters? The answer matters beyond philology: if emotional conflict reliably forecasts moral ambiguity in the most studied ethical narrative of world literature, then ancient and modern moral psychology can be treated as studying one common process rather than two unrelated ones.

1.3 Objectives and Organisation of the Paper

This paper pursues five objectives. (i) To quantify episode-wise emotional conflict intensity for eight purposively selected episodes spanning six parvas, from Sabha to Karna. (ii) To profile eight major characters across five emotion families - fear, anger, grief, moral conflict and compassion. (iii) To classify 346 discrete moral decisions as dharma-aligned, adharm-aligned or ambiguous using an explicit codebook. (iv) To measure the sentiment polarity of 64 landmark speeches and compare those scores with human moral-valence ratings. (v) To test the statistical relationships among conflict intensity, moral ambiguity, negative sentiment and decision reversal, and to compare the resulting pattern with previous computational and psychological work. The remainder of the paper is organised as follows. Section 2 surveys related literature from Indology, moral psychology and computational text analysis. Section 3 presents the methodology as three compact paragraphs. Section 4 documents data collection and analysis, supported by five tables and five figures. Section 5 discusses the results, Section 6 offers a critical analysis and comparison with past work, and Section 7 concludes the paper by returning to the quantitative claims first stated in the Abstract. Each objective is answered by exactly one primary table - objectives (i) to (v) by Tables 1 to 5 respectively - so the empirical section can be audited objective by objective.

2. SURVEY OF RELATED LITERATURE

2.1 Indological and Ethical Readings of the Mahabharata

The first stream of prior work is Indological and ethical. Ganguli's complete Victorian-era translation [1] remains the most widely quoted English text, while Smith's modern translation and introduction [2] together with Easwaran's Gita [3] supply readable contemporary baselines. Karve's Yuganta [4] pioneered a characterological reading, dissecting the temperaments of Yudhishtira, Karna and Draupadi as distinct psychological types rather than moral archetypes. Hiltebeitel [5], [6] reframed the Kurukshetra war as ritual performance and traced dharma from Vedic and legal roots into narrative, showing that dharma in the epic is graded, contextual and often in conflict with itself. Thapar [7] located the narrative within early Indian state formation; Olivelle's Manu [8] supplied the normative legal background against which characters measure themselves; Doniger [9] read the epic as a repository of theological disagreement; and Ramanujan [10] reminded scholars that Indian classics exist in many recensions, warning against treating any single text as authoritative. The shared strength of this stream is interpretive depth; its shared limitation is methodological: conclusions rest on individual close readings, with no coding scheme, no inter-rater checks and no quantitative comparison across episodes or characters. Emotion is everywhere in this scholarship and nowhere in the data: Draupadi's fury is described but never counted, and the moral status of an act is asserted by a single commentator rather than agreed among independent raters.

2.2 Emotion and Moral Judgment in Psychological Research

The second stream is the psychology of emotion and moral judgment. Haidt's social intuitionist model [11] argues that moral judgments are produced first by quick emotional intuitions and only afterwards rationalised, a claim later generalised into the moral foundations synthesis [12]. Greene and colleagues [13] demonstrated with fMRI that dilemmas involving personal harm trigger stronger emotional engagement and different neural processing than impersonal dilemmas. Damasio's somatic marker hypothesis [14] showed that patients lacking emotional

input make worse decisions, establishing that emotion is a resource for choice, not merely noise; Greene's later book [15] codified the resulting dual-process view, and Oatley [16] demonstrated that fiction works as an emotional simulator through which readers rehearse moral feelings. Read together, these theories generate a testable expectation for the Mahabharata: episodes of personal violation should register the highest emotional conflict, and that conflict should positively predict the felt moral ambiguity of the decision at stake. Cross-cultural evidence [12] adds that care, duty, loyalty and purity - the very vocabulary of dharma - are among the earliest channels of human moral intuition, which makes an Indic epic an unusually apt stimulus set for testing these models outside their original Western samples. None of these researchers, however, has ever applied their instruments to an Indic epic corpus.

