Inpatient Rehabilitation Facility Prospective Payment System Reform: Primary Diagnosis Technical Memo

February 2026

Acumen, LLC
500 Airport Blvd., Suite 365
Burlingame, CA 94010

OVERVIEW

This memo summarizes the IRF PPS Reform: Primary Diagnosis Evaluation analysis exploring alternative payment approaches to the primary diagnosis portion of the current Inpatient Rehabilitation Facility (IRF) Prospective Payment System (PPS). Section 1 of this document provides the policy context and motivation for the analysis. Section 2 outlines the methodological approach. Section 3 summarizes the main results. Finally, the Appendix contains figures and tables referenced throughout the memo.

TABLE OF CONTENTS

LIST OF TABLES AND FIGURES

[This page is intentionally left blank.]

1 BACKGROUND

CMS is exploring options to modernize the IRF PPS by moving away from using the current Impairment Group Code (IGC) and Rehabilitation Impairment Category (RIC) system and leveraging existing classification models such as the Clinical Category (CC) system under the Skilled Nursing Facility Patient-Driven Payment Model (SNF PDPM). The primary objective of the Primary Diagnosis Evaluation analysis is to simplify the diagnosis classification and improve payment accuracy while addressing issues related to the incomplete mapping of the International Classification of Diseases, version 10 (ICD-10) codes to IGCs and RICs.

The current IRF PPS requires providers to report an IGC and corresponding etiologic diagnoses for a patient on the IRF Patient Assessment Instrument (IRF-PAI), after which the IGC is mapped to a RIC. One of the main challenges of the current system is the complexity in classification, requiring a three-stage mapping process, which can be mitigated by simplifying the classification to a two-stage mapping process. Another major challenge of the current system is the ambiguity created by the potential use of up to three ICD-10 codes in the etiologic diagnosis field, which complicates the identification of the true reason for an IRF admission. This complexity can lead to unreliable reporting and potential misalignment between the diagnosis and patient care, as providers may be incentivized to adjust the coding of the etiologic diagnosis to justify their choice of IGC. As a result, there is a need to create a more transparent and efficient model for diagnosing and classifying IRF patients. To address these issues, Acumen conducted the IRF Primary Diagnosis Evaluation analysis to assist CMS in aligning aspects of the IRF PPS with the SNF PDPM. The analysis focuses on assigning PDPM CCs based on IRF primary diagnoses.

In aligning the two payment systems, Acumen considered several key similarities and differences between the two settings. Both IRFs and SNFs provide post-acute care (PAC) treatment that address therapy and nursing needs. However, IRFs provide more intensive therapy while SNFs focus more on skilled nursing care. The two systems also differ in how they report primary diagnoses – IRF stays report etiologic diagnoses and IGC on IRF-PAI items 22A/B/C and 21 respectively, whereas SNF stays report a single ICD-10 based primary diagnosis through item I0020B on the Minimum Data Set (MDS). Moreover, IRFs do not require a prior 3-day inpatient stay for admission, unlike SNFs, which suggests that the case mix of patients in both settings might differ given that IRF stays can be admitted from the community as well. Additionally, the payment model also differs between the two as IRFs are paid on a per stay basis, covering all costs, while SNFs follow a per-diem payment model, which break costs into six components (Physical Therapy, Occupational Therapy, Speech-Language Pathology, Non-Therapy Ancillary, Nursing, and Non-Case-Mix).

2 METHODOLOGY

Section 2 highlights the methodology for the Primary Diagnosis Evaluation analysis broken down by the relevant topics. Subsection 2.1 outlines the study population, subsection 2.2 describes the process of cost standardization, subsection 2.3 details the methodological approach used to assign PDPM clinical categories based on IRF primary diagnoses, subsection 2.4 outlines the primary diagnosis model evaluation criterion, and subsection 2.5 identifies the major methodological deviations from the SNF PDPM methodology.

