
Introduction: Economic Properties of Creative Activities

Economists have studied a number of industrial sectors for their special and distinctive features: pharmaceuticals and computer chips for competition to innovate; chemical process industries for rivalry to install new capacity; food processing for product differentiation and the rise of dominant brands. Indeed, few sectors have escaped notice that display some distinctive form of competitive behavior or pose distinctive problems for public policy. One has been largely missed, however—the “creative” industries supplying goods and services that we broadly associate with cultural, artistic, or simply entertainment value. They include book and magazine publishing, the visual arts (painting, sculpture), the performing arts (theatre, opera, concerts, dance), sound recordings, cinema and TV films, even fashion and toys and games. So far, economists exploring this area have mainly focused on public subsidy for the elite performing arts.

While economists have been game to think about public policy toward creative activities, they have largely ignored questions about why those activities are organized the way they are. Artists of all types engage in creative processes and tasks that come to completion only with the collaboration of “humdrum” (or “ordinary”) partners, and perhaps of other artists as well. The painter needs an art dealer, the novelist a publisher. The cinema film requires a number of actors, a director, screenwriter, cinematographer, production designer, make-up specialist, and many others who see themselves in some measure as artists (along with teamsters and accountants, who likely do not). These collaborations rest on deals and contracts—perhaps of the “handshake” variety, perhaps elaborately drawn. Some artistic production occurs within continuing organizations—a major symphony orchestra, or a Hollywood film studio in the 1930s. Other creative wares emerge from one-time deals. This book is about the organization of creative activities: why deals and contracts are structured the way they are; why some creative activi-

ties occur in ongoing organizations ("firms"), and others in one-off deals ("the market").

Economic analysis has tools for answering these questions, but they have seen little application to the creative industries. The theory of contracts addresses why self-interested parties structure their deals as they do. The field of industrial organization studies (among other questions) why transactions occur within continuing firms or between independent parties and why those firms are few or many, operating in one market or several. In this book, I attempt an integrated attack on the organization of creative activities using these tools. There is much to be done, and the goal could be pursued in several ways. We could identify a few choice, distinct questions and wring them dry. Or we could block out the bedrock properties of creative activities and show how they drive organizational patterns in the many markets with substantive creative elements. The latter procedure is the weapon of choice. This chapter proposes some fundamental properties and sketches how they will be used throughout the book to explain why creative activities are organized as they are. A final section provides a primer on contract theory, the most important branch of economics for this study and the one least well represented in most introductory courses.

Economists writing about the arts (for example, William D. Grampp in his *Pricing the Priceless*) often take up the burden of proving that artistic workers make the same sorts of economic choices as do humdrum mortals who profess no creative urge or skill.¹ Here I assume that creative workers are purposive and intendedly rational in their activities—like everyone else. They may have different tastes. They may be less entrepreneurial or less skilled in negotiation, but that remains to be seen. What I stress instead is that creative goods and services, the processes of their production, and the preferences or tastes of creative artists differ in substantial and systematic (if not universal) ways from their counterparts in the rest of the economy where creativity plays a lesser (if seldom negligible) role. These differences rest on the bedrock properties of these activities that distinguish them from other sectors of the economy, and in some cases distinguish creative activities from one another.

Basic Economic Properties of Creative Activities

Demand Is Uncertain

There is great uncertainty about how consumers will value a newly produced creative product, short of actually producing the good and placing it before them. It might meet acclaim and bring in revenue far exceeding its cost of production, or it might find few customers who place any positive value on it. If the creative product is costly to produce (a movie, rather than a painting),

producers will try however they can to learn whether buyers' valuations will be high or low, before all of the product's costs have been incurred. Research and pretesting are largely ineffective, however, because a creative product's success can seldom be explained even *ex post* by the satisfaction of some pre-existing need.² The problem worsens when the costs are sunk, as they usually are, and cannot be retrieved once a disaster is evident. This property implies that the risk associated with any given creative product is high, and that ways of allocating or sharing it will be important for the organization of production. This is what one Hollywood observer called the *nobody knows* property.³ One study of the success of TV series demonstrated it dramatically. The TV networks use a number of signals about new programs to attract advertisers, and the networks themselves heed those signals in deciding which pilot series to select for the season's schedule. Nonetheless, none of these signals has any statistical ability to predict what series will actually succeed.⁴