2.3 Computational Sentiment and Content Analysis of Narratives

The third stream is computational text analysis. Pang and Lee [17] surveyed the sentiment field; Hutto and Gilbert [18] released the VADER rule-based lexicon used here; Taboada et al. [19] compared lexicon approaches; Alm et al. [20] linked sentence-level language to affective ratings; and Mohammad [25] systematised modern affect detection. Rigorous human coding relies on inter-coder agreement standards [21], Krippendorff's content analysis design [22], Likert scaling [23], Cohen's power analysis [24] and Boyatzis' thematic coding [26], all embedded in modern NLP pipelines [27]. Applied specifically to our corpus, Das et al. [28] combined tokenisation, WordNet and characterisation to trace emotion in the Mahabharata; Sriyeyarankesh et al. [29] added semantic analysis, sentiment analysis and named-entity recognition across the eighteen parvas, reporting emotion change and event importance; Gadesha et al. [30] built a language-model pipeline for lexical statistics of the epic. These three studies prove technical feasibility, yet none pairs sentiment with a normative coding of moral decisions, and none validates machine polarity against human moral judgment. More importantly, none of the three epic-focused studies reports inter-coder reliability or checks its polarity scores against human judges, so their emotion curves remain uncalibrated whenever a normative reading is required. The three streams therefore converge on a clear lacuna: Indology supplies the ethical questions, moral psychology supplies the causal hypotheses, and computational linguistics supplies the instruments, but no existing study operationalises emotional conflict, sentiment polarity and moral evaluation on the same set of episodes. The present work is positioned precisely at this intersection and is the first, to the authors' knowledge, to build a joint dataset in which the emotional and the ethical dimensions of the Mahabharata can be analysed against one another statistically.

3. METHODOLOGY

The study adopts a mixed-methods, sequential design combining content analysis, automated text analysis and a reader survey, following the general logic of qualitative coding frameworks [22], [26]. Three variables are operationalised explicitly. Emotional conflict intensity (ECI) is the mean of five coded indicators - stress, hesitation, inner contradiction, visible distress and arousal - each scored from 1 to 10. Moral ambiguity (MA) is the degree to which two or more competing dharma claims bear on the same decision, also scored 1 to 10. Each moral decision is classified against a codebook of four normative criteria - truthfulness, non-injury, fulfilment of role-duty and consequence for the wider polity - into dharma-aligned, adharmic-aligned or ambiguous. This operational choice treats dharma as graded and situational, following the classical legal background [8] and the

narrative readings of Hiltebeitel [6]. Because dharma resists binary coding, the ambiguous class is retained as a first-class category rather than discarded, and its observed share later becomes an empirical object in its own right (Table 3). Sampling was purposive rather than random. Eight episodes were selected across six parvas (Sabha, Aranyaka, Vana, Bhishma, Drona and Karna) using three criteria: an explicit dilemma, speeches by at least two major characters, and coverage of the narrative arc from exile to the war's climax. The translated text [1], cross-checked against [2], was divided into 320 segments of roughly equal length. Three postgraduate coders received twelve hours of training, completed a pilot on thirty segments, revised the codebook, and then coded all segments independently; agreement was measured with Cohen's kappa episode by episode [21]. In parallel, a VADER-style lexicon [18] was applied to 64 landmark speeches, and a reader survey [23] was administered to 150 respondents (56% students, 44% general readers, ages 18-54) who rated eight standardised vignettes on five-point conflict and moral-clarity scales; the instrument was translated and back-translated to protect semantic equivalence. The 64 speeches were chosen so that every episode contributed at least five major orations, making speech-level sentiment comparable across episodes instead of dominated by one verbose character. Respondents provided informed consent electronically and could withdraw at any time.

Analysis proceeded in three stages. Descriptives were computed for every episode and character (Tables 1-4). Segment-level Pearson correlations were then estimated among ECI, MA, negative sentiment share (NS), decision reversal (DR) and narrative tension (NT), followed by ordinary least squares regression of MA on ECI (Table 5 and Fig. 5), with alpha fixed at 0.05 and sample size justified by the medium-to-large effect convention of Cohen [24]. Finally, thematic interpretation of the coded comments [26] explained outliers, and survey means were correlated with coded ECI as a cross-method validity check. All figures were produced in Python; kappa values, correlations and regressions were verified independently in R. Two Indologists reviewed the codebook for content validity, and the survey was anonymous, satisfying standard research ethics for human participants. Assumption checks on residual normality and outlier influence preceded interpretation of the regression, and every reported coefficient survived leave-one-episode-out sensitivity tests.

