

# The MSK Cost Crisis Doesn't Start in the Operating Room. It Starts With the Radiology Report.



[By Richard Herzog, M.D., FACR](#)

Chief Medical Officer, Covera Health | Former Director of Spinal Imaging, Hospital for Special Surgery | Attending Radiologist Emeritus, Hospital for Special Surgery | Former Professor of Radiology, Weill Cornell Medical College

March 19, 2026

I recently had the opportunity to reflect on two patients whose stories have stayed with me, not because they had unusual conditions, but because what happened to them happens every day, to workers in every industry, in every city in this country. Both patients were active, employed adults. Both had musculoskeletal pain that was interfering with their lives and their work. Both underwent MRI. Both were nearly sent down a care path that would have made them worse, not better. What saved them wasn't a better surgeon. It wasn't a more sophisticated treatment protocol. It was a second opinion of their radiology study. Most patients never get one.

Musculoskeletal conditions, spine, knee, shoulder, and hip, are among the most significant drivers of workforce disability, lost productivity, and healthcare spend for large employers. And yet, as employers invest in care navigation programs, surgical centers of excellence, and physical therapy platforms, one variable consistently escapes scrutiny: the quality of the radiology study that initiates and shapes the entire care journey. Before a surgeon forms a plan. Before a physical therapist designs a program. Before an authorization is issued. The MRI report has already set the trajectory. If that report is wrong, everything that follows is built on a false foundation.

This is the MSK quality problem that shows up everywhere in your outcomes data.

## Why MSK Imaging Is Uniquely Unforgiving

To understand why the radiology report carries this much weight, it helps to understand what makes MSK imaging so much harder to get right than most people assume. Consider a screening mammogram or a lung CT. The clinical question is essentially binary: is there a lesion, and does it cross a threshold of concern? The interpretation space is bounded. The stakes of any single read, while important, are constrained by the nature of the question.

Now consider a lumbar spine MRI, or a knee MRI, or a shoulder study ordered for a 45-year-old warehouse worker with six weeks of worsening pain. In a patient of that age, degenerative changes are universal. There will be multiple findings on virtually every study. Disc bulges, signal changes, partial tears, joint space narrowing, labral irregularities. The radiologist's task is not simply detection. It is clinical relevance ranking: identifying, among a constellation of findings, the one most likely responsible for this specific patient's symptoms, and communicating that with enough precision that the treating physician can act on it correctly.

That requires two things that vary enormously, and largely invisibly, across the imaging landscape.

The first is technical protocol quality. The sequences chosen during image acquisition, the field strength, the number of cuts, whether fat suppression is applied, these are not administrative details. They determine what is visible and what is not.

Protocol standards exist, but imaging centers under throughput pressure face persistent economic incentives to skip sequences or narrow coverage to complete the study faster. The result is studies that meet accreditation requirements on paper but fall short of what a complete MSK evaluation demands. In spinal imaging, as an example, a study that uses only ten sagittal cuts instead of sixteen to twenty, or that only angles axial cuts through the disc spaces rather than using a complete stack sequence, will simply not visualize parts of the spine. A disc herniation compressing a nerve root in unimaged territory cannot be reported. The patient's actual diagnosis remains invisible.

The second is interpretive expertise. A subspecialty MSK radiologist reads a knee study differently than a general radiologist, not marginally differently, but consequentially differently. Subtle tendon pathology, early cartilage defects, the distinction between a clinically significant labral tear and a normal variant, these are findings that require subspecialty expertise to identify reliably. And they are findings that, when missed or mischaracterized, could send patients and their physicians in entirely the wrong direction.

The evidence on the magnitude of that variability is striking.

## The Data That Should Change How You Think About Imaging

In 2017, colleagues and I published a study in *The Spine Journal* that I believe every employer health benefits leader should read. We sent a single patient, a 63-year-old woman with documented low back pain and right-sided radicular symptoms, to ten different MRI centers over a three-week period. All ten centers were accredited by the American College of Radiology. All ten were using standard imaging equipment. The patient and her reported symptoms were the same. The pathology was the same. The results were not.

Across the ten examinations, the average miss rate was 43.6%. Not a single finding was identified unanimously across all ten reports. Nearly a third of all reported findings appeared in only one report out of ten. Seventy-three percent of centers missed nerve root involvement. The average examination contained 12.5 interpretive errors.

I want to be precise about what this means, and what it doesn't. ACR accreditation establishes that a facility meets defined operational and equipment standards. It was never designed to measure ongoing interpretive quality or subspecialty expertise. Our ten centers were all accredited. Their reads varied enormously. Meeting a standard is not the same as delivering quality.

While this study focused on the lumbar spine, the phenomenon it documents is not unique to the spine. It reflects a systemic condition across MSK imaging, one that I have observed across 40 years of clinical practice.

## **What Happens When the Report Is Wrong**

Let me return to my two patients, because their stories illustrate something that data alone cannot fully convey.

