The Price of Better Dosing: FDA Regulation and the Architecture of Oncology Trials

Zachary Santamaria*

November 3, 2026

Abstract

Regulators face a constrained optimization problem over the quantity, timing, and form of evidence required from firms. We examine how pharmaceutical sponsors respond to evolving evidentiary requirements using Project Optimus, a late 2021 U.S. Food and Drug Administration (FDA) initiative that increased emphasis on dose optimization during oncological drug clinical development. We assemble a novel ClinicalTrials.gov data panel covering 2000 to 2026, by matching and combining 75,233 early-phase drug and biologic trials from 2015 to 2026, historical registry versions for 20,247 trials, and 58,481 additional studies from 2000 to 2014 used to reconstruct long-run sponsor-asset histories. Our synthesized data panel distinguishes project entry, initial experimental design, subsequent registry amendment, enrollment, and recruitment. We find that relative to contemporaneous non-oncology development by the same sponsors, oncology trials become 5 percentage points more likely to incorporate randomized dose comparison into their initial registered design. Explicit dose-optimization language rises by 2.5 percentage points. These design changes are already present at trial launch rather than emerging through later registry amendments. We find weaker evidence of proportional increases in planned enrollment or slower recruitment. Conversely, adjustment is evident on the extensive margin of project selection. Following the August 2024 final guidance, the probability that an industry sponsor initiates a new oncology asset is approximately 3.5 percentage points lower relative to its contemporaneous non-oncology development. This estimate is equivalent to approximately 19 percent of the pre-final oncology entry rate, even though oncology entry continues to rise in absolute terms. Our findings show that evidentiary regulation can reshape both which projects enter clinical development and the architecture of the experiments used to study them.

Keywords: medical innovation; pharmaceutical regulation; clinical trials; Project Optimus; oncology; dose optimization; evidence production; project selection.

*Corresponding author: Zachary Santamaria, SUNY Purchase, 735 Anderson Hill Road, Purchase, NY 10577. Email: zachary.santamaria@purchase.edu.

1 Introduction

Regulators shape the production path of evidence that pharmaceutical companies must follow insofar getting their critical biologic assets to market. These companies face a constrained optimization problem over how much information firms must produce, when in the development process it must be generated, and what form that evidence should take. These choices can affect more than the cost of an individual clinical trial. They can alter which projects firms choose to pursue and how firms design the experiments used to study them. We examine these responses in the context of Project Optimus, an FDA initiative that increased the emphasis placed on dose optimization during oncology drug development. The initiative provides a useful setting because it changes the informational expectations surrounding a specific part of the development process. Rather than asking firms simply to conduct more research, the policy places greater value on resolving uncertainty about dosage before approval U.S. Food and Drug Administration (n.d., 2024). This creates several possible margins of adjustment. Sponsors may change which oncology assets enter clinical development, reorganize the comparisons embedded in trials that do proceed, increase the number of patients enrolled, lengthen recruitment, or combine these responses.

Our analysis firmly delineates these margins using a new ClinicalTrials.gov dataset spanning 2000–2026. Our primary trial data panel contains 75,233 early-phase drug and biologic studies beginning between 2015 and 2026. We then reconstitute historical registry versions for 20,247 trials, allowing us to observe experimental design near study launch rather than relying exclusively on the current registry record. Further, we incorporate 58,481 early-phase studies from 2000–2014. We then use these early-phase studies to reconstruct long-run sponsor-asset histories and distinguish genuine clinical-development projects from compounds already present in a sponsor's pipeline. The historical data is critical for empirical measurement. ClinicalTrials.gov trial records evolve as studies progress. Enrollment targets can change, trial arms can be revised, and new language can appear long after recruitment begins. Current records combine original design choices with subsequent amendments and

different stages of trial maturity. Our diligence in recovering earlier versions allows us to identify what sponsors planned to compare when an experiment began and to follow enrollment and recruitment at common horizons. Our empirical design benchmarks oncology development against contemporaneous non-oncology activity by the same sponsors. We compare oncology and non-oncology studies (trial-level) initiated within common sponsor-time environments. We reorganize the data into a sponsor-by-domain-by-quarter panel that explicitly includes quarters with zero new projects in order to identify project entry. We identify new sponsor-drug assets using intervention names, registry aliases, and experimental-arm assignments, and use the pre-2015 history to reduce left-censoring in the measurement of first clinical entry.

The clearest intensive-margin response appears in trial architecture. Following Project Optimus, oncology trials become roughly five percentage points more likely to incorporate randomized dose comparison into their initial registered design relative to contemporaneous non-oncology development by the same sponsors. Explicit dose-optimization language rises by roughly two and one-half percentage points. Historical versions show that the randomized-dose difference is already present near trial launch, while the probability that comparable language is added later changes little. The response appears in the formation of the experiment rather than through retrospective modification of the registry. The evidence for proportional expansion of trial scale is considerably weaker. Historical enrollment targets reconstructed at launch and at common post-launch horizons show little systematic increase, and recruitment histories provide no clear evidence of a persistent slowdown. Larger differences visible in current enrollment records diminish as trials are required to have more time to mature. These results suggest that sponsors can increase the informational content of an experiment without mechanically increasing patient scale.

We find evidence of adjustment at the project-selection margin. Around the final guidance, the probability that an industry sponsor initiates a new oncology asset is approximately 3.5 percentage points lower relative to its contemporaneous non-oncology development. The

magnitude is equivalent to roughly 19 percent of the pre-final oncology entry rate. Moreover, oncology entry continues to increase in absolute terms; the result reflects slower growth in oncology entry relative to the same firms' non-oncology portfolios rather than an absolute contraction in oncology development. Pre-policy placebo exercises and local event-study diagnostics provide substantially cleaner support for this asset-entry result than broader measures of trial-start frequency.

We contribute to the economics of medical innovation by showing that firms respond to changing evidentiary expectations through the regulatory architecture of scientific production (Project Optimus). The regulatory response is visible in what sponsors design trials to compare, not merely in how many patients they enroll. Further, we provide ample evidence that the adjustment can extend to project selection. Evidentiary regulation may influence both which projects cross the threshold into clinical development and how firms organize research conditional on proceeding. This distinction matters for evaluating regulatory burden. The information requirements of a regulatory policy will not necessarily increase trial size and duration or research expenditure. Firms can respond by shifting between margins, changing the experimental design, shifting resources into different therapeutic areas or being more selective about the marginal projects they take into clinical trials. The economic impact will depend on how much research and development is undertaken as well as what type of research is initiated.

The remainder of the paper proceeds as follows. In section 2 we describe Project Optimus and the related literature. Section 3 presents the ClinicalTrials.gov data, historical registry reconstruction, and sponsor-asset identification. Section 4 develops the empirical framework for project entry, trial architecture, enrollment, and recruitment. Section 5 presents the results and robustness exercises. Section 6 concludes.

2 Literature Review

2.1 Dose Optimization and Project Optimus

Misgivings surrounding oncology dose selection has preceded the Project Optimus era. Minasian et al. (2014) advocated for increased evidence beyond toxicity propounding other factors such as exposure, effectiveness, tolerability, and patient heterogeneity. They argue that better dose characterization can improve effectiveness and reduce toxicity. Shah et al. (2021) question the traditional "more is better" paradigm used in oncology dose selection. This approach reflects the historical cytotoxic-chemotherapy practice of advancing doses near the maximum tolerated dose, a strategy that is less appropriate for many targeted therapies Samineni et al. (2024). They note that oncology doses and schedules are sometimes inadequately characterized before registration trials, particularly for targeted therapies. Sotorasib provides a concrete example. At its 2021 accelerated approval, the FDA required a postmarketing comparison of the approved 960 mg dose with a 240 mg dose because the agency did not consider the approved dosage fully optimized. Brock et al. (2021) offer empirical by analyzing dose-level outcomes from 115 oncology dose-finding studies. They report that toxicity increases more consistently with dose than efficacy does. U.S. Food and Drug Administration (FDA) formalized the concern through their Project Optimus. Project Optimus encourages early characterization of the dose-response properties U.S. Food and Drug Administration (n.d.). Further, Project Optimus promotes randomized comparisons among clinically relevant doses, rather than assuming that the highest tolerated dose should advance. Randomized trials assist with improving comparability between patients on different dosing schedules, providing improved estimation of dose to exposure response relationships for activity, safety and tolerability U.S. Food and Drug Administration (n.d.). The August 2024 final guidance provides further details regarding what evidence should be provided to demonstrate optimal dosage prior to submission for a new oncology indication. Including evidence of safety, tolerability, efficacy, pharmacokinetics, pharmacodynamics and exposure-

response relationships. The guidance additionally promotes using direct comparisons of multi-dose groups and model-informed strategies such as adding further dose cohorts or backfill cohorts if they improve dose characterization Rahman et al. (2024). Other propositions include the consideration of all relevant evidence in relation to dosage, not just maximum tolerated dose. It is encouraged by the regulators to review non-clinical and clinical data throughout drug development as new safety/efficacy/exposure/tolerability data becomes available. Quantitative methods can be used to help identify the optimal doses to progress to late stage and registrational trials.