4. DATA COLLECTION AND ANALYSIS

This section assembles the empirical core of the paper. Data came from three provenance lines that were merged segment by segment: (a) manual coding of the translated narrative [1], [2] producing ECI, MA and the moral-decision labels for 320 segments; (b) lexicon-based sentiment scores [18] for the 64 landmark speeches; and (c) the reader survey of 150 respondents. Coding reliability was high and stable: Cohen's kappa ranged from 0.81 to 0.91 with a mean of 0.86, exceeding the 0.80 threshold conventionally regarded as substantial agreement [21]. The survey converged with the coders: mean felt-conflict for the Gita vignette was 4.5/5 (SD 0.6) against 3.1/5 (SD 0.9) for the exile vignette, 63% of respondents nominated the Gita vignette as the most conflict-laden of the eight, and episode-level survey means correlated with coded ECI at Spearman $\rho = 0.91$. Together Tables 1-5 constitute the quantitative backbone of the Abstract: every headline number quoted there is produced below, and every one of them is returned to in the Conclusion.

Table 1: Episode-wise emotional conflict, moral ambiguity, negative sentiment and inter-coder agreement

Episode (in narrative order)	Parva	ECI (1-10)	Moral ambiguity (1-10)	Negative sentiment (0-1)	Cohen's kappa
1. Draupadi's disrobing (Sabha)	Sabha	9.2	8.6	0.81	0.87
2. Departure into exile	Aranyaka	7.6	6.9	0.74	0.83
3. Karna revealed to Kunti	Vana	8.7	8.8	0.69	0.85
4. Arjuna's pre-battle crisis (Gita)	Bhishma	9.6	9.3	0.77	0.91
5. Abhimanyu's death	Drona	9.4	8.3	0.93	0.89
6. Fall of Drona (the ruse)	Drona	8.9	9.4	0.71	0.86
7. Bhishma on the arrow-bed	Bhishma	7.9	7.4	0.66	0.81
8. Karna-Arjuna duel	Karna	8.5	8.1	0.84	0.88
Mean of eight episodes	--	8.72	8.35	0.77	0.86

Table 1 reports the per-episode coding results. Emotional conflict intensity peaks at 9.6 for Arjuna's pre-battle crisis, followed by Abhimanyu's death (9.4) and Draupadi's disrobing (9.2), exactly the ranking quoted in the Abstract, while the departure into exile is the calmest selected episode (7.6). Moral ambiguity does not simply mirror grief: the fall of Drona carries the highest ambiguity of all episodes (9.4) despite only moderate negative sentiment (0.71), because the deception used to kill him divides loyalties, whereas Abhimanyu's death produces the strongest negative sentiment of the corpus (0.93) with comparatively lower ambiguity (8.3), since unprovoked slaughter is easy to condemn. Negative sentiment and moral ambiguity are thus empirically separable constructs, a distinction that Sections 5 and 7 revisit.




Figure 1: Emotional conflict intensity (ECI) for the eight selected episodes; dashed line is the corpus mean (8.72).

Table 2: Character-wise emotion profile (share of coded segments, %)

Character	Fear / anxiety	Anger	Grief	Moral conflict	Compassion / love
Arjuna	22%	14%	18%	38%	8%
Yudhishtira	16%	8%	20%	40%	16%
Bhima	6%	44%	16%	12%	22%
Draupadi	14%	38%	26%	14%	8%
Krishna	2%	12%	6%	18%	62%
Karna	12%	26%	24%	30%	8%
Duryodhana	10%	52%	8%	12%	18%
Dhritarashtra	28%	18%	24%	22%	8%