Economics pays a great deal of attention to sellers and buyers who are stuck with problems of asymmetrical information: the buyer of a used car who suspects that the seller conceals his superior knowledge about its defects; the seller of a gourmet food product who can prove its excellence only by begging customers to buy and try it. A creative product is an "experience good" like these, but the buyer's satisfaction will be a subjective reaction. The producer's intimate knowledge of the good's production process still leaves him in the dark about whether customers will like it: *nobody knows*. The organizational problem is to deal with symmetrical ignorance, not asymmetrical information. A different sort of asymmetry, however, proves very important. Complex creative products (cinema films, popular music albums) proceed from conception to finished product in a series of stages, with costs at each stage completely sunk when the product moves to the next stage. But fresh news does arrive along the way about the good's market prospects. When Joe hands the unfinished good along to Mary, Joe's input to the project is a sunk cost, but Mary can still adjust or even withhold her input. If Joe and Mary at that moment both receive fresh news about the project's ultimate success, it is efficient for Mary to hold decision rights about whether and how to proceed (Joe has no more choices open).⁵ This asymmetry explains the pervasive use in creative industries of the option contract, by which Mary buys from Joe the option to proceed if she wishes, once she has digested any fresh news.

Creative Workers Care about Their Product

Economists normally assume that workers hired for some job do not care about the traits and features of the product they turn out. They care about their pay and working conditions and how much effort they must exert, but

not the output's style, color, or features. Skilled craftspersons often do express pride in or concern for the quality of their work and the goods they turn out, but economists seldom see this interest as affecting the organization of production. In creative activities, however, the creator (artist, performer, author) cares vitally about the originality displayed, the technical prowess demonstrated, the resolution and harmony achieved in the creative act. While these concerns with artistic achievement bear some relation to consumers' ultimate reception of the product, the relationship need not be close. The musician may value the achievement of some finesse of execution that will elude the typical concertgoer, although it *will* be noticed by fellow music professionals. Hence, the artist may divert effort from aspects of the task that consumers will notice (thus affecting their willingness to pay) to those they will neither notice nor value. Throughout this book we distinguish sharply between creative and humdrum inputs. Wherever they work, humdrum inputs demand a wage at least equal to what they earn in the outside market for inputs of their type. They do not care who employs them or what task (within their competence) they are asked to undertake. They are just in it for the money.

The problem of coupling creative effort with humdrum commerce goes still deeper. The view of creative inspiration inherited from romanticism holds that the artist creates out of inner necessity. Imagination and passion carry their own warrant and should not compromise with reason and established practice. Successful imitation of a master, once considered a worthy achievement, becomes an act of cowardice and sloth. Art claims a superior reality that separates the artist from the craftsman.⁶ Asked to cooperate with humdrum partners in some production process, the artist is disposed to forswear compromise and to resist making commitments about future acts of artistic creation or accepting limitations on them. The rub is that resources are scarce, and compromise is hence often unavoidable. Rejecting it on principle distracts one's mind from making the best deal available.

The artist's tastes bear not only on how the creative task is performed, but also on how much effort is allocated to creative work. The artist typically has the alternative of a day job, and indeed is likely to possess valuable skills that could command high pay. A taste for creative work increases the amount of effort supplied by diverting it from humdrum tasks: the "starving artist" syndrome. The prevalence and strength of tastes that affect the qualities and quantity of creative effort we call the *art for art's sake* property.