2.1 Study Population

The study population for this analysis includes all Medicare Part A IRF stays in FY 2023 after the necessary study restrictions have been applied. The restrictions are implemented to ensure the results of the analysis would be applicable to the general Medicare fee-for-service (FFS) IRF population. The restriction types include Medicare Part A restrictions, matching restrictions, validity restrictions, atypical stays restrictions, and primary diagnosis restrictions. The detailed breakdown of the study restrictions are provided in Table 1, which summarizes the frequencies of each restriction applied during the selection process.

2.2 Standardizing Cost

For each provider, Acumen calculates cost center-specific cost-to-charge ratios (CCRs) using data from Healthcare Provider Cost Reporting Information System (HCRIS). Next, for each stay, Acumen estimates cost-center-specific costs by multiplying the most relevant CCR to each cost-center-specific charge on the IRF claim. Then, Acumen sums up the costs for each cost center to estimate IRF stay total costs. Costs are standardized by removing the rural, teaching status, low-income proportion, and wage index/labor share adjustment.

2.3 Assigning IRF Stays to Clinical Categories

Acumen uses the following hierarchical approach to assign CCs to IRF stays based on the availability of a prior Inpatient (IP) stay or the IRF stay's etiologic diagnoses:

- For stays with a prior IP stay (within a 60-day lookback window from IRF admission), Acumen assigns a PDPM clinical category and an IRF clinical category using the prior IP stay's principal diagnosis on claims by applying the FY 2023 PDPM ICD-10 Clinical Category Mappings.
- For stays without a prior IP stay (within a 60-day lookback window from IRF admission), Acumen assigns a PDPM clinical category and an IRF clinical category to each of the three existing etiologic diagnoses (i.e., IRF PAI Items 22A/B/C), whenever available, by applying the FY 2023 PDPM ICD-10 Clinical Category Mappings.
    - For stays without a prior IP stay within the lookback window and with multiple clinical categories (i.e., multiple etiological diagnoses): Acumen selects the category assigned to the first etiological diagnosis (i.e., IRF PAI Item 22A), unless the remaining etiologic

diagnoses have a better match with the recorded admission IGC on the IRF-PAI, in which case, Acumen selects that corresponding category.

During the development of SNF PDPM, CMS leveraged the claim's principal diagnosis on the prior proximal IP stay to ensure accuracy in the reason for treatment. Along those lines, Acumen found that approximately 88% of IRF stays in the study population have a prior IP stay within the 60-day lookback window. The remaining 12% of the stays rely on the etiologic diagnosis as a classification mechanism under the primary diagnosis assignment hierarchy detailed in subsection 2.3 above. Due to the substantial proportion of IRF stays with a prior IP stay, the reliance on ambiguous IRF-PAI data to prospect payment was reduced.

After assigning stays to their appropriate SNF PDPM CCs (11 in total, 10 valid categories plus a "Return to Provider" category), it became evident that stays were being categorized disproportionately into specific CCs, as highlighted in Figure 1, due to the differing patient case mix in the IRF setting. To mitigate the concerns over the distribution patterns, Acumen divided the CCs with proportionally higher stays into more granular CCs (See detailed breakdown of CCs under subsection 2.5) based on the data and clinical input to ensure a more even distribution. As a result, this process yields 17 uncollapsed IRF CCs.

Acumen then uses the 17 uncollapsed IRF CCs and runs a simplification process through the Classification and Regression Tree (CART) method, using the model "Total Cost per Stay ~ IRF Clinical Category (Uncollapsed)". The results of the CART analysis demonstrate that most options of binning CCs share similar model performance (i.e. average squared error [ASE]). As a result of the analysis and to ensure interpretability and clinical sense, the IRF CCs are finalized into 16 bins as shown in Figure 2b.

2.4 Evaluating IRF Methodology

Acumen evaluates performance of the new CCs model against the current RIC model to assess differences in cost prediction accuracy. Model performance is compared using adjusted R-squared values. A detailed breakdown of the model performance comparison can be found in Table 2.