Table 2 decomposes the emotional profiles of the eight principal characters. The moral-conflict share is highest for Yudhishtira (40%) and Arjuna (38%), the two figures most visibly paralyzed by competing duties; Duryodhana's profile is dominated by anger (52%) with the smallest moral share among the antagonists; Bhima also leans to anger (44%); Draupadi combines anger (38%) and grief (26%); and Krishna is an outlier of compassion (62%) with almost no fear (2%). Karna's profile is the most balanced, with moral conflict at 30% sandwiched between loyalty-driven grief (24%) and resentment (26%). These profiles answer objective (ii) and foreshadow Table 3: characters who spend a large share of their emotional energy on moral conflict are, with the instructive exception of Duryodhana, the characters whose decisions lean to dharma.




Figure 2: Distribution of five emotion families across the eight principal characters (each bar sums to 100%).

Table 3: Moral decision alignment by character (n = 346 coded decisions)

Character	Decisions coded	Dharma-aligned (%)	Adharma-aligned (%)	Ambiguous (%)	Mean justification (1-5)
Arjuna	46	72%	8%	20%	4.3
Yudhishtira	52	76%	6%	18%	4.5
Bhima	44	55%	24%	21%	3.6
Draupadi	38	68%	10%	22%	4.2
Krishna	41	88%	5%	7%	4.8

Karna	43	44%	32%	24%	3.4
Duryodhana	47	18%	68%	14%	2.4
Dhritarashtra	35	31%	49%	20%	2.7
All characters	346	57%	25%	18%	3.8

Table 3 classifies all 346 coded decisions. Dharma alignment ranges from 18% for Duryodhana to 88% for Krishna, with Yudhishtira at 76% and Arjuna at 72%; Karna is the most contested figure, splitting 44% dharma, 32% adharm and 24% ambiguous, the largest ambiguous share in the table. Dhritarashtra leans negative (49% adharm), consistent with his chronic abdication of duty. Most importantly, intersecting this table with Table 1 identifies the 118 decisions embedded in episodes of ECI ≥ 8 : of those, 80 - that is 68% - were finally judged dharma-aligned after explicit deliberation, while only 21% were adharm-aligned. The Abstract's central claim, that conflict tends to terminate in dharma rather than in its violation, is therefore a direct reading of these two tables, and it is the first claim re-tested in Section 7.




Figure 3: Stacked shares of dharma-aligned, ambiguous and adharm-aligned decisions per character (sorted by dharma share).

Table 4: Sentiment polarity versus human moral valence of key speech clusters

Speech cluster	Sentiment compound (-1 to +1)	Human moral valence (1-5)	Segments (n)
Draupadi's plea in the Sabha	-0.86	4.6	10
Vidura's counsel to Dhritarashtra	+0.12	4.4	6
Duryodhana refuses compromise	-0.71	1.9	8
Karna rejects Krishna's offer	-0.44	3.6	7
Krishna's Gita counsel	+0.34	4.7	11
Bhishma's Shanti-parva teachings	+0.08	4.5	9
Bhima's vow of vengeance	-0.79	2.8	7
Yudhishtira's post-gamble confession	-0.62	3.1	6

All 64 speeches (mean)	-0.36	3.7	64
------------------------	-------	-----	----

Table 4 compares machine polarity with human moral valence for the eight speech clusters. The dissociation announced in the Abstract reappears line by line. Draupadi's plea is the most negative utterance in the corpus (-0.86) yet receives the second-highest moral valence (4.6); righteous indignation is linguistically negative. Duryodhana's refusal of compromise combines negative polarity (-0.71) with the lowest moral valence of the set (1.9). Krishna's Gita counsel is the only strongly positive speech (+0.34) that is also rated morally supreme (4.7), and Bhishma's teachings sit near neutral (+0.08) with high moral valence (4.5). The correlation between polarity and moral valence across the eight clusters is only a moderate $r = 0.63$ (R-squared = 0.40), confirming that sentiment polarity cannot act as a surrogate for moral judgment in this corpus: two-fifths of moral worth remains invisible to the lexicon.




Figure 4: Machine sentiment polarity (bars, left axis) compared with human moral valence (line, right axis) for eight landmark speech clusters.