This property implies that artists turn out more creative product than if they valued only the incomes they receive, and on average earn lower pecuniary incomes than their general ability, skill, and education would otherwise warrant. Less obviously, when the worker cares deeply about the traits of the

product, the problem of organizing its production is fundamentally changed. The entrepreneur who organizes purely humdrum production faces this problem: Can I recruit inputs needed to turn out a given product at a cost less than what buyers will pay for it? In creative activities the good news for the entrepreneur is that creative inputs come cheaply. The bad news is that the traits of the product and the terms of employment of the creative inputs must be negotiated at the same time, and with persons unwilling and perhaps unable to precommit their creative choices. Artists find it difficult after the event to explain their aesthetic choices, let alone stipulate them in advance.⁷

The fact that the artist works outward to realize and reify an inner vision partly explains why *nobody knows*. The artist does not know and cannot pretest whether her creative vision will prove equally compelling to others. Still worse, she cannot tell whether her conception has been successfully extracted from her inner vision and turned into an external creative product. The quality of the vision and the effectiveness of its realization are both up for grabs.

Some Creative Products Require Diverse Skills

Some creative outputs need only a single creative worker: the artist who paints a canvas. Many, however, require diverse skilled and specialized workers, each bringing personal tastes with regard to the quality or configuration of the product. A cinema film results from the efforts of many different artists, each with different skills and aesthetic values and thus with potentially conflicting priorities and preferences that are pressed upon the director who coordinates their efforts. The diverse tastes and preferences obviously complicate the deal for organizing the activity, as implied by the *art for art's sake* property. Formal contracts are largely infeasible, and we find in complex creative activities that conflicting preferences of collaborating artists often get resolved by the rank-order mechanism of "muscle."

A further vital complication results, however, if all of these inputs must perform at or above some level of proficiency and conformance for a viable product to result. The creative activity then involves what economists call a multiplicative production function. A production function is simply a recipe for combining specified inputs to obtain an output. We usually suppose that inputs are substitutable: you can produce a gadget with two units of labor and one of capital, or one of labor and two of capital; perhaps four units of labor and no capital will do the job. In a multiplicative production relationship, however, every input must be present and do its job—or at least perform at or above some threshold level of proficiency—if any commercially valuable output is to result. A large number multiplied by zero is still zero. This is what Michael Kremer called the O-rings theory of production.⁸ Here we call

it the *motley crew* property. It has a number of implications about selecting the team of inputs for a given creative product and sustaining all team members' cooperation while production is under way.

Differentiated Products

When a creative product goes to consumers for their assessment, they typically value it in comparison to similar creative products. Seeing a movie and reading a book are alternative ways to spend the evening. More important for the organization of creative activities, one movie gets chosen among all those that are playing (or on rental at the video store). No two are identical, but they can and do differ in two ways that themselves have diverse consequences. Once they have experienced both, buyers may agree that product A is better than product B, what an economist calls vertically differentiated. If A and B are sold at the same price, nobody will buy product B. Cultural outputs also can differ in their traits, moods, styles, whatever, in ways independent of buyers' overall judgments about their quality levels. Two songs, two paintings, two "action" movies may be quite similar in the character and quality that consumers see in them, but they are not identical. In economic terms they are horizontally differentiated. When horizontally differentiated products are sold at the same price, some people will prefer one, some another. Creative products generally exhibit some mixture of vertical and horizontal differentiation. Most creative products can differ from one another in many ways: paintings, for example, may vary in size, color, type of imagery, skill of draftsmanship. The more dimensions, the more likely it is that differentiation is horizontal: everyone might agree that the leading actor's performance was better in film B than in A, but some people like A better for other reasons.

This is the *infinite variety* property. We use it to invoke either the universe of possibilities from which the artist chooses, or the array of actual creative products from which consumers or intermediaries choose. The paintings that could be painted are infinite. The canvases that have been painted are very numerous, and many pairs of them can be found that most people would regard as quite similar in style and about equal in merit. That is, they are differentiated, but not much differentiated. The *infinite variety* property has many implications for organizing creative activities. For instance, when many examples of a creative product are available for consumption, most consumers may be nearly indifferent, yet somebody must choose which Scarlatti sonata to play, or which low-budget slasher film to show. The choice may no longer matter to Domenico Scarlatti, but it matters a lot to holders of copyrights in slasher films.