- *Current IRF RIC Model: log(Standardized Cost) = CMG + (Comorbidity Tier * RIC)*
- *New IRC CC Model: log(Standardized Cost) = Clinical Category + Age Bin + Motor Score Bin + (Comorbidity Tier * Clinical Category)*

2.5 Major Deviations from SNF PDPM Methodology

Due to the different setting and payment structures, there are substantial methodological differences between the IRF PPS and SNF PDPM. Acumen made several key methodological adjustments to account for the unique characteristics of IRFs:

Use of Total Cost per Stay

Since IRFs are reimbursed on a per-stay basis, total costs per stay is used as the outcome variable, in contrast to the per-diem structure used in SNFs, where component-specific costs per day are used as the outcome variable.

Use of Principal Diagnosis on Prior Inpatient (IP) Claims

The principal diagnosis on prior IP claims is used as the source for identifying the primary reason for an IRF stay. This decision was based on an evaluation of various alternative options for assigning a primary diagnosis, including IGCs on IRF-PAI, etiological diagnoses on IRF-PAI and principal diagnosis on IRF Claims. Prior IP claims provide the highest combination of relevance and data accuracy for determining the reason for IRF care.

Exclusion of Surgical Spilt

The surgical split option when classifying IRF stays is excluded as it could potentially affect only 400 stays (~0.13% of the entire study population), and the affected stays are not cost outliers. Due to the minimal frequency and payment impact on IRF stays, expanding Section J in the IRF-PAI to accommodate the surgical split option is not needed.

Refinement of Clinical Categories

Initially, Acumen worked with the 11 clinical categories including a "Return to Provider" category derived from SNF PDPM. However, after reviewing the stay distributions of IRF by PDPM clinical categories (initial iterations of Figure 1), Acumen found that "Acute Neurologic" and "Non-Surgical Orthopedic/Musculoskeletal" were too broad. To respond to a high proportion of IRF stays assigned to broad categories, certain clinical categories are refined. This led to a further breakdown of "Acute Neurologic" into 4 CCs and "Non-Surgical Orthopedic/Musculoskeletal" into 4 CCs below:

- "Acute Neurologic - Brain"
- "Acute Neurologic - Cranial Nerve"
- "Acute Neurologic - Peripheral Nerve/Muscle"
- "Acute Neurologic - Spinal Cord"
- "Non-Surgical Orthopedic/Musculoskeletal - Infection"
- "Non-Surgical Orthopedic/Musculoskeletal - Injury"
- "Non-Surgical Orthopedic/Musculoskeletal - Peripheral Nerve/Muscle"
- "Non-Surgical Orthopedic/Musculoskeletal - Rheumatoid/Structural"

Additionally, we combined "Non-Surgical Orthopedic/Musculoskeletal – Cancer" with the broader "Cancer" clinical category. As a result of these changes, the total number of uncollapsed IRF CCs increased to 17, allowing for a more tailored and meaningful distribution of IRF stays.

Final Number of IRF Clinical Categories

As noted, the initial number of IRF CCs was 17, but after interpreting the results of the CART analysis and clinical consultation, Acumen proceeded with 16 IRF CCs. Specifically, while the default CART results produced 15-binned IRF CCs as the optimal model which grouped "Medical Management" and "Non-Orthopedic Surgery" into one CC, the 16-bin version was chosen as the final model as it allows

for greater interpretability and keeps distinct clinical categories separate without compromising on model performance.

3 RESULTS

This section summarizes key insights from the analysis, with corresponding figures and a table provided in the appendix.

Figure 1 compares stay distributions between SNF and IRF setting using PDPM clinical categories. The distribution of the stays between the two settings is generally comparable, with any notable variations attributable to expected population differences. Specifically, categories such as "Acute Infections," "Acute Neurologic," "Cancer," "Cardiovascular and Coagulation Disorders," and "Medical Management" show similar proportions of stays. In fact, about half of the clinical categories differ by only 1-2% in terms of stay percentages, indicating that both systems align closely in terms of how patients are categorized and treated. "Non-Surgical Orthopedic/Musculoskeletal" is the CC where the two systems differ the most and therefore, has been broken down into more specific CCs due to being too broad as mentioned in subsection 2.5.