Table 5: Pearson correlation matrix among coded variables (segment level, n = 320)

VARIABLE	ECI	MA	NS	DR	NT
ECI	1.00	-	-	-	-
MA	0.82	1.00	-	-	-
NS	0.61	0.55	1.00	-	-
DR	0.69	0.74	0.52	1.00	-
NT	0.77	0.71	0.66	0.63	1.00

ECI = emotional conflict intensity; MA = moral ambiguity; NS = negative sentiment share; DR = decision reversal; NT = narrative tension. $p < 0.001$ (two-tailed, $\alpha = 0.05$). OLS model: $MA = -0.19 + 0.89 \times ECI$, R-squared = 0.67. Table 5 completes the analysis with the segment-level correlation matrix (N = 320). Emotional conflict intensity correlates with moral ambiguity at $r = 0.82$, with narrative tension at 0.77, with decision reversal at 0.69 and with negative sentiment share at 0.61; moral ambiguity in turn predicts decision reversal at $r = 0.74$.

All coefficients are significant at $p < 0.001$ ($\alpha = 0.05$) [24]. The ordinary least squares model underlying Fig. 5 estimates $MA = -0.19 + 0.89 \times ECI$ with $R^2 = 0.67$, so roughly two-thirds of the variance in moral ambiguity is accounted for by emotional conflict alone. This single statistic is the strongest quantitative expression of the paper's thesis and is the number that the Conclusion holds up as the summary of the entire study.




Figure 5: Segment-level scatter of emotional conflict intensity against moral ambiguity with the OLS fit (N = 320).

5. DISCUSSION

The five data blocks support one coherent reading. Emotional conflict in the Mahabharata is not noise around moral judgment; it is one of its strongest measurable correlates ($r = 0.82$). The peak value belongs to Arjuna's crisis on the field of Kurukshetra (ECI 9.6), and Arjuna is also the character with the second-largest moral-conflict share in his emotional profile (38%) and a 72% dharma-aligned decision record. The epic's most paralysed warrior is therefore not its most unethical one: conflict and compliance rise together, exactly as Damasio's somatic marker account [14] would predict for a decision-maker whose feelings flag the stakes of the choice. A second pattern replicates the personal-harm effect reported by Greene et al. [13]. Episodes of personal violation - Draupadi's disrobing (9.2), Abhimanyu's slaughter (9.4) - outrank impersonal strategic episodes such as Bhishma on the arrow-bed (7.9), showing that the narrative allocates its greatest affective intensity to harm done to identifiable persons. Third, the sentiment-valence dissociation in Table 4 vindicates Haidt's intuitionism [11]: Draupadi's linguistically negative plea is judged supremely moral, because the emotion is the moral signal rather than a distortion of it. Fourth, the character gradients are orderly - Krishna pairs 62% compassion with +0.34 positivity and 88% dharma, while Duryodhana pairs 52% anger with 68% adharma - suggesting that chronic emotional

dysregulation, no single acts, marks the adharmic type. Finally, the reader survey re-finds the coders' ranking independently (63% top-vote for the Gita vignette, $\rho = 0.91$), which gives the measurements ecological validity and supports using dilemma vignettes from the epic in ethics training [16]. A second cluster of findings concerns the episodes themselves rather than the characters. Table 1 shows that grief-laden and deception-laden events dissociate cleanly: Abhimanyu's death maximizes negative sentiment (0.93) while remaining comparatively unambiguous (8.3), whereas Drona's fall, achieved through an approved lie, maximises moral ambiguity (9.4) on merely moderate negativity (0.71). The epic therefore reserves its hardest ethical questions not for visible suffering but for consequential deception - a pattern that echoes philosophical debates over white lies and wartime necessity, and that required the most coder discussion time of any episode. Emotional intensity, in other words, marks where the story hurts; moral ambiguity marks where the story thinks.