The methodological literature has changed along with the regulatory environment. Thall et al. (2024) point out that the dose selected in a conventional Phase I trial does not always end up being the dose clinicians prefer to use. Capecitabine is a working example. In an ASCO survey they discuss, 41 percent of oncologists treating breast or gastrointestinal cancers said they regularly started patients below the drug's original labeled dose. Physicians having been making these adjustments in routine practice suggests that a dose chosen largely on the basis of toxicity may not provide the best balance between efficacy and tolerability. Yuan et al. (2024) examine the DREAMM-2 trial, where 97 patients received 2.5 mg/kg of belantamab mafodotin and 99 received 3.4 mg/kg. The two dose regiments produced similar response rates, 31 percent and 34 percent. The lower dose led to fewer serious adverse events, treatment interruptions, and dose reductions. Those differences buttressed the selection of the 2.5 mg/kg dosing. Reaching that conclusion required collecting comparative evidence on efficacy, tolerability, safety, and exposure across multiple doses earlier in development. Murphy et al. (2023) expounds that a review of molecularly targeted therapies in Phase III trials found that 48 percent of patients needed their doses modified even though those doses had already been recommended in Phase I. Their discussion points to a basic problem: identifying a tolerable dose is not the same thing as identifying the dose that will work best in clinical practice. It is crucial that an institutional and regulatory framework can allow trial architecture to take this information into account.

Evidence from industry indicates that these design changes are happening. A survey of 18 pharmaceutical companies showed that 46% of companies preferred randomized dose-comparison expansion cohorts, whilst 36% preferred adaptive dose selection trials Samineni et al. (2024). 75% of companies preferred to establish proof of concept prior to dose optimization. Further, a review of FDA approvals showed that there is considerable post-marketing demand for further dose evidence. It is important to note that due to the lack of patients in rare cancers it can be difficult to conduct large randomized dose-ranging studies, so strategies such as strategic backfill and more efficient dose exploration within trials should be used. Overall there is evidence to indicate that many oncology products have needed more dose optimization after approval. Glasmacher et al. (2025) came to similar conclusions. Evidence shows that when choosing the optimum biological dose as the recommended Phase II dose 83% of programs had the final approved dose matching the optimum biological dose. When choosing the maximum tolerated dose as the recommended Phase II dose only 58% of programs had the final approved dose matching the optimum biological dose. This workshop supports the importance of integrating pharmacokinetics, pharmacodynamics, biomarkers and patient-reported tolerability earlier on.

Recent industry evidence points toward continued implementation. Ananthakrishnan et al. (2026) survey oncology statisticians across biopharmaceutical companies. Evidence shows the approach is moving away from using max tolerated dose regiments. Respondents reported moving toward more integrated dose-optimization strategies that jointly consider efficacy, safety, pharmacokinetics, and patient experience, consistent with the aims of Project Optimus. Evidence shows there has also been a change in statistical practice and structure within early oncology development through Project Optimus. Bhamidipati et al. (2025) review 367 industry-sponsored Phase I protocols activated at eight Sarah Cannon Research Institute drug-development units between 2021 and 2024. Thirty percent contained an explicit dose-optimization plan. Moreover, they report that Bayesian dose-escalation designs became substantially more common over the period. Among protocols containing a dose-

optimization plan, 75 percent specified randomization across more than one dose level, with a median planned incremental enrollment of 40 patients. Their evidence establishes that protocol practice evolved during the Project Optimus era. Our analysis addresses a different empirical question. We ask whether oncology development changed differentially relative to contemporaneous non-oncology development conducted by the same sponsors. We also reconstruct the first registered version of each trial, which allows us to distinguish a change in the experimental blueprint at launch from later amendments to the registry.

2.2 Regulation and the Production of Evidence

Further, our analysis relates to an economics literature showing that regulation can reshape the incentives, organization and timing of research. Stern Stern (2017) finds that pioneer devices, defined as the first entrants within an FDA product code, spend substantially longer in regulatory review than follow-on products. They find that pioneer products spend 34 percent longer in regulatory review relative to follow-on entrants, a difference of approximately 7.2 months and estimate that this delay can exceed 7 percent of the total cost of bringing a high-risk device (subject to the FDA's Premarket Approval - PMA) process to market. When the regulator publishes formal guidance clarifying application and evaluation requirements, subsequent approval times fall by approximately 6.1 months, or 41 percent. These results provide evidence that uncertainty about evidentiary and procedural expectations can itself constitute a substantial development cost. The FDA's Breakthrough Therapy Designation (BTD) provides a closer parallel. Hermosilla (2024) finds that the program changed industry-wide development incentives, not merely the treatment of designated products. They find evidence that a one-standard-deviation increase in exposure to BTD incentives is associated with approximately 30 percent more new indications entering clinical development. The increase reaches 13 percent during 2016–2018 and 19 percent during 2019–2021. New product launches subsequently rise by about 7.4 percent during 2019–2021. Concomitantly, the evidence does not indicate greater risk-taking in project selection. Chandra et al. (2026)

examine a different margin of the same program. They report evidence that BTD shortens clinical development by approximately 23 percent without worsening the ex post safety profile of designated drugs. The gains are largest among less experienced firms. BTD is also associated with simpler clinical-trial designs, suggesting that regulatory flexibility can reduce both development time and experimental complexity. Project Optimus operates in a different informational direction. BTD provides regulatory flexibility that can economize on development time and experimental complexity. Rather, Project Optimus raises expectations for resolving a specific uncertainty before approval: which dosage should advance through development. This contrast makes oncology especially useful for studying how firms reorganize research when regulation calls for richer premarket evidence rather than faster evidence.

2.3 Clinical-Trial Complexity

The preceding research is clear in reporting that clinical trial enrollment does not fully define experimentation costs. Getz and Campo (2018) analyze 9,737 clinical protocols and 130,601 investigative-site contracts. Between 2001–2005 and 2011–2015, nearly every dimension of Phase I–III protocol complexity increased significantly. Trials required more procedures, site visits, and site work. Costs per planned participant visit rose significantly. Phase III trials experienced the largest growth in overall protocol complexity. Protocol changes create an additional source of development burden. Getz et al. (2016) examine 836 Phase I–IV protocols and find that 57 percent underwent at least one substantial amendment. Nearly 45 percent of those amendments were judged avoidable. Phase II and Phase III protocols averaged 2.2 and 2.3 global amendments. Amended protocols had longer recruitment periods and longer overall study durations. Getz et al. (2016) finds that direct median implementation costs reached $141,000 per Phase II substantial amendment and $535,000 per Phase III substantial amendment. These findings motivate a distinction between experimental scale and experimental architecture. Scale captures how much research is conducted, such as enrollment or recruitment duration. Architecture captures how that research is organized

such as treatment arm utilization, comparisons, and procedures among others. Enrollment measures the number of participants exposed to an experiment. Architecture determines the procedures, comparisons, visits, treatment components, and revisions required to generate evidence. Let clinical information be produced according to:

$$I = F(N, A, D, M, S, T), \tag{1}$$

where N denotes enrolled patients, A treatment arms, D intervention components, M measurement intensity, S site infrastructure, and T development time. A regulator that raises the required information content of development need not induce an equal increase in each input. A sponsor could enlarge enrollment, but it can also repartition a similar patient population across more informative comparisons, alter treatment structure, collect more or richer set of biomarkers, or shift evidence generation across stages of development. The cost of better evidence may appear in the organization of experimentation even when patient scale moves little. Our results fit this distinction. The clearest response to Project Optimus is a change in randomized dose comparison in the protocol initially registered at trial launch. Evidence of larger initial enrollment targets and slower recruitment is considerably weaker.

3 Data

3.1 ClinicalTrials.gov Trial Data Panel

Our data panel comes from ClinicalTrials.gov. We retrieve interventional studies involving drugs or biological products and beginning between January 1, 2015 and September 15, 2026. We retain all studies classified as Early Phase I, Phase I, Phase I/II, or Phase II. Our cross-tabulated classification includes: 75,233 early-phase drug and biologic trials with 32,074 studies described as oncology trials and 43,159 as non-oncology trials. Our retained non-oncology studies are used as a contemporaneous control for determining changes in

clinical development that are independent of oncology. For each clinical trial study, the registry provides the following details: sponsor, phase, study dates, allocation, enrollment, study arms, interventions, site locations, recruitment status, outcomes, and free-text protocol descriptions. Each empirical question requires a different view of the registry. Analyses of current trial characteristics use the full trial panel. Analyses of initial experimental design instead require historical registry versions near study launch, because enrollment, trial arms, interventions, and protocol language can change after registration. The extensive-margin analysis is more restrictive. Identifying project entry requires industry sponsors active in both oncology and non-oncology, which permits comparisons within the same development portfolio. It requires a long prehistory of sponsor-asset activity, since a compound first observed in the main panel may have entered clinical development years earlier. We therefore use pre-2015 records to distinguish genuinely new assets from previously active development programs.