The first patient came to me with medial knee pain. His referring physician suspected a meniscal tear, a reasonable hypothesis, and ordered an MRI specifically to rule it out. The report came back confirming a medial meniscal tear. Surgery was scheduled. What the report did not mention, and what I identified upon reviewing the images, was that the patient also had an MCL sprain, a second, independent potential source of the exact same medial knee pain. The original radiologist found the meniscal tear, confirmed the hypothesis, and stopped there. The MCL sprain, while visible in the images, was not originally identified.

The surgeon who scheduled that operation was not a bad doctor. He had complete confidence in the information he was given, and he acted on it logically. He simply never knew there was a differential to consider. When I reviewed the images and identified the sprain, another physician re-examined the patient with that differential in mind and concluded that the MCL sprain was the pain generator. The patient was treated

conservatively. Within a month, he had returned to full function and to work, without surgery.

The second patient presented with shoulder pain of several weeks' duration. His surgeon examined him, developed a clinical hypothesis of a superior labral tear, and ordered an MRI for confirmation. The report came back describing a labral tear. Surgery was scheduled. When I reviewed the images, there was no labral tear but a normal, anatomic variant of the labrum that can mimic a tear. This is not as uncommon as it sounds. Distinguishing between a tear and an anatomic variant requires subspecialty expertise that a general radiologist may not have. When my findings were communicated to the surgeon, he took a closer look at the images and agreed; the surgery was canceled. The patient underwent physical therapy and became fully asymptomatic within a few weeks.

These two cases represent the full spectrum of interpretive failure. The first is a false negative: a real finding, present in the images, omitted from the report, leaving the treating physician without the differential that would have changed everything. The second is a false positive: a finding documented in the report that simply was not there, accepted as fact because no one independently verified it.

In both cases, the mechanism is the same. The physician conducted a clinical evaluation, formed a hypothesis, and ordered imaging to confirm it. When the report appeared to validate that hypothesis, the clinical reasoning stopped. The report confirmed the working diagnosis. Sometimes the physician has no way of knowing. Sometimes the process makes it too easy not to look harder. In both cases, the system has failed the patient.

## **The Cost That Never Appears in Your Analysis**

For employers focused on managing healthcare spend, here is the calculation that is almost never made.

The unit cost difference between a high-quality MSK imaging center and a bottom-tier accredited center may be two hundred to four hundred dollars. Against a typical employer's total healthcare budget, this feels like a meaningful lever. It is not.

What is meaningful is what happens downstream when the report is wrong. An unnecessary meniscectomy. A missed nerve root compression routed to six weeks of physical therapy that cannot address the actual pathology. A labral repair performed on a labrum that didn't need repair. Repeat imaging when the first treatment fails. A second specialist consultation. A second surgery. Extended disability. Lost productivity for weeks or months.

And then there is a dimension that is rarely discussed in benefits conversations but is documented clearly in the clinical literature. The landmark Spine Patient Outcomes Research Trial, SPORT, one of the most comprehensive spine outcomes studies ever conducted, identified six months as a critical threshold.

Patients with symptom duration of a disc herniation under six months had better long-term outcomes than those whose correct diagnosis was delayed beyond that window. This finding held regardless of whether patients ultimately received surgery or conservative treatment. The physiology of chronic pain is unforgiving: neural sensitization, muscular compensation patterns, functional adaptation, these changes accumulate over time and become progressively harder to reverse.

What this means practically is this: a diagnostic error that delays correct treatment past the six-month mark is not merely an inconvenience that gets corrected later. In a meaningful subset of patients, it is a permanent outcome event. The employee does not just have a difficult few months. They may never return to their prior functional level.

The downstream cost of that outcome, long-term disability claims, lost productivity, replacement hiring, retraining, dwarfs any cost savings captured at the imaging stage by an almost incomprehensible margin.

The employer claims data tells the same story. A large, self-insured employer, already running best-in-class orthopedic and rehabilitation programs, changed one variable: radiology quality. Routing MSK patients to imaging centers with rigorous protocols and subspecialty expertise produced a 26% reduction in downstream in-episode costs compared to a control group receiving imaging through standard network channels. One variable. One change. Twenty-six percent.

## **What Good Actually Looks Like**

I want to be clear about what I am not arguing. I am not arguing that the most expensive imaging is the best imaging. Hospital-based imaging is frequently overpriced by a substantial margin, driven by facility fees and institutional overhead that have nothing to do with diagnostic quality. Steering employees away from hospital imaging may be entirely appropriate and when done can generate meaningful savings.

What I am arguing is that price and quality are independent variables, and in MSK imaging, both must be measured. The highest-value imaging for an MSK patient is not the cheapest, and it is not the most expensive. It is the study performed with rigorous, comprehensive protocols by a radiologist with demonstrated subspecialty expertise, delivered at a fair market price. That combination exists. It simply requires the infrastructure to identify it.

For employers who are serious about MSK outcomes and costs, the operating room is not where the work begins. The radiology study is. And right now, for employees, that study is the weakest link in the entire chain, invisible, unmeasured, and quietly determining everything that follows.