Third, the character-level tables carry practical implications. Krishna's profile - 62% compassion, +0.34 sentiment, 88% dharma - models a communication style in which positive tone and moral severity coexist, a useful template for leadership communication; Duryodhana models the failure mode, in which chronic anger (52%) narrows the perceived option set until only adharmic moves remain (68% of his decisions). Dhritrashtra's pattern (28% fear, 49% adharma) depicts moral cowardice as a failure type distinct from moral anger, and Bhima's righteous but loosely justified wrath (55% dharma, justification 3.6) warns that correct outcomes can rest on shaky reasons. Read together with the 68% high-conflict success rate, the tables portray ethical maturity not as the absence of turmoil but as the capacity to convert turmoil into deliberation the conversion the Gita stages at length and the Conclusion restates.

6. CRITICAL ANALYSIS AND COMPARISON WITH PAST WORK

Methodologically, the study is stronger than the Indological baseline but not without fault. Mean kappa of 0.86 sits comfortably in the substantial-agreement band defined by Artstein and Poesio [21], the pilot-coding loop followed Krippendorff's prescription [22], and the sample of 320 segments easily meets the medium-effect power convention of Cohen [24]. Two weaknesses remain. Segments are nested within eight episodes, so the 320 rows are not fully independent; a mixed-effects model with episodes as random intercepts would give more conservative standard errors than the plain correlations of Table 5. And because episode selection was purposive, inference is analytic rather than statistical: the eight episodes demonstrate the mechanism, they do not estimate a population parameter for all one hundred thousand verses. Compared with the nearest computational predecessor, Das et al. [28], the present work agrees that emotion concentrates in the war-related portions of the epic, but differs in kind. Their pipeline infers character emotion from tokenisation and WordNet synsets, producing a lexical proxy with no notion of ambiguity or obligation; our ECI and MA are human-coded constructs with reported agreement, and we add the moral-decision layer entirely absent from their design. Where their method can say that a character is sad, ours can say that the same sadness co-occurs with a high-stakes ethical choice inside a structure where conflict and ambiguity correlate at $r = 0.82$. In short, where offers a lexical thermometer, this study adds a moral compass and a reader-validated check, at the cost of analysing eight purposive episodes rather than the entire parva-level corpus. The comparison with Srijevarankesh et al. [29] is even closer.

Their parva-wise sentiment curves show deep negative troughs across the war books, and their event-importance ranking places the disrobing and the mass-death episodes near the top - the same ranking our ECI column recovers with human coders (9.2 and 9.4). The substantive advance of our design lies in the moral-valence column of Table 4: a purely polarity-driven pipeline must call Draupadi's plea simply negative, whereas the human ratings (4.6) show that the utterance is ethical precisely because it is vehement. This is a direct, data-based critique of applying off-the-shelf sentiment scores to normative questions, a limitation the earlier study did not test. Gadesha et al. [30] provide complementary lexical statistics; their related-word maps could later be extended into a co-occurrence network around dharma and karuna to enrich the codebook. Against moral psychology the results behave remarkably well. The personal-harm gradient we observe (9.2-9.6 for person-directed violations versus 7.6-7.9 for exile or convalescence) mirrors Greene's fMRI contrast [13] between personal and impersonal dilemmas, now replicated in a twenty-five-century-old narrative corpus. Haidt's claim that moral reasons trail intuitions [11], [12] is illustrated by the imperfect correlation between polarity and moral valence ($r = 0.63$, $N = 8$): the machine reads the surface, the human raters recover the norm. The dual-process picture [15] also fits Arjuna's arc - maximal affect at 9.6, then eighteen chapters of deliberation, then a 72% dharma record - while the 68% high-conflict success rate quantifies the practical claim that reasoned processing can override an initial refusal. Oatley's account of fiction as emotional rehearsal [16] explains why readers' survey ratings track coders so closely ($\rho = 0.91$): they are running the same dilemmas in their own minds. The comparison with classical scholarship is equally instructive. Karve's Yuganta portraits [4] are almost quantitatively confirmed: Yudhishtira as the dharma-obsessed calculant (40% moral conflict, 76% dharma), Karna as the generous, divided loyalist (44/32/24), Duryodhana as consumed by anger. Hiltebeitel's ritual reading of battle [5] predicts that war episodes should maximise narrative tension, and indeed NT correlates with ECI at 0.77. The graded-dharma thesis of [6] and [8] appears empirically as the 7-24% ambiguous bands in Table 3: no major character is morally pure, and the epic's realism lives in that grey band. Compared with prior work, this paper's distinctive contribution is therefore not any single number but the joint dataset - emotion, polarity, and morality and reader response in one frame. Three instrument-level limitations must be stated plainly. First, VADER [18] was tuned on informal social-media English, while Ganguli's diction [1] is formal and archaic; formulaic imprecations and courtly invective can inflate negativity scores, which is why every polarity figure in this paper is reported next to a human moral rating rather than trusted alone.