Table 1: Construction of the Primary Dataset

Sample	Observations
Early-phase drug and biologic trials, 2015–2026	75,233
Oncology trials	32,074
Non-oncology trials	43,159
Guidance-covered oncology proxy	28,290
Excluded-modality oncology proxy	3,784
Historical same-sponsor-quarter core	20,247
Reliable launch-architecture records	17,784
Industry trials among dual-domain sponsors	19,952
Industry sponsors active in both domains	983

3.2 Phase Composition

The composition of the analytical sample across early development phases is stable over the regulatory period. Table 2 reports phase shares separately for oncology and non-oncology trials. Among oncology studies, Phase I trials account for 29.9 percent of the sample before 2021 and 25.9 percent after final guidance. Phase I/II trials remain close to 17 percent throughout the period. Non-oncology trials exhibit similarly modest changes in phase composition. The

stability of these shares suggests that the post-Project Optimus changes documented below are unlikely to be explained solely by a broad shift in the mix of development phases.

Table 2: Trial Phase Composition by Regulatory Period and Domain

Period	Domain	Early Phase I	Phase I	Phase I/II	Phase II
Pre-2021	Oncology	3.8%	29.9%	16.0%	50.2%
Pre-2021	Non-oncology	5.5%	38.6%	9.1%	46.7%
Optimus LTD	Oncology	4.3%	30.2%	16.9%	48.6%
Optimus LTD	Non-oncology	6.3%	38.7%	9.4%	45.6%
Draft to final	Oncology	5.0%	28.2%	16.9%	49.8%
Draft to final	Non-oncology	7.6%	39.5%	9.6%	43.3%
Post-final	Oncology	5.5%	25.9%	16.7%	51.9%
Post-final	Non-oncology	7.6%	37.5%	10.5%	44.4%

3.3 Historical Registry Versions

ClinicalTrials.gov records are revised throughout a study's life cycle. Sponsors can alter enrollment targets, modify arm descriptions, update intervention information, change recruitment status, and revise free-text fields. The current record need not reproduce the experimental plan registered when a trial began. We reconstruct the version history of the 20,247 trials in the historical study. Our launch sample requires the first usable record to have been submitted no more than 30 days after Study Start Date. This leaves 17,784 trials, including 8,857 oncology and 8,927 non-oncology studies. Coverage of the principal design variables is high. Among the reliable launch records, 17,724 report at least one study arm, 17,773 report at least one intervention, and planned enrollment is recovered for 17,744 studies. The historical sample remains substantial throughout the regulatory period. There are 7,456 reliable launch records before 2021, 6,095 between the launch of Project Optimus and the final guidance, and 4,233 after final guidance.

Table 3: Historical Registry Reconstruction

	Oncology	Non-oncology	Total
Panel A: Historical coverage			
Historical same-sponsor-quarter core			20,247
Reliable launch record within 30 days	8,857	8,927	17,784
Launch record with at least one arm			17,724
Launch record with at least one intervention			17,773
Launch record with planned enrollment			17,744
Panel B: Reliable launch sample by period			
Pre-2021	3,548	3,908	7,456
2021 through August 7, 2024	3,044	3,051	6,095
Post-final guidance	2,265	1,968	4,233
Total	8,857	8,927	17,784

3.4 Long-Run Sponsor-Asset Histories

Our extensive-margin analysis requires distinguishing genuinely new projects from drugs already present in a sponsor's development portfolio. Restricting attention to the 2015–2026 analytical panel would create substantial left-censoring: an asset appearing in 2017 could have entered clinical development many years earlier. We retrieve an additional 58,481 early-phase interventional studies beginning between 2000 and 2014. These observations are not used as the main treatment-period trial panel. Instead, they provide a historical lookback for reconstructing prior sponsor-asset activity. Combining the pre-2015 history with the current panel gives sponsor-asset histories spanning 2000–2026. For the industry dual-domain sample, the combined records generate 24,829 strict sponsor-asset components. Under this strict definition, an asset is counted only when its normalized drug alias appears in an experimental arm and not solely as a comparator or background treatment. We classify 13,888 sponsor-assets as eligible first observed entries from 2017 onward. Only three first-entry records have unresolved same-date oncology and non-oncology domain ties and are excluded.

Table 4: Construction of Sponsor-Asset Histories

Measure	Count
Prehistory studies retrieved, 2000–2014	58,481
Industry dual-domain sponsors	983
Industry trials among dual-domain sponsors, 2015–2026	19,952
Drug/biologic intervention alias rows, 2015–2026	55,456
Drug/biologic intervention alias rows, 2000–2014	39,230
Strict experimental-only alias rows, 2015–2026	47,599
Strict experimental-only alias rows, 2000–2014	29,118
Strict sponsor-asset components, 2000–2026	24,829
Eligible first observed sponsor-assets, 2017 onward	13,888
Ambiguous same-date cross-domain entries excluded	3
Post-final new-asset entries available for audit	3,194

3.5 Identifying Focal Experimental Assets

A drug making an appearance in a trial record is not evidence insofar being the asset under development by the sponsor. Oncology trials frequently contain active comparators, background therapies, and combination partners. Counting every listed drug as a sponsor asset would overstate project entry. We use ClinicalTrials.gov arm-intervention mappings to identify focal experimental treatments. A drug or biologic enters the strict asset definition when it is assigned to an experimental arm and is not used solely as a comparator in that trial. Intervention aliases reported by ClinicalTrials.gov are then linked within sponsor so that development codes, generic names, and other explicitly connected names can contribute to a common sponsor-asset identity. The role restriction materially narrows the intervention data. In the 2015–2026 industry sample, approximately 91 percent of drug or biologic alias rows appear in at least one experimental arm, while about 86 percent satisfy the stricter experimental-exclusive definition. Comparator use is more common in the pre-2015 records.

Table 5: Intervention Roles Used in Sponsor-Asset Construction

Period	Alias rows	Trials	Experimental exclusive	Comparator use
2015–2026	55,456	19,930	85.8%	13.1%
2000–2014	39,230	13,752	74.2%	19.5%

The historical lookback reduces the likelihood that an established asset is mistaken for a new clinical-development project simply because its first appearance in the main analytical

panel occurs after 2015. Our extensive-margin outcome is the first observed strict experimental sponsor-asset entry after this long prehistory has been incorporated.

3.6 Sponsor-Quarter Entry Panel

Trial-level data only contains studies that were initiated. Studying the extensive margin requires observations for periods in which a sponsor could have initiated an oncology project but did not. We reorganize the industry sample into sponsor-by-domain-by-quarter records. Each sponsor contributes separate oncology and non-oncology observations. Within the sponsor's active development window, quarters with no new asset or trial are retained explicitly as zeros. This construction yields 38,552 sponsor-year-quarter-domain observations in the primary extensive-margin panel. A narrower composition sample retains quarters in which the sponsor has observed development activity and contains 17,618 sponsor-quarter-domain observations. For each sponsor-domain-quarter we measure the number of new strict sponsor-assets, whether any new asset enters, the number of early-phase trial starts, whether any trial starts, and analogous trial-start measures restricted to studies containing a strict focal experimental asset.

Table 6: Structure of the Extensive-Margin Panel

Panel component	Count
Industry dual-domain sponsors	983
Current-period industry trials	19,952
Primary sponsor-year-quarter-domain observations	38,552
Composition sponsor-quarter-domain observations	17,618
Eligible new sponsor-assets, 2017 onward	13,888
Post-final new sponsor-assets	3,194

3.7 Project Entry and Trial Activity

Table 7 describes project entry before any regression adjustment. This is useful for distinguishing absolute changes in oncology activity from changes relative to non-oncology

development. Before 2021, the probability that a sponsor-domain-quarter contains a new asset is approximately 12.8 percent in oncology and 13.6 percent outside oncology. Both oncology and non-oncology entry rates rise during the Project Optimus period. After final guidance, 20.8 percent of oncology sponsor-domain-quarter observations contain at least one new asset. The corresponding share is 22.3 percent for non-oncology development. Thus, oncology entry does not decline in absolute terms. Oncology grows less rapidly than new-asset entry elsewhere in the same sponsors' portfolios. The difference is larger for trial initiation. Following final guidance, we observe that at least one early phase trial has started for 25.9% of oncology sponsor-domain-quarter observations compared to 34.5% of non-oncology observations. This shows the average values are not extreme, and therefore indicates that we are investigating whether there is a difference in oncology development relative to the non-oncology development by the same sponsors. It does not indicate that oncology research is being discontinued or reduced.