While creative possibilities are always abundant, creative realizations are

sometimes not. Songs are cheap to compose, but productions of Wagner's *Ring* cycle are very costly (relative to opera lovers' bankrolls) and hence few. When the fixed (and sunk) costs soar to a level approaching consumers' combined willingness to pay, a different organizational problem arises.⁹ Possibly no ticket price that the impresario can charge will cover the *Ring*'s costs, even though music lovers' total willingness to pay will suffice (because what some pay for tickets is less than their valuation of the experience). This problem accounts for the deficits bemoaned in the nation's opera houses and concert halls. It explains the nonprofit donor-supported organizations that arise to manage and finance these activities (addressed in Part IV).

Vertically Differentiated Skills

Cultural products differ unpredictably in the quality levels that consumers see in them. The artists who supply individual creative inputs also differ in skill, originality, and/or proficiency, though less unpredictably. These talent differences are observed when the skilled agent performs in turning out the creative product, when the finished product goes on display, or both. One artist's skill may be apparent to peers trained to supply the same creative input, to persons specialized in the coordination of this and other creative inputs, and perhaps to others qualified as teachers, critics, and the like. Artists may raise their skills by training and/or practice, but nonetheless trained and mature creative agents settle on different plateaus of proficiency. That artists (however proficient) can have good days and bad days does not undercut this ranking process, although it does wrap some uncertainty around it. For example, Hollywood's screenwriters, directors, and producers will largely agree at any one time on who are the "A list" and the "B list" screenwriters. In the terminology of economics, these creative inputs are themselves vertically differentiated. This is the *A list/B list* property.

Quality-differentiated inputs raise a number of questions about the organization of creative activities. Why do the B-list inputs get used at all? Can a film producer successfully economize with a B-list composer after paying a star's ransom for an A-list screenwriter? To what extent do consumers make their own A lists and B lists, and are these the same as the artists' and producers'? One may also wonder why the ranking of talents absorbs so much of the small change of conversation in many creative circles. Must we invoke a human propensity to establish pecking orders, or is there a clear economic explanation? A potential B-list artist with a good pedigree has something to lose from refining the ranking. Howard Becker argued that the ranking of artists and their work is intrinsic to the problem-solving process that underlies all forms of creativity.¹⁰

One reason why rank matters is the money at stake. The relevant economic concept is the differential rent—the extra total amount that people will pay to see a movie with an A-list star over the same film with a B-list star.¹¹ That differential rent limits the maximum pay that the A-list star can demand. To ask less is to leave money on the table. The rent concept also explains why the B-list artist might find it difficult to sell her services at any price. No matter how cheaply she works, the resulting film's revenue might not cover its other costs. Or, given *infinite variety*, she may face long waits between films in which her services can be a cost-effective substitute for an A-list artist.

Time Is of the Essence

The performing arts and creative activities involving complex teams—the *motley crew* property—obviously require close temporal coordination of their activities. The concert must be announced and prepared to take place at a particular time. The cinema film is efficiently shot in a certain sequence over a few weeks, during which all creative inputs must be available when needed. Although they might expect actors' performances to suffer, directors usually shoot scenes of a film out of sequence in order to group together those involving particular sets, locations, or actors. This problem of coordination in time flags the fact that the selection of a creative input depends not just on its qualities but also on when it is available. A less competent performer who is free at the right time might get the nod. Along with the *motley crew* property, temporal coordination implies a hold-up problem: an indispensable input demanding better terms on the threat of withholding its services at the last moment.

The problem of temporal coordination interacts with the conventional effect of time's passage on the value of a project. An investment that costs \$100 and will bring revenue of \$200 tomorrow is of course more profitable than one that will bring \$200 in ten years. When a motion picture with a \$100 million budget is first being organized, its expected profitability need not suffer much if the start is postponed. No substantial costs have yet been incurred, and besides, the movie might flop. But with the project under way and \$95 million already spent, any delay that postpones the influx of revenue is very costly. That reliance of the economic profitability of creative activities on close temporal coordination of production and the prompt realization of revenues I call the *time flies* property.