Figure 2a and Figure 2b compare stays and associated costs distribution by RIC and IRF CC, respectively. The figures are sorted such that the average IRF costs per stay increase from left to right, with the least expensive categories on the left and the costliest on the right. The results shows that no single RIC or IRF CC dominates the distribution of stays, indicating a relatively balanced assignment across categories. The difference in average IRF cost per stay between the most expensive and least expensive RIC is $12,559. The cost differential across the IRF CCs is smaller, at $9,893. The results also suggest that the percentage of stays do not correlate with average IRF costs, meaning that cost differences exist independent of the volume of stays within each category.

Table 2 compares model performance across different IRF primary diagnosis models. The results demonstrate that the new IRF CC model achieved a similar adjusted R-squared value of 0.2684, compared to 0.2997 of the current RIC model. This indicates that the new model performs comparably well to the current RIC model while simplifying and modernizing the classification process.

APPENDIX – TABLES AND FIGURES

Table 1: Frequencies of Study Population Restrictions

Restrictions	Frequency		Cumulative Frequency	
	# of Stays	% of Stays	# of Stays	% of Stays
All	542,402	100.0%	542,402	100.0%
Part A Restrictions				
Beneficiary is enrolled in Part A at admission	412,850	76.1%	412,850	76.1%
Stay has positive utilization days	514,679	94.9%	402,464	74.2%
Stay has positive Medicare payment	406,341	74.9%	402,365	74.2%
Matching Restrictions				
Stay is matched to an IRF-PAI assessment	540,507	99.7%	402,134	74.1%
Provider of stay can be found in the FY 2023 IRF PPS Rate Setting File (RSF)	516,972	95.3%	381,094	70.3%
Provider of stay can be found in CASPER or POS	536,720	99.0%	377,190	69.5%
A cost report can be found for the provider	533,095	98.3%	377,190	69.5%
Validity Restrictions				
Stay has nonzero charges	542,402	100.0%	377,190	69.5%
Stay log standardized costs are not missing	483,430	89.1%	371,262	68.4%
Stay does not have an estimated outlier cost (0.5th <= costs per stay <= 99.5th)	478,594	88.2%	367,975	67.8%
Atypical Stays Restrictions				
Stay does not have a Disaster-Relief (DR) waiver	523,293	96.5%	354,815	65.4%
Stay is not a short-stay transfer paid on a per diem rate nor a stay that receives a blended transfer payment (Pricer return code on claims = 02, 03, 06, 07, 12, 13, 16, or 17)	468,179	86.3%	304,508	56.1%
Stay length is longer than 3 days	524,884	96.8%	301,833	55.6%
Stay's CMG does not equal 5001, 5101, 5102, 5103, nor 5104	542,373	100.0%	301,819	55.6%
Primary Diagnosis Restrictions				
Stay has at least one etiological diagnosis	540,507	99.7%	301,819	55.6%
Stay has a valid clinical category from the FY 2023 PDPM ICD-10 Mappings	542,048	99.9%	301,753	55.6%
Study Population	301,753	55.6%	301,753	55.6%