Table 7: Project Entry and Trial Activity by Period and Domain

Period	Domain	Mean new assets	Any new asset	Mean trial starts	Any trial start
Pre-2021	Non-oncology	0.321	13.6%	0.651	29.9%
Pre-2021	Oncology	0.257	12.8%	0.394	23.8%
Optimus LTD	Non-oncology	0.401	17.5%	0.565	28.2%
Optimus LTD	Oncology	0.365	18.6%	0.399	25.0%
Draft to final	Non-oncology	0.395	18.2%	0.571	28.1%
Draft to final	Oncology	0.386	18.8%	0.374	24.1%
Post-final	Non-oncology	0.508	22.3%	0.731	34.5%
Post-final	Oncology	0.483	20.8%	0.454	25.9%

The entry series show that oncology development continues to expand after final guidance. The empirical question is not whether oncology entry disappears, but whether its growth differs from contemporaneous development in the same sponsors' non-oncology portfolios.

3.8 Trial-Architecture Measures

For trials that enter development, we characterize the experiment. Our architectural outcome identifies randomized dose comparison in the initial registered protocol. The measure combines structured evidence of randomized allocation with language indicating comparison

of alternative doses or dosing strategies. This is intended to capture whether dose comparison is built into the experimental design near trial launch. We construct another indicator for explicit dose-optimization language. This indicator equals one when the trial record contains phrases such as "dose optimization," "optimal dose," or equivalent language linking optimization explicitly to dosage. We search the brief and detailed study descriptions, arm and intervention descriptions, and registered outcome text. For historical analyses, the same rule is applied to the initial registry version, allowing us to distinguish dose-optimization language present at launch from language introduced later. Broader architecture measures include registered arms and intervention components. Intervention count captures treatment-component intensity. Historical versions permit a further distinction between initial design and subsequent amendment. For randomized-dose comparison, we observe whether the feature appears in the initial version, whether it is present in the current record, and whether it was added after trial initiation.

Table 8: Construction of Principal Empirical Measures

Measure	Construction	Interpretation
New sponsor-asset entry	First observed strict experimental sponsor-drug asset after incorporating the 2000–2014 development history	Entry of a new observable project into human clinical development
Any new sponsor asset	Indicator that at least one new strict experimental asset enters a sponsor-domain-quarter	Primary extensive-margin outcome
Trial starts	Number of early-phase trials beginning in a sponsor-domain-quarter	Intensity of clinical-development activity
Randomized-dose comparison	Initial protocol combines randomized allocation with language indicating comparison of alternative doses or dosing strategies	Primary trial-architecture outcome
Explicit dose optimization	Initial record directly references dose optimization or optimal dose selection	Textual measure of dose optimization as a development objective
Launch arms	Number of registered arm groups in the initial historical version	Breadth of trial architecture at launch
Launch interventions	Number of registered intervention components in the initial version	Treatment-component intensity
High-precision different-dose arms	Same active intervention appears in separate arms with distinct explicit arm-specific doses	Conservative validation of literal multidose architecture
Randomized-dose addition	Randomized-dose comparison absent at launch but present in a later record	Post-launch alteration of the primary design feature
Launch planned enrollment	Estimated enrollment in the first reliable historical record	Patient scale planned near trial initiation
Fixed-horizon enrollment	Most recent enrollment target available at 6, 12, or 18 months after study start	Evolution of planned trial scale
Recruitment exit	First observed transition out of an active recruiting state	Reconstructed end of active participant recruitment

3.9 High-Precision Dose Validation

Dose information embedded in registry text can be noisy. Numeric values may refer to commercial strengths, titration, within-arm escalation, or different drugs within a combination regimen. We construct a validation measure requiring the same active intervention to appear in different arms at different explicit doses. Among 17,784 reliable launch records, 7,350 map

the same active intervention to multiple trial arms. Of those 478 satisfy the high-precision different-dose definition, of which 378 are randomized. This measure focuses on specificity and is used as validation.

Table 9: High-Precision Dose-Architecture Validation

Measure	Trials	Share
Reliable launch records	17,784	100.00%
Same active intervention mapped to multiple arms	7,350	41.33%
High-precision different-dose architecture	478	2.69%
Randomized high-precision different-dose architecture	378	2.13%

3.10 Enrollment Reconstruction

Enrollment requires separate treatment because its interpretation changes as a trial matures. Early registry versions typically report an estimated target. Later records may contain revised targets or an actual participant count. We recover planned enrollment near launch and reconstruct the most recent target available by 6, 12, and 18 months after study start date. Eligibility at each horizon requires sufficient calendar time to have elapsed before the September 15, 2026 data truncation. This produces a common-maturity panel that separates initial patient scale from subsequent revisions and avoids comparing recently launched trials with studies that have already completed recruitment.

Table 10: Historical Enrollment Reconstruction

Enrollment horizon	Eligible trials	Target recovered
Launch	17,744	17,744
6 months	16,879	16,322
12 months	15,840	15,199
18 months	14,727	14,060
Complete 0-, 6-, 12-, and 18-month history	14,059	14,059

3.11 Recruitment Histories

ClinicalTrials.gov does not provide a universal participant recruitment-end date. We reconstruct recruitment cessation from historical status changes. A trial is considered actively

recruiting when its status is "Recruiting" or "Enrolling by Invitation." We identify the first subsequent transition to "Active, not recruiting," "Completed," "Terminated," "Suspended," or "Withdrawn." Our high-confidence measure requires a recruiting state to have been observed before the exit. We distinguish successful recruitment cessation from interrupted exits. Transitions to Active, Not Recruiting or Completed are classified as successful exits. Terminated, Suspended, and Withdrawn trials are treated separately. This prevents premature termination from being interpreted as unusually rapid recruitment. The resulting histories permit comparisons at 6, 12, and 18 months after study start, placing recent and older trials on common observation horizons.

3.12 Trial-Level Descriptive Statistics

Table 11 describes current-record trial characteristics before Project Optimus and after final guidance. The statistics are descriptive rather than identifying comparisons. Non-oncology trials are substantially more likely to use randomized allocation. Oncology trials display a pronounced rise in explicit dose-optimization language after the regulatory transition. Median enrollment rises in both domains, while average arm counts decline slightly.

Table 11: Trial-Level Summary Statistics

	Oncology		Non-oncology	
Characteristic	Pre-2021	Post-final	Pre-2021	Post-final
Trials	14,243	7,066	20,443	8,658
Industry lead sponsor	33.4%	33.3%	51.9%	54.2%
Randomized allocation	22.6%	22.6%	70.8%	66.9%
Randomized-dose language	13.1%	14.6%	37.5%	38.2%
Explicit dose-optimization language	1.6%	4.7%	1.5%	1.9%
Registered enrollment, median	34	46	40	50
Reported arms, mean	2.00	1.88	2.75	2.65
Two or more reported arms	44.7%	43.9%	80.5%	76.6%
Registered interventions, mean	2.50	2.44	2.20	2.17

The oncology and non-oncology trial level domains differ markedly in baseline experimental structure, and the composition and maturity of registered trials change over time. These

differences are evidence to motivate us from a simple comparison of oncology and non-oncology trial levels which would be inadequate. The combined data allow us to observe the research process at several distinct points. The long sponsor-asset history captures whether a new project enters human development. Historical registry versions recover the architecture specified near trial launch. Subsequent versions reveal amendments and enrollment revisions, while status histories trace recruitment. This structure allows the empirical analysis to distinguish project selection from the design, scale, and execution of research conditional on entry.

4 Model Specification

4.1 Research Design

The empirical analysis studies several margins along which firms can respond to changing evidentiary expectations. Sponsors may alter which projects enter clinical development, change the architecture of experiments conditional on entry, revise those experiments after launch, expand patient scale, or change the time required to execute them. A simple comparison of oncology development before and after Project Optimus would be difficult to interpret. Clinical research changed substantially over the same period for reasons unrelated to oncology dose regulation. This issue is addressed by benchmarking oncology activity against contemporaneous non-oncology development conducted by the same sponsors. For the trial-level analysis, the data sample retains sponsor-quarter cells in which a sponsor initiates both oncology and non-oncology trials. This comparison places trials within a common organizational and calendar environment while allowing their exposure to oncology-specific regulatory expectations to differ. For trial i, sponsored by firm s, belonging to research domain d, and beginning in calendar quarter q, we estimate:

$$Y_{isdq} = \beta \left(Oncology_d \times Post_q \right) + \mathbf{X}_i' \gamma + \mu_{sq} + \eta_{sd} + \varepsilon_{isdq}, \tag{2}$$

where Y_{isdq} denotes a trial-design, enrollment, or recruitment outcome. $Oncology_d$ identifies oncology development, while $Post_q$ identifies trials beginning after the regulatory threshold under consideration. $\mathbf{X}_i$ contains development-phase controls. The specification includes sponsor-quarter and sponsor-domain fixed effects. The coefficient β measures the differential change in oncology trials relative to non-oncology trials initiated by the same sponsor during the same period. For binary outcomes, coefficients are interpreted as percentage-point changes. For logged outcomes, they are interpreted approximately as proportional changes. Project Optimus developed over several years rather than through a single isolated announcement. Via the FDA cited timeline: the FDA began the initiative in 2021, published draft dosage-optimization guidance in 2023, and issued final guidance in August 2024 U.S. Food and Drug Administration (n.d., 2024). The FDA describes the initiative as an effort to improve dose optimization and dose selection during oncology drug development, with greater attention to efficacy, safety, and tolerability. We estimate specifications using both the draft- and final-guidance periods rather than interpreting the final guidance date as the first moment at which sponsors could respond.