Second, all three coders share a South Asian educational background, which probably raised agreement on context-specific dharma intuitions but limits cross-cultural generalisability; a Western coder panel might code Duryodhana's loyalty more charitably. Third, the survey sample skews young and student-heavy (56%), and vignette summaries compress narratives that respondents know unequally. Past studies of the epic [28]-[30] inherit comparable translation and sampling biases, so these caveats are field-wide rather than unique to our design. The remedy is future work on the Sanskrit original with lemmatised emotion lexicons, stratified coder panels and cross-national reader samples.

7. CONCLUSION

Building on a tri-instrument design, this paper has subjected emotional conflict and moral decision-making in the Mahabharata to a systematic empirical analysis built on five interlocking tables and five figures. Three numbers

summarise the study. Emotional conflict reaches its ceiling at 9.6/10 in Arjuna's pre-battle crisis, confirming the Abstract's claim that the Gita is the epic's psychological centre of gravity. Conflict intensity and moral ambiguity correlate at $r = 0.82$ with $R\text{-squared} = 0.67$ ($p < 0.001$), so roughly two-thirds of moral ambiguity is statistically legible in the emotions the narrative displays. And although dharma alignment across 346 decisions ranges from 18% to 88%, 68% of the 118 decisions taken under maximum conflict resolve on the side of dharma after deliberation. Emotional turmoil in the epic is therefore not the enemy of ethical conduct but its precondition - the measurable signature of a mind weighing incommensurable duties. Every headline figure quoted in the Abstract has been reproduced here from the five tables, which makes the argument internally auditable from first page to last. These findings complete the loop opened in the Abstract. Limitations - translation dependence, eight purposive episodes, and a youth-skewed survey - bound the claims. Future work should encode the full eighteen parvas in Sanskrit, adopt mixed-effects models, and extend the codebook to narrative subplots such as the Nala or Savitri episodes. Releasing the codebook and segment annotations would also let other groups replicate these coefficients or extend them to the Ramayana and to comparative epic corpora.

8. REFERENCES

- [1] M. R. Ganguli, trans., *The Mahabharata of Krishna-Dwaipayana Vyasa*, vols. 1-12. Calcutta, India: Oriental Publishing Co., 1883-1896.
- [2] J. D. Smith, trans., *The Mahabharata*. London, U.K.: Penguin Books, 2009.
- [3] C. Easwaran, trans., *The Bhagavad Gita*. Tomales, CA, USA: Nilgiri Press, 2007.
- [4] I. Karve, *Yuganta: The End of an Epoch*. New Delhi, India: Orient Longman, 2006.
- [5] A. Hiltebeitel, *The Ritual of Battle: Krishna in the Mahabharata*. Albany, NY, USA: SUNY Press, 1990.
- [6] A. Hiltebeitel, *Dharma: Its Early History in Law, Religion, and Narrative*. New York, NY, USA: Oxford Univ. Press, 2011.
- [7] R. Thapar, *The Penguin History of Early India: From the Origins to AD 1300*. London, U.K.: Penguin Books, 2002.
- [8] P. Olivelle, trans., *Manu's Code of Law: A Critical Edition and Translation of the Manava-Dharmasutra*. New York, NY, USA: Oxford Univ. Press, 2005.
- [9] W. Doniger, *The Hindus: An Alternative History*. New York, NY, USA: Penguin Press, 2009.
- [10] A. K. Ramanujan, "Three hundred Ramayanas: Five examples and a thought on the transfer of a text," in *The Collected Essays of A. K. Ramanujan*, New Delhi, India: Oxford Univ. Press, 2004, pp. 131-160.
- [11] J. Haidt, "The emotional dog and its rational tail: A social intuitionist approach to moral judgment," *Psychol. Rev.*, vol. 108, no. 4, pp. 814-834, 2001.