Figure 1: Stay Distribution by PDPM Clinical Categories

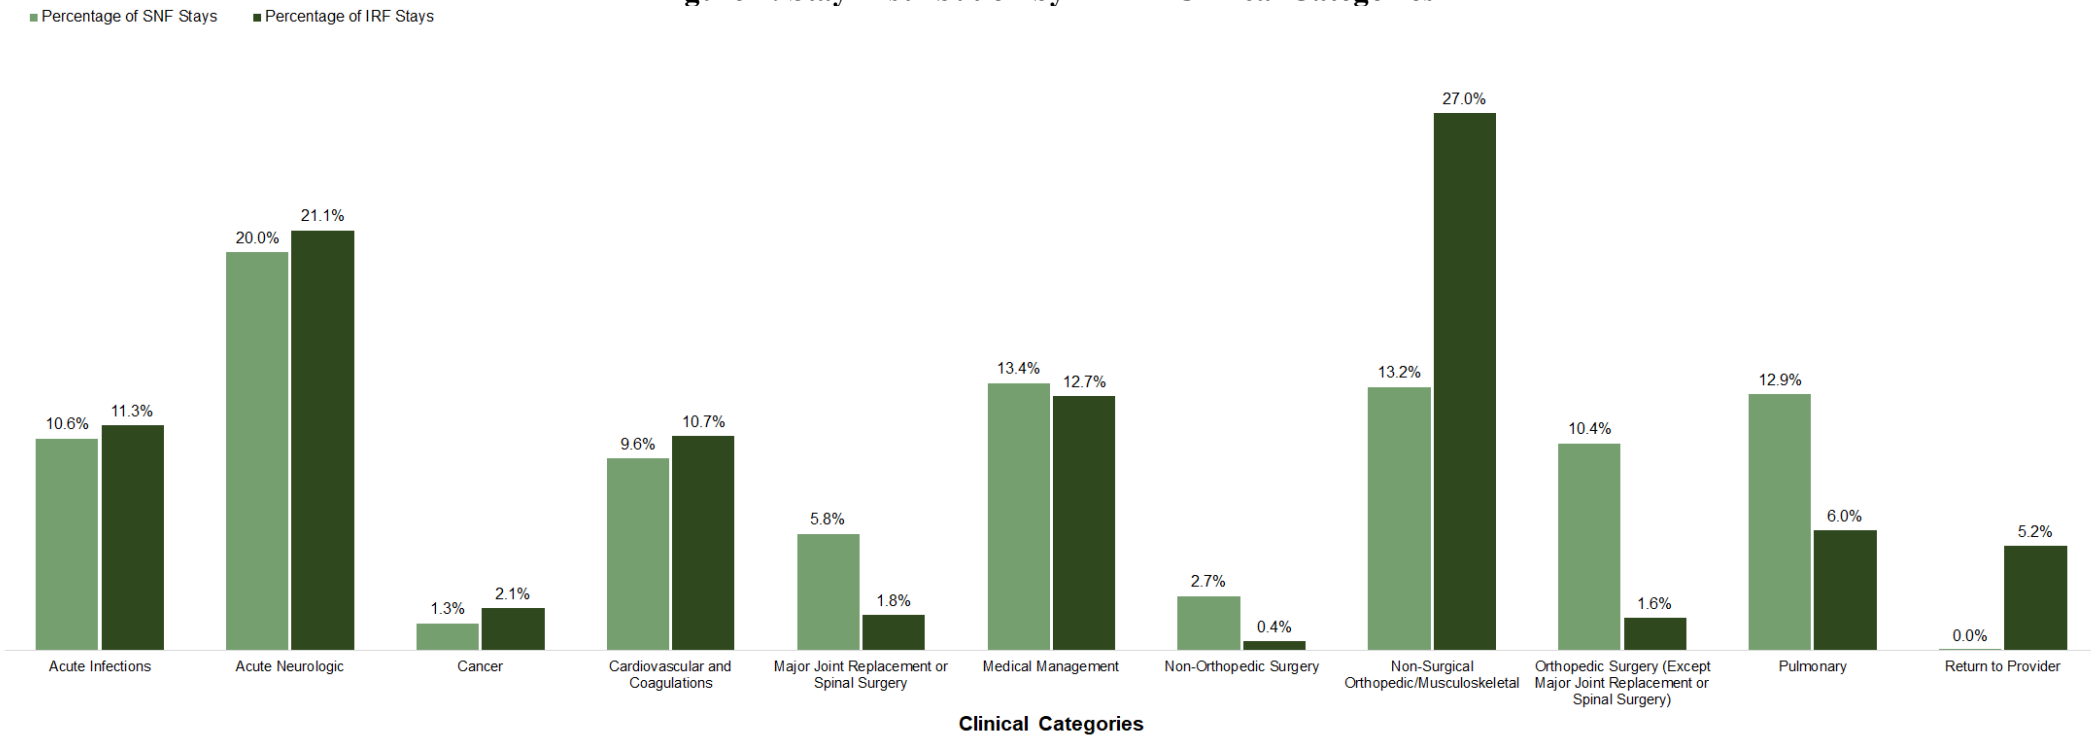

Figure 2a: Stay and Cost Distribution by RIC

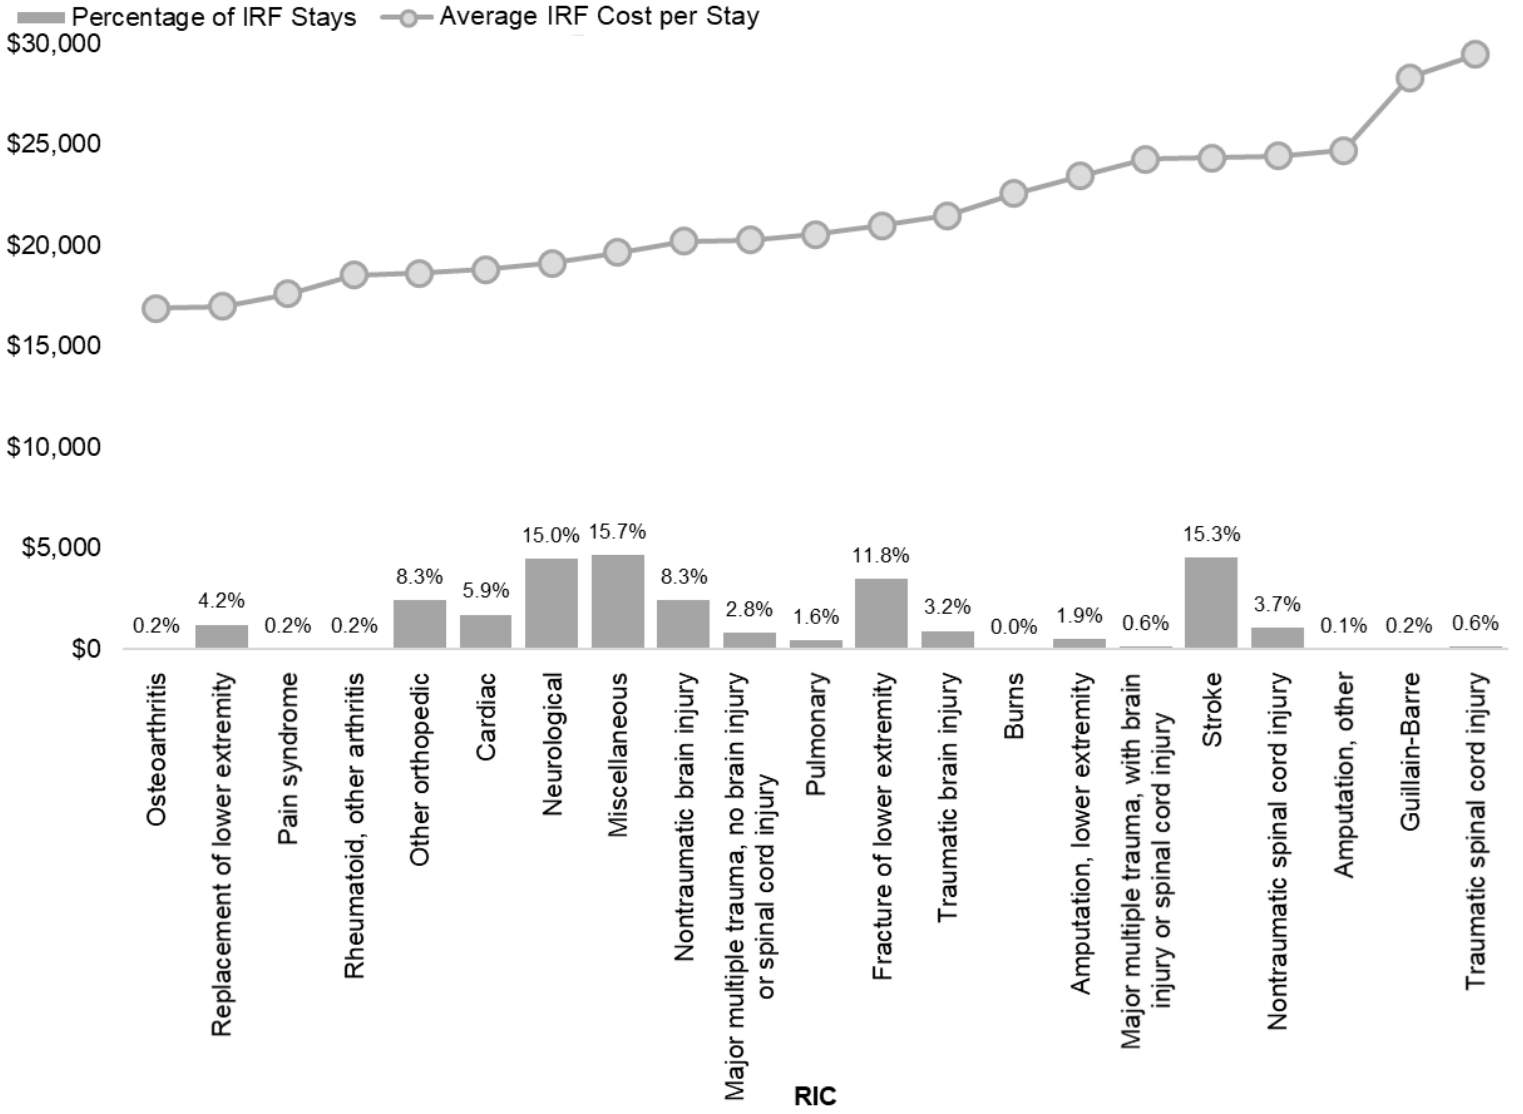

Figure 2b: Stay and Cost Distribution by IRF Clinical Categories