4.2 Project Entry and the Extensive Margin

Our trial-level analysis conditions on an experiment having been initiated. Regulatory expectations can affect an earlier decision: whether an asset enters clinical development at all. More demanding premarket evidence requirements may increase the expected cost of developing an oncology asset, causing marginal projects to fall below the firm's investment threshold. Further, greater regulatory clarity or earlier resolution of dosage uncertainty may lower downstream development risk taking. The effect on project entry is clearly an empirical question rather than a maintained assumption. We study this margin by reorganizing the ClinicalTrials.gov data into a sponsor-by-domain-by-quarter panel. Sponsors observed in both oncology and non-oncology development define the comparison universe. Within each sponsor's observed development window, we create separate oncology and non-oncology

records for every calendar quarter, including quarters in which no new trial or asset appears. Explicitly retaining these zeros allows us to distinguish changes in project selection from changes in trial design conditional on launch. Our extensive-margin outcome is the entry of a new sponsor-drug asset into registered human clinical development. We normalize drug and biologic intervention names and aliases within sponsor and identify the first registered clinical appearance of each sponsor-drug pair. From these records we construct the number of newly observed assets and an indicator for whether any new asset enters a given domain in a sponsor-quarter. Secondary outcomes include the number of new early-phase trial starts, an indicator for any trial initiation, and the oncology share of new clinical-development activity. These measures distinguish entry of new assets from changes in the number of trials conducted around assets already under development. For sponsor s, research domain d, and quarter q, we estimate:

$$Y_{sdq} = \beta^{E} \left(Oncology_{d} \times Post_{q} \right) + \mu_{sq} + \eta_{sd} + \varepsilon_{sdq}, \tag{3}$$

where Y_{sdq} measures project entry or trial-start activity. The sponsor-quarter component accounts for changes in the sponsor's overall development activity at a given point in time. Sponsor-domain components account for differences in their propensity to conduct oncology and non-oncology research. The parameter β^{E} measures whether new clinical-development activity shifts differentially toward or away from oncology after the regulatory transition. We estimate count outcomes using a Poisson specification and binary entry outcomes using the linear framework. We apply the same alternative regulatory windows and pre-policy placebo exercises used in the trial-level analysis. This interpretation is intentionally limited to observable clinical development. ClinicalTrials.gov does not reveal compounds abandoned during discovery or preclinical development before registration. The extensive-margin analysis measures selection into registered human development rather than the complete pharmaceutical research portfolio.

4.3 Dynamic Response

The static specifications summarize mean post-policy differences but do not show when those differences emerge. We estimate a dynamic version of the trial-level model:

$$Y_{isdq} = \sum_{k \neq -1} \beta_k \left[Oncology_d \times \mathbf{1}\{EventQuarter_q = k\} \right] + \mathbf{X}_i' \gamma + \mu_{sq} + \eta_{sd} + \varepsilon_{isdq}, \qquad (4)$$

where k indexes calendar quarters relative to the regulatory threshold and the quarter immediately preceding the threshold is omitted. The sequence of β_k coefficients provides a diagnostic for differential pre-period movement and traces the evolution of trial design after the regulatory transition. This is particularly useful in the Project Optimus setting because firms could begin adapting development plans before the final guidance was issued. We complement the event study with alternative windows around the regulatory threshold and with specifications beginning at the draft-guidance period. The objective is not to isolate a single announcement effect, but to establish whether the observed change emerges during the period in which dose-optimization expectations were being formalized.

4.4 Initial Design versus Subsequent Amendment

ClinicalTrials.gov is a living registry. A feature observed in the current record may have been part of the experiment when the trial began or may have appeared through a later amendment. Historical versions allow us to distinguish these possibilities. Let Y^L_{isdq} denote a design outcome measured in the initial reliable registered version and Y^C_{isdq} the corresponding outcome in the current record. We estimate Equation 2 separately for the two measures. For binary outcomes, we then define and estimate both:

$$Add_i = \mathbf{1}\left\{ Y_i^L = 0 \text{ and } Y_i^C = 1 \right\}, \qquad (5)$$

$$Add_{isdq} = \beta^{A}\left(Oncology_{d} \times Post_{q}\right) + \mathbf{X}_{i}'\gamma + \mu_{sq} + \eta_{sd} + \varepsilon_{isdq}. \tag{6}$$

This decomposition separates a change in experimental design from a change in registry content after launch. A post-policy difference already present in the initial protocol, accompanied by little change in Add_i, indicates that the response is incorporated into the experimental plan rather than generated by subsequent registry editing.

4.5 Trial Architecture

The extensive-margin analysis asks which projects enter clinical development. We next condition on trial initiation and examine how sponsors organize the experiments they choose to undertake. Our primary architectural outcome identifies randomized dose comparison in the initial registered protocol. The measure combines randomized-design information with language indicating comparison of alternative doses or dosing strategies. It captures whether randomized dosage comparison forms part of the experiment near launch rather than whether the study ultimately succeeds in identifying an optimal dose. We estimate Equation 2 for explicit dose-optimization language and several broader characteristics of trial architecture. These include the number of registered arms, whether the study contains multiple arms, and the number of intervention components. Intervention count measures the number of registered treatment components, which may include experimental drugs, combination partners, comparators, or background therapies. It therefore captures broader treatment complexity rather than the number of distinct dose levels tested. Further, we construct a conservative arm-level measure requiring the same active intervention to appear in distinct trial arms at different explicit doses.

4.6 Enrollment and Experimentation Scale

A change in experimental architecture need not produce a proportional change in the number of patients used to conduct the experiment. We study trial scale separately. For each historical trial, we reconstruct planned enrollment at launch and at 6, 12, and 18 months after Study Start Date. For horizon $h \in \{0, 6, 12, 18\}$, we estimate:

$$\log\left(Enrollment_{isdq}^{(h)}\right) = \beta_h \left(Oncology_d \times Post_q\right) + \mathbf{X}_i'\gamma + \mu_{sq} + \eta_{sd} + \varepsilon_{isdq}^{(h)}. \tag{7}$$

Comparing β_h across horizons distinguishes differences already embedded in the initial enrollment target from changes that emerge as a study matures. We then measure revisions relative to the launch target:

$$Revision_i^{(h)} = \log\left(\frac{Enrollment_i^{(h)}}{Enrollment_i^{(0)}}\right). \tag{8}$$

We estimate the baseline specification using this revision measure and separate indicators for upward and downward changes. These outcomes test whether sponsors begin with similar enrollment plans but subsequently expand or contract them. The fixed-horizon approach addresses an important feature of current registry data. Recent studies are more likely to report estimated enrollment, whereas older studies have had more opportunity to revise their targets or convert enrollment to actual counts. Comparing studies at common maturity horizons reduces this lifecycle mismatch. Studies registered in different calendar periods are observed after the same amount of elapsed time, limiting the possibility that apparent differences reflect unequal opportunities for enrollment revisions, protocol amendments, recruitment progress, or other post-launch changes rather than differences associated with the regulatory period.

4.7 Recruitment Dynamics

Evidentiary requirements can affect development time even when patient targets remain unchanged. The complexities associated with gathering evidence can increase the length of time taken for trial development. Trials requiring more complex experimental designs may require more screening, dosing, monitoring and coordination. There is no single 'end date' provided for recruitment on ClinicalTrials.gov. We recreate the recruitment timeline based upon status changes. A trial is identified as starting active recruitment if it is observed to be Recruiting or Enrolling by Invitation. Changes in status to Active, Not Recruiting; Completed; Terminated; Suspended or Withdrawn are then identified. Successful and unsuccessful and abandoned recruitment are identified to ensure trials that are stopped early are not automatically assumed to have a high recruitment rate. Over the horizon of $h \in \{6, 12, 18\}$ months, we define:

$$R_i^S(h) = \mathbf{1}\left\{\text{successful recruitment exit by } h\right\}, \tag{9}$$

$$R_i^F(h) = \mathbf{1}\left\{\text{failed or interrupted exit by } h\right\}. \tag{10}$$

We estimate Equation 2 separately for the two outcomes at each horizon. The fixed-horizon design places recent and older trials on more comparable observation windows without imposing a constant relationship over the full duration of recruitment.