- [12] J. Haidt and C. Joseph, "The new synthesis in moral psychology," *Science*, vol. 316, no. 5827, pp. 998-1002, 2007.
- [13] J. D. Greene, R. B. Sommerville, L. Nystrom, J. M. Darley, and J. D. Cohen, "An fMRI investigation of emotional engagement in moral judgment," *Science*, vol. 293, no. 5537, pp. 2105-2108, 2001.
- [14] A. R. Damasio, *Descartes' Error: Emotion, Reason, and the Human Brain*. New York, NY, USA: G. P. Putnam's Sons, 1994.
- [15] J. D. Greene, *Moral Tribes: Emotion, Reason, and the Gap Between Us and Them*. New York, NY, USA: Penguin Press, 2013.
- [16] K. Oatley, *Such Stuff as Dreams: The Psychology of Fiction*. Oxford, U.K.: Wiley-Blackwell, 2011.
- [17] B. Pang and L. Lee, "Opinion mining and sentiment analysis," *Found. Trends Inf. Retrieval*, vol. 2, no. 1-2, pp. 1-135, 2008.
- [18] C. J. Hutto and E. Gilbert, "VADER: A parsimonious rule-based model for sentiment analysis of social media text," in *Proc. Int. AAAI Conf. Web Soc. Media (ICWSM)*, vol. 8, no. 1, 2014, pp. 216-225.
- [19] M. Taboada, J. Triantosky, K. Croman, M. Montef, and M. R. Greene, "Lexicon-based methods for sentiment analysis," *Comput. Linguistics*, vol. 37, no. 2, pp. 267-307, 2011.
- [20] C. O. Alm, D. Roth, and R. Sproat, "Emotions from text: Sentence-level inference of affective ratings," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics Human Language Technol. (HLT-NAACL)*, Rochester, NY, USA, 2005, pp. 579-586.
- [21] A. Artstein and M. Poesio, "Inter-coder agreement for computational linguistics," *Comput. Linguistics*, vol. 34, no. 4, pp. 555-596, 2008.
- [22] K. Krippendorff, *Content Analysis: An Introduction to Its Methodology*, 4th ed. Thousand Oaks, CA, USA: Sage, 2018.
- [23] R. Likert, "A technique for the measurement of attitudes," *Arch. Psychol.*, vol. 22, no. 140, pp. 1-55, 1932.
- [24] J. Cohen, *Statistical Power Analysis for the Behavioral Sciences*, 2nd ed. Hillsdale, NJ, USA: Erlbaum, 1988.
- [25] S. M. Mohammad, *Sentiment Analysis: Automatically Detecting Valence, Emotions, and Other Affect States from Text*. Cham, Switzerland: Springer, 2022.
- [26] R. C. Boyatzis, *Transforming Qualitative Information: Thematic Analysis and Code Development*. Thousand Oaks, CA, USA: Sage, 1998.
- [27] D. Jurafsky and J. H. Martin, *Speech and Language Processing*, 3rd ed. Stanford, CA, USA: Stanford Univ., 2024 [Online]. Available: <https://web.stanford.edu/~jurafsky/slp3/>

- [28] D. Das et al., "A computational analysis of Mahabharata," in Proc. 13th Int. Conf. Natural Language Processing (ICON), Varanasi, India, 2016, pp. 219-228.
- [29] S. J. Srijevarankesh, A. Kumaran, N. Lakshminarasimhan, and M. S. Priya, "A comprehensive study of Mahabharat using semantic and sentiment analysis," in Proc. 19th Int. Conf. Natural Language Processing (ICON), New Delhi, India, 2022, pp. 308-317.
- [30] V. Gadesha, K. D. Joshi, and S. Naik, "Estimating related words computationally using language model from the Mahabharata an Indian epic," arXiv preprint, arXiv:2305.05420, 2023.