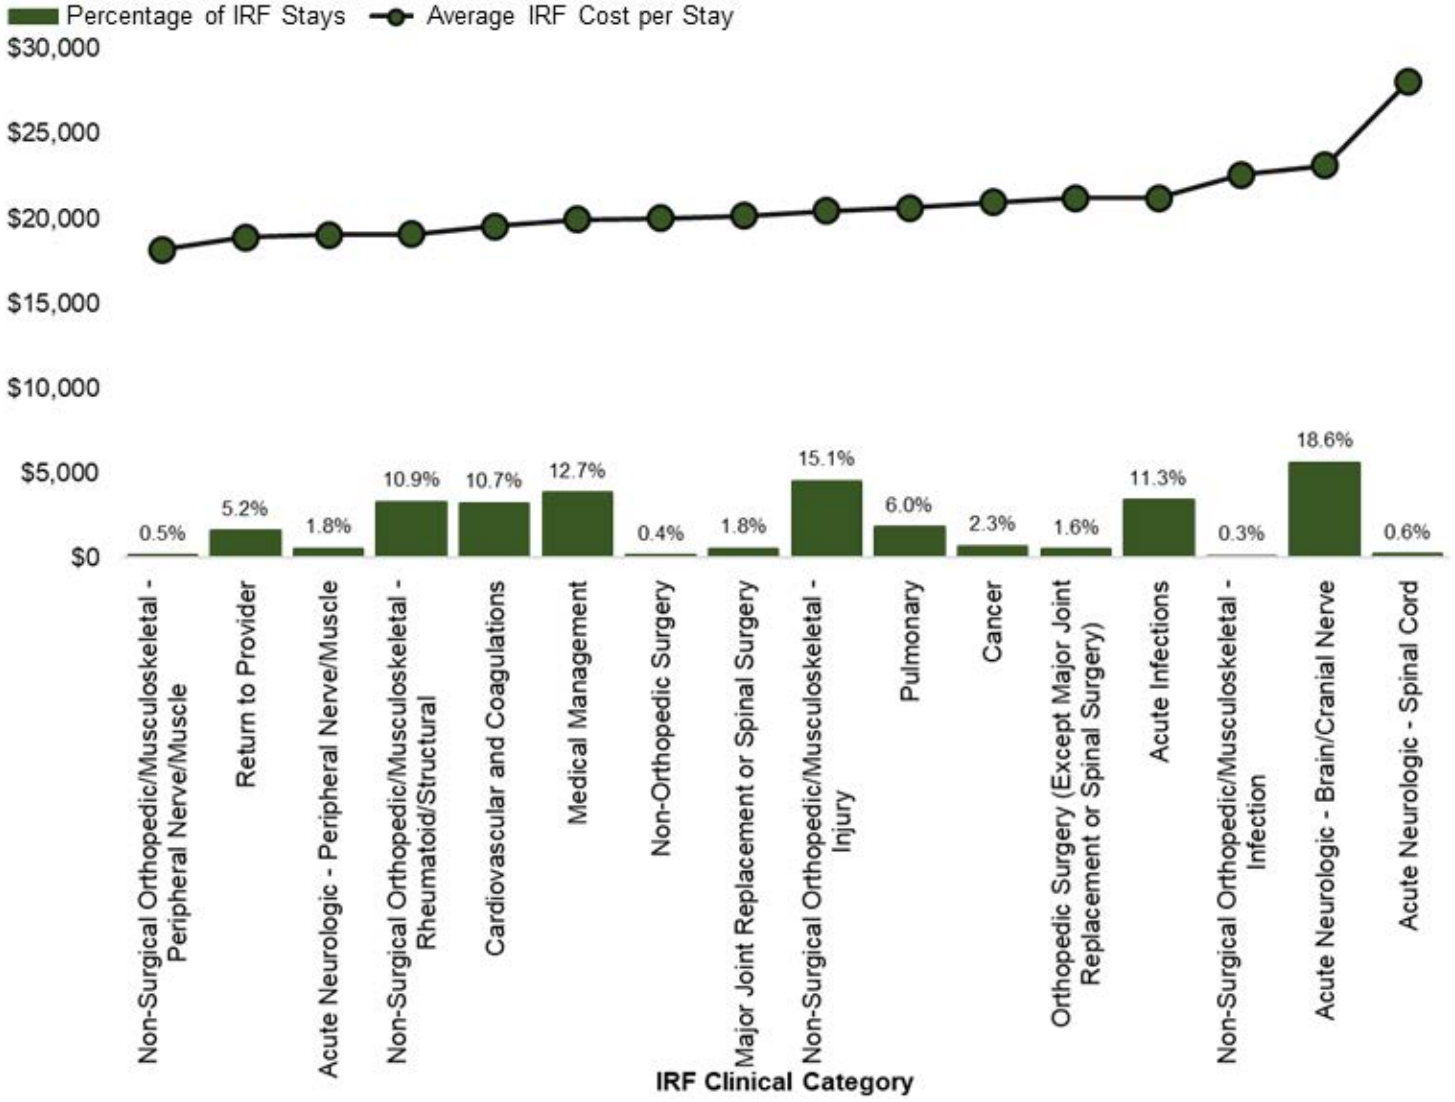

Table 2: System Comparison

System	System Specification	Adjusted R-Squared
Base Model	Log(Std Cost) = Intercept + Age Bin + Motor Score Bin + Comorbidity Tier	0.2487
Current RIC Model	Log(Std Cost) = Intercept + CMG + (RIC * Comorbidity Tier)	0.2997
New PDPM CC Model	Log(Std Cost) = Intercept + CC + Age Bin + Motor Score Bin + (CC * Comorbidity Tier)	0.2642
New IRF CC Model (Uncollapsed)	Log(Std Cost) = Intercept + CC + Age Bin + Motor Score Bin + (CC * Comorbidity Tier)	0.2686
New IRF CC Model	Log(Std Cost) = Intercept + CC + Age Bin + Motor Score Bin + (CC * Comorbidity Tier)	0.2684