4.8 Pre-Policy Placebo Tests

The interpretation of the estimates depends on the absence of comparable oncology-specific breaks before Project Optimus. We re-estimate the principal specifications using 2017, 2018, 2019, and 2020 as pseudo-treatment dates. For placebo year p, we estimate:

$$Y_{isdq} = \beta_p^P \left(Oncology_d \times Post_{pq}^P\right) + \mathbf{X}_i'\gamma + \mu_{sq} + \eta_{sd} + \varepsilon_{isdq}. \tag{11}$$

These exercises are useful for registry-derived architecture measures. Clinical trial design changes over time even without a discrete regulatory intervention, and some candidate outcomes may contain pre-existing oncology-specific trends. We assign the strongest interpretive weight to outcomes that move during the Project Optimus period without generating comparable breaks at earlier pseudo-treatment dates. The same logic applies to the extensive margin. If oncology project entry was already moving differentially relative to non-oncology development before 2021, a post-Project Optimus change in entry would be difficult to attribute to the new evidentiary environment.

4.9 Guidance-Scope Restriction

The August 2024 guidance does not cover every oncology product category. FDA states that the guidance does not address dosage optimization for radiopharmaceuticals, cellular and gene therapy products, oncolytics, microbiota products, or cancer vaccines and does not specifically address pediatric drug development. We construct a narrower oncology sample that removes these categories using registered interventions and protocol text. We use this restriction as a robustness exercise rather than a sharp treatment boundary. Regulatory expectations may extend beyond the literal scope of a guidance document, and registry-based modality classification is necessarily imperfect. The relevant question is whether the principal results weaken materially when the oncology sample is restricted toward the products most directly addressed by the guidance. The models trace adjustment at several stages in the production of clinical evidence. The extensive-margin specification asks whether changing evidentiary expectations alter which oncology assets enter registered human development. Conditional on entry, the architecture models examine what sponsors choose to compare when a trial begins and whether those choices are subsequently revised. The enrollment and recruitment models then test whether changes in experimental design are accompanied by greater patient scale or longer execution. A higher expected information requirement can change the profitability of undertaking a project, the organization of the experiment used to

study it, or the resources devoted to its execution. These responses need not move together. Firms may absorb a new evidentiary burden by redesigning experiments without abandoning projects, or project selection itself may change if the expected costs and benefits of clinical development shift sufficiently. The empirical framework defines four margins of adjustment: entry into clinical development, experimental architecture, patient scale, and execution. This allows the analysis to ask not merely whether oncology development changed during Project Optimus, but where in the production of clinical evidence the response occurred.

5 Results

5.1 Project Entry and Development Activity

We begin our analysis with the extensive margin of clinical development. The sponsor-domain-quarter panel allows us to observe not only trials and assets that enter development, but quarters in which a sponsor initiates no new project in a given domain. Our analysis focuses on the local window around the August 2024 final guidance. Table 12 reports the post-final estimates. Following final guidance, the probability that an industry sponsor initiates at least one new oncology asset is estimated to be 3.47 percentage points lower relative to its contemporaneous non-oncology development. Relative to the 18.8 percent oncology entry rate during the draft-to-final period, the magnitude is equivalent to approximately 19 percent of the pre-final entry probability. The corresponding estimate for the number of new assets is negative as well. The coefficient implies approximately 11 percent fewer new oncology assets relative to non-oncology development, although its confidence interval includes zero. We place greater weight on the binary entry result.

Table 12: Post-Final Changes in Project Entry and Trial Initiation

Outcome	Estimate	95% CI	Implied change	Observations
New strict experimental sponsor-assets	-0.1196	[-0.3403, 0.1011]	-11.3%	11,454
Any new strict experimental sponsor-asset	-0.0347	[-0.0613, -0.0080]	-3.47 pp	18,002
Trials with a strict experimental focal asset	-0.1908	[-0.3397, -0.0420]	-17.4%	14,558
Any trial with a strict experimental focal asset	-0.0574	[-0.0863, -0.0285]	-5.74 pp	18,002
All early-phase trial starts	-0.1872	[-0.3302, -0.0442]	-17.1%	15,431
Any early-phase trial start	-0.0593	[-0.0885, -0.0302]	-5.93 pp	18,002

These estimates describe a relative change rather than an absolute contraction in oncology development. The data show that the probability of new oncology asset entry rises from 18.8 percent during the draft-to-final period to 20.8 percent after final guidance. Over the same periods, however, the corresponding non-oncology probability rises from 18.2 percent to 22.3 percent. The estimated extensive-margin response reflects slower growth in oncology entry relative to the same sponsors' non-oncology portfolios. Trial initiation exhibits a larger relative difference. The estimated number of early-phase oncology trial starts is approximately 17 percent lower relative to non-oncology development after final guidance, while the probability of initiating any trial is lower by about 5.9 percentage points. We interpret these trial-frequency results more cautiously than the new-asset measure because some earlier pseudo-treatment dates produce evidence of differential movement in trial initiation. The asset-entry outcome performs substantially better in the pre-policy diagnostic exercises.

5.2 Randomized Dose Comparison

The most conspicuous evidence of change in oncology trial architecture appears in randomized dose comparison. Table 13 reports estimates using the initial registered version of each trial. Following the August 2024 final guidance, oncology trials are 4.95 percentage points more likely to incorporate randomized-dose comparison relative to contemporaneous non-oncology

development by the same sponsors. This is the primary measure of whether sponsors changed the comparisons built into early-stage oncology experiments. The response is not confined to the period following publication of the final guidance. Using January 2023 as the beginning of the post-draft period produces an estimate of 5.07 percentage points. The similar magnitudes across the draft and final-guidance specifications suggest that trial design began adjusting during the broader Project Optimus transition rather than only after the final guidance was issued. The result is similarly robust to narrowing the oncology sample toward products covered more directly by the final guidance. Within this guidance-covered proxy, the post-final estimate is 5.07 percentage points. The primary result does not appear to be driven by oncology modalities outside the central scope of the guidance.

Table 13: Post-Final Changes in Initial Trial Architecture

Outcome	Estimate	95% CI	Implied change
Randomized-dose comparison	0.0495	[0.0103, 0.0887]	4.95 pp
Explicit dose-optimization language	0.0239	[0.0025, 0.0453]	2.39 pp
Log registered interventions	0.0693	[0.0153, 0.1232]	7.17%
Two or more registered arms	0.0406	[0.0012, 0.0801]	4.06 pp
High-precision different-dose arms	0.0065	[-0.0055, 0.0186]	0.65 pp
High-precision randomized different-dose arms	0.0074	[-0.0035, 0.0184]	0.74 pp

Explicit dose-optimization language provides a separate measure of the same shift in development strategy. Oncology trials become 2.39 percentage points more likely to state dose optimization directly in their initial registered protocol. Unlike several broader measures of trial complexity, this outcome does not exhibit a comparable break at the pre-Project Optimus placebo dates. The result corroborates the randomized-dose finding using a measure based on the stated objective of the trial rather than its allocation structure. The remaining architecture measures describe how the broader experimental design changes, but they provide less direct evidence of dose comparison. Registered intervention count increases by 7.2 percent following final guidance. An intervention in ClinicalTrials.gov can represent a drug, biologic, combination component, or other treatment element rather than a distinct dosage. This estimate is evidence of increased treatment-component complexity, and is not a

direct count of additional dose levels. The probability of registering at least two arms rises by approximately 4.1 percentage points. This measure is less persuasive as a policy response because a comparable positive break appears around the 2018 placebo threshold. It is useful for describing the structure of post-Optimus oncology trials.

The high-precision dose measures provide a conservative check. These measures require the same active intervention to appear in separate trial arms at explicitly different doses. The estimated post-final changes are positive but small: 0.65 percentage points for different-dose architecture and 0.74 percentage points when the design is additionally required to be randomized. Both confidence intervals include zero. As shown in our summary statistics, these definitions identify only a small subset of trials. The strongest evidence does not suggest a uniform expansion in every observable dimension of trial complexity. It is concentrated in whether sponsors explicitly design trials to compare doses and whether dose optimization appears as a stated development objective. Project Optimus-era oncology trials therefore appear to differ most clearly in whether they are designed to distinguish among clinically relevant doses and characterize dose-response relationships, rather than simply to identify a tolerable dose or establish preliminary activity.

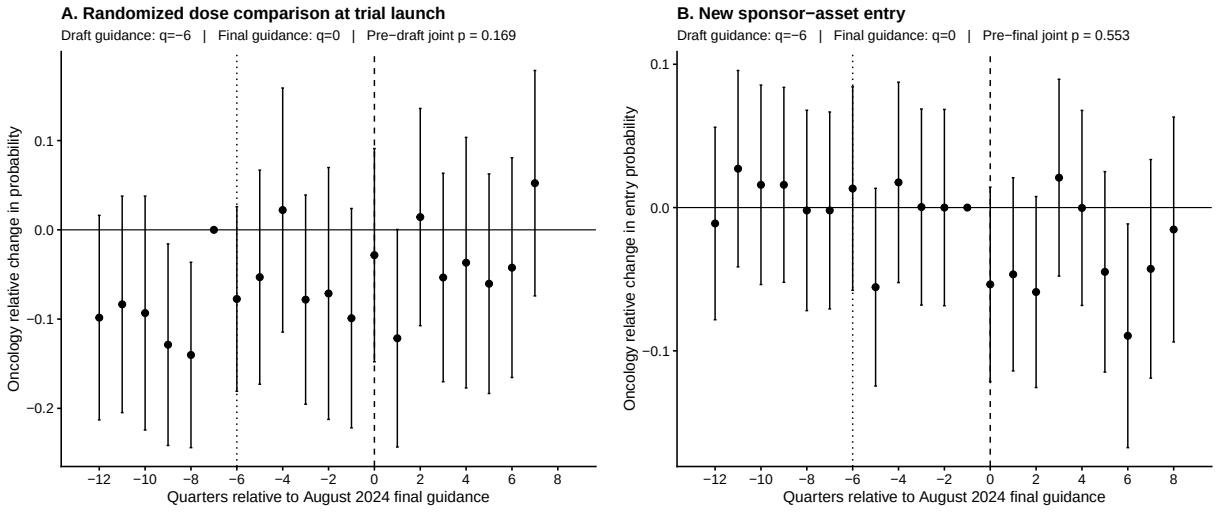


Figure 1: Project Selection and Experimental Design during Project Optimus

Figure 1 places the two principal responses on a common regulatory timeline. The dotted

vertical line delineates the January 2023 draft guidance while the dashed line delineates the August 2024 final guidance. Panel A is normalized to 2022Q4, the quarter immediately preceding the draft guidance, while Panel B is normalized to 2024Q2, the quarter immediately preceding final guidance; vertical bars show 95% confidence intervals. Randomized-dose comparison begins to change during the period surrounding the draft guidance, before publication of the final guidance. The relative change in new oncology asset entry becomes more apparent around the final-guidance period. The dynamic evidence comports with adjustment occurring at different points in the development decision: experimental architecture responds during the broader regulatory transition. The project-selection response emerges later.

5.3 Trial Launch

The historical registry allows us to determine whether randomized-dose comparison was incorporated into the trial when it began or appeared through later amendments. Table 14 presents this decomposition. The post-final estimate in the initial registered protocol is 4.95 percentage points. Measuring the same outcome in the current version of those records produces an estimate of 4.82 percentage points. Further, there is essentially no change in the probability that randomized-dose language is added after launch.

Table 14: Initial Design versus Subsequent Registry Amendment

Outcome	Estimate	95% CI
Randomized-dose comparison in initial protocol	0.0495	[0.0103, 0.0887]
Randomized-dose comparison in current record	0.0482	[-0.0009, 0.0973]
Randomized-dose language added after launch	0.0016	[-0.0338, 0.0371]

The timing sharpens the interpretation of the architecture result. Sponsors do not appear to generate the post-Project Optimus difference by adding randomized-dose terminology to existing registry records after study initiation. The differential is already present in the experimental plan registered near launch. The principal adjustment occurs when the experiment is constructed: sponsors change what the trial is designed to compare rather than

merely changing how an existing study is described.

5.4 Trial Architecture versus Trial Scale

The increase in randomized-dose comparison does not translate cleanly into a proportional increase in patient scale. Current ClinicalTrials.gov records at first appear to suggest substantially larger oncology trials after the final guidance. That relationship weakens, however, as the sample is restricted to trials with greater opportunity to mature. Table 15 reports the post-final current-enrollment estimate after imposing successively longer minimum follow-up requirements. With no maturity requirement, the estimate corresponds to approximately 18.9 percent higher enrollment. Requiring at least 6 months of maturity reduces the estimate to 14.7 percent. At 12 months the difference falls to 9.7 percent and is no longer precisely estimated. At the 18 month mark it falls to 5.7 percent.

Table 15: Current Enrollment and Trial Maturity

Minimum maturity	Log estimate	Implied change	95% CI	Trials
None	0.1730	18.89%	[8.60%, 30.15%]	19,576
6 months	0.1370	14.68%	[4.15%, 26.26%]	18,693
12 months	0.0926	9.71%	[-2.26%, 23.14%]	17,565
18 months	0.0551	5.67%	[-9.06%, 22.78%]	16,387

The monotonic attenuation suggests that the large current-record estimate is partly a feature of recent-cohort composition. Newly launched trials are more likely to retain estimated enrollment targets, while older trials have had more time to revise those targets or convert enrollment to actual counts. Historical enrollment records provide a cleaner comparison. We place trials on common maturity horizons and reconstruct the enrollment target registered at launch, 6 months, 12 months, and 18 months after study start. Table 16 reports these estimates. None indicates a meaningful differential expansion in planned enrollment following the final guidance. The launch estimate is approximately 1 percent, and the estimates at 6, 12, and 18 months remain close to zero.

Table 16: Post-Final Differences in Historical Enrollment Targets

Enrollment horizon	Log estimate	95% CI	Trials
Launch	0.0103	[-0.1260, 0.1470]	12,366
6 months	-0.0201	[-0.1590, 0.1190]	12,366
12 months	-0.0010	[-0.1460, 0.1440]	12,366
18 months	-0.0500	[-0.2070, 0.1070]	12,366

The same conclusion emerges from the recorded enrollment revisions. Relative to the launch target, the post-final difference is approximately -1.8 percent at 6 months, -0.8 percent at 12 months, and -5.9 percent at 18 months. None provides evidence of systematic post-launch target expansion. The enrollment results change the interpretation of the current-record association. Recent oncology trials may display larger registered enrollment values, but the historical records do not indicate a comparable increase in planned patient scale among trials observed at common stages of maturity.

5.5 Recruitment Dynamics

A more information-intensive design could increase development burden without changing the enrollment target. We examine whether oncology trials became slower to exit active recruitment. Table 17 reports fixed-horizon estimates. Successful recruitment exit becomes slightly less common by six and twelve months, but the differences are small and imprecise. By 18 months the estimated difference is essentially zero. Failed or interrupted exits show no clear pattern at 6 or 12 months and a positive estimate at 18 months.

Table 17: Post-Final Recruitment Outcomes at Common Horizons

Horizon	Successful exit	Failed or interrupted exit	Trials
6 months	-0.0235	-0.0006	12,048
12 months	-0.0292	0.0040	11,313
18 months	-0.0048	0.0338	10,609

The recruitment evidence reported does not show a clear slowdown accompanying the change in trial architecture. Longer horizons become progressively harder to study because

the post-final cohorts have not yet accumulated sufficient follow-up. For this reason, we do not use the small number of trials observed for twenty-four months to describe longer-run recruitment effects.

5.6 Pre-Policy Placebos

The paper's two principal margins perform well in pre-Project Optimus placebo exercises. Randomized-dose comparison remains close to zero when pseudo-treatment is assigned to 2017, 2018, 2019, or 2020. The new-asset entry measure likewise produces no statistically distinguishable break at the 2018, 2019, or 2020 placebo thresholds.

Table 18: Pre-Project Optimus Placebos for Main Outcomes

Outcome	Pseudo-treatment year	Estimate	95% CI includes zero
Panel A: Randomized-dose comparison			
	2017	-0.0179	Yes
	2018	-0.0096	Yes
	2019	-0.0025	Yes
	2020	0.0131	Yes
Panel B: Any new strict experimental sponsor-asset			
	2018	-0.0143	Yes
	2019	-0.0174	Yes
	2020	-0.0286	Yes

The local event-study diagnostics around final guidance are similarly reassuring for project entry. A joint test of the eleven pre-treatment coefficients yields a p-value of 0.553 for the probability of any new asset and 0.583 for the number of new assets. These diagnostics reinforce the pre-policy placebo evidence: the cleaner sponsor-asset construction does not display an obvious differential break before the final-guidance period. The broader trial-start outcomes warrant more caution. Several earlier pseudo-treatment dates generate differential changes in trial initiation even though the local pre-final event-study diagnostics are considerably cleaner. For this reason, we place greater weight on new-asset entry when interpreting the extensive-margin response.

The architecture outcomes exhibit a similar hierarchy. Explicit dose-optimization language remains quiet across the pre-2021 placebo dates, while the high-precision randomized different-dose measure produces no comparable pre-policy break. In contrast, the indicator for two or more arms generates a positive difference around the 2018 placebo threshold, and the biopsy measure rejects the null around 2020. Historical intervention count performs better, although its 2018 placebo estimate approaches conventional significance thresholds. The evidence provided by the placebo data further supports the differences between the outcomes. The most specific evidence was obtained using the randomised dose comparison, dose optimisation and new sponsor assets. The higher level information about the number of trials and complexity of experiments provides useful information on how the development activity has changed, however is less impactful as standalone evidence for a regulatory response.

6 Conclusion

We examine Project Optimus which provides a useful setting for studying how firms reorganize clinical research when regulators raise the informational expectations placed on premarket development. The evidence points to adjustments along two distinct margins. Following the final guidance, new oncology asset entry grows more slowly than contemporaneous non-oncology development within the same industry sponsors. The probability that a sponsor initiates at least one new oncology asset is estimated to be approximately 3.5 percentage points lower relative to its non-oncology portfolio, a differential equal to roughly 19 percent of the pre-final oncology entry rate. Importantly, oncology entry continues to rise in absolute terms. The result reflects relative reallocation across development portfolios rather than an absolute contraction in oncology research. Further, we report that conditional on entering development, oncology trials become more likely to incorporate dose comparison into their initial experimental design. Randomized-dose comparison rises by roughly 5 percentage points relative to contemporaneous non-oncology development by the same sponsors, while

explicit dose-optimization language increases by roughly 2.5 percentage points. Historical registry records show that the randomized-dose difference is already present near trial launch. The probability that comparable language is added through subsequent registry amendment changes little. The principal response appears when the experiment is constructed rather than through later modification of the registry.

The change in experimental architecture is not accompanied by equally clear growth in patient scale. Large enrollment differences visible in current ClinicalTrials.gov records attenuate as trials are required to have more time to mature. Enrollment targets reconstructed at common post-launch horizons show little evidence of proportional expansion, and recruitment histories provide no clear evidence of a persistent slowdown. More intensive dose comparison therefore does not appear to require a commensurate increase in planned enrollment or an obvious extension of active recruitment. Evidentiary regulation need not operate only by increasing the quantity of research that firms conduct. It can alter both the composition of projects entering clinical development and the structure of the research conducted within those projects. Higher information requirements may lead firms to reallocate development effort across projects while redesigning experiments to answer different questions. In our setting, these adjustments occur without a comparably clear increase in the number of patients planned for study. This distinction matters for regulatory burden measurement. Larger samples, longer trials, and greater research expenditures capture part of the adjustment available to firms. Project selection and experimental design provide additional margins through which regulatory expectations can reshape the production of evidence. Under Project Optimus, the clearest changes appear in which oncology projects advance relative to competing development opportunities and in whether trials are designed to compare clinically relevant doses, rather than in a simple expansion of trial size.

The post-final cohorts are still immature. This limits what can be learned about realized enrollment, completion, development cost, approval and postmarketing evidence requirements. The present analysis identifies changes in observed clinical-development choices rather than

their welfare consequences. This could be explored further in future research by focusing on whether more detailed dose information leads to better decisions on which doses to use, less uncertainty, altered probabilities of approval and/or different times and costs. Our findings provide reliable evidence that FDA regulation can shape scientific output by influencing which projects firms pursue and how experiments are designed, thereby affecting the amount, allocation, and purpose of research conducted by firms.

References

Ananthakrishnan, Revathi, Patrick Mitchell, Dong Xi, Gu Mi, Gina D'Angelo, Haijun Ma, Yujun Wu, and Philip He. 2026. "Implementing Project Optimus in Oncology Dosage Optimization: Where Are We Now?" *Therapeutic Innovation & Regulatory Science* 60(5): 1542–1559. https://doi.org/10.1007/s43441-026-01010-1.

Bhamidipati, Deepak, Suzanne Jones, David Spigel, Gerald Falchook, Manish Patel, Howard Burris, and Vivek Subbiah. 2025. "Impact of Food and Drug Administration Project Optimus Guidance on Design of Early-Phase Clinical Trials." *JCO Oncology Advances* 2(1): e2500084. https://doi.org/10.1200/OA-25-00084.

Brock, Kristian, Victoria Homer, Gurjinder Soul, Claire Potter, Cody Chiuzan, and Shing Lee. 2021. "Is More Better? An Analysis of Toxicity and Response Outcomes from Dose-Finding Clinical Trials in Cancer." *BMC Cancer* 21: 777. https://doi.org/10.1186/s12885-021-08440-0.

Chandra, Amitabh, Jennifer Kao, Kathleen L. Miller, and Ariel D. Stern. 2026. "Regulatory Incentives for Innovation: The FDA's Breakthrough Therapy Designation." *Review of Economics and Statistics* 108(2): 470–484. https://doi.org/10.1162/rest_a_01434.

Getz, Kenneth A., and Rafael A. Campo. 2018. "New Benchmarks Characterizing Growth in Protocol Design Complexity." *Therapeutic Innovation & Regulatory Science* 52(1): 22–28. https://doi.org/10.1177/2168479017713039.

Getz, Kenneth A., Stella Stergiopoulos, Mary Short, Leon Surgeon, Randy Krauss, Sybrand Pretorius, Julian Desmond, and Derek Dunn. 2016. "The Impact of Protocol Amendments on Clinical Trial Performance and Cost." *Therapeutic Innovation & Regulatory Science* 50(4): 436–441. https://doi.org/10.1177/2168479016632271.

Glasmacher, Axel, Elena Garralda, Chad Gwaltney, Katrin Rupalla, Chunze Li, and Harald Weber. 2025. "Dose Optimization in Cancer Drug Development: Review and Outcome

of a Multi-Stakeholder Workshop." *European Journal of Cancer* 226: 115593. https://doi.org/10.1016/j.ejca.2025.115593.

Hermosilla, Manuel. 2024. "Regulating Ethical Experimentation: Impacts of the Breakthrough Therapy Designation on Drug R&D." *Journal of Health Economics* 94: 102855. https://doi.org/10.1016/j.jhealeco.2023.102855.

Minasian, Lori, Owen Rosen, Dalia Auclair, Atiqur Rahman, Richard Pazdur, and Richard L. Schilsky. 2014. "Optimizing Dosing of Oncology Drugs." *Clinical Pharmacology & Therapeutics* 96(5): 572–579. https://doi.org/10.1038/clpt.2014.153.

Murciano-Goroff, Yonina R., Sean M. Devlin, Alexia Iasonos, and Alexander Drilon. 2024. "Optimus-Era Dose Finding for Rare Cancers." *Cancer Discovery* 14(6): 909–914. https://doi.org/10.1158/2159-8290.CD-24-0368.

Murphy, Ravindhi, Sarah Halford, and Stefan Nicholas Symeonides. 2023. "Project Optimus, an FDA Initiative: Considerations for Cancer Drug Development Internationally, from an Academic Perspective." *Frontiers in Oncology* 13: 1144056. https://doi.org/10.3389/fonc.2023.1144056.

Rahman, Atiqur, Mirat Shah, and Stacy S. Shord. 2024. "Dosage Optimization: A Regulatory Perspective for Developing Oncology Drugs." *Clinical Pharmacology & Therapeutics* 116(3): 577–591. https://doi.org/10.1002/cpt.3373.

Samineni, Divya, Karthik Venkatakrishnan, Ahmed A. Othman, Yazdi K. Pithavala, Srinivasu Poondru, Chirag Patel, Pavan Vaddady, Wendy Ankrom, Saroja Ramanujan, Nageshwar Budha, Michael Wu, Nahor Haddish-Berhane, Holger Fritsch, Azher Hussain, Jitendra Kanodia, Meng Li, Mengyao Li, Murad Melhem, Apurvasena Parikh, Vijay V. Upreti, and Neeraj Gupta. 2024. "Dose Optimization in Oncology Drug Development: An International Consortium for Innovation and Quality in Pharmaceutical Development White Paper."

Clinical Pharmacology & Therapeutics 116(3): 531–545. https://doi.org/10.1002/cpt.3298.

Shah, Mirat, Atiqur Rahman, Marc R. Theoret, and Richard Pazdur. 2021. "The Drug-Dosing Conundrum in Oncology—When Less Is More." *New England Journal of Medicine* 385(16): 1445–1447. https://doi.org/10.1056/NEJMp2109826.

Stern, Ariel Dora. 2017. "Innovation under Regulatory Uncertainty: Evidence from Medical Technology." *Journal of Public Economics* 145: 181–200. https://doi.org/10.1016/j.jpubeco.2016.11.010.

Thall, Peter F., Elizabeth Garrett-Mayer, Nolan A. Wages, Susan Halabi, and Ying Kuen Cheung. 2024. "Current Issues in Dose-Finding Designs: A Response to the U.S. Food and Drug Administration's Oncology Center of Excellence Project Optimus." *Clinical Trials* 21(3): 267–272. https://doi.org/10.1177/17407745241234652.

U.S. Food and Drug Administration. n.d. "Project Optimus: Reforming the Dose Optimization and Dose Selection Paradigm in Oncology." Oncology Center of Excellence. Accessed September 23, 2026. https://www.fda.gov/about-fda/oncology-center-excellence/project-optimus.

U.S. Food and Drug Administration. 2024. *Optimizing the Dosage of Human Prescription Drugs and Biological Products for the Treatment of Oncologic Diseases: Guidance for Industry*. Oncology Center of Excellence, Center for Drug Evaluation and Research, and Center for Biologics Evaluation and Research. August 2024. https://www.fda.gov/media/164555/download.

Yuan, Ying, Heng Zhou, and Suyu Liu. 2024. "Statistical and Practical Considerations in Planning and Conduct of Dose-Optimization Trials." *Clinical Trials* 21(3): 273–286. https://doi.org/10.1177/17407745231207085.