Durable Products and Durable Rents

Many creative products are durable: not the symphony performance itself, which dies with the last reverberation in the hall, but the symphony's pub-

lished score, and its performance recorded by a particular orchestra and conductor. Long after its composition, orchestra managers may be willing to pay royalties to perform the score, or music lovers to pay a royalty-inclusive price for the classic recording. The legal duration of the copyright determines how long the original creator or performers can collect these royalties, which are rents to the creator.¹² This durability is the *ars longa* property.

That some of these streams of rent arrive as numerous small “lumps” is important for organizing creative activities. When a recorded song is played by a radio station, on a jukebox, or by a background-music service in a public space, not many people hear it, and each may value the experience at most in pennies. Yet those pennies may add up to “real money”—that is, an amount worth pursuing if collection costs permit.⁷ The efficient collection of numerous small-value rents poses an organizational problem for owners of copyrights. It also creates a problem of contracting with the creative inputs when the durable creative product is first produced: shares in expected future rents must be defined and bargained over, and some agent assigned the responsibility for collecting and distributing them.

The other group of issues raised by the *ars longa* property concerns the warehousing and retrieval of creative durables. Some repose in public or non-profit organizations—museums and libraries. Others, such as cinema films or master tapes of music recordings, remain in the hands of their creators or those firms’ commercial successors. The Metropolitan Museum holds a vast inventory of Old Master paintings; the Wildenstein Gallery’s inventory is much smaller though still impressive. How tasks are divided between them is an economic issue.

Notice that the preceding seven properties make no distinction between what Herbert J. Gans called “popular culture” and “high culture.”¹³ That distinction is treasured in the realm of high culture, exemplified by Clement Greenberg’s famous essay on “Avant-Garde and Kitsch.”¹⁴ Serious producers of high culture engage in a cutting-edge process of problem solving and experimentation in which form transcends and dissolves content, and the consumer’s direct, subjective response matters for little or naught relative to her appreciation of the creative product’s formal relation to the aesthetic frontiers of yesteryear. Kitsch is the easy-to-like mature products descended from previous creative innovations, enjoyed by consumers who have made no intellectual investment to engage with the frontiers of creative advance. What stands unclear in Greenberg’s distinction is the mapping of avant-garde and kitsch to economic and social realms of high culture and popular culture. Is all popular culture kitsch? Or does each line of cultural products—novels, popular songs, cinema films—have its own ecology of avant-garde and kitsch products and purveyors? Gans, a sociologist, argued for the latter position. The social processes and organizational structures surrounding high and low

forms of culture do not fundamentally differ. The sophistication and self-awareness of avant-garde artists, their number relative to producers of routine creative goods, and the distribution of consumers' interests between experimental products and the equivalent of "easy listening" may vary from one creative realm to the next, but all of these elements appear in each realm. If this complex ecology does prevail in each creative activity, it should certainly imprint itself on the organization of firms and contracts, so it is a working hypothesis for this study.

Plan of the Book

This book addresses the play of these properties of creative industries as factors explaining the structures of organizations, deals, and contracts found therein. The empirical focus is on the here and now, but some telling evidence arrives from distant shores and long-ago times. The focus jumps freely from one creative activity to another, exploiting the fact that we all have a working sense of how movies are made and exhibited, books published, concerts organized, and the like. Each creative sector makes its appearance where its special features stand out most sharply for displaying some pattern, or to provide contrast and comparison.

Part I uses simple creative goods (one artist's product) to explore the economic situations of the artist striving for recognition and the gatekeeper who might select her work for development and promotion. Their relationship exposes the many difficulties of contracting on creative activities. Part II switches to complex creative goods—motion pictures, plays—that require teams of artists with diverse talents. They are assembled through contracts that try to harmonize their interests in the face of the constraint that *nobody knows*. Part III turns to consumers of creative goods and the certifiers on whom they rely for advice. Consumption decisions about creative goods share some distinctive properties. The consumer's benefit increases with accumulated experience, and consumption choices are made in a social context. Part IV addresses markets in which *infinite variety* fails, and the supplier of a single creative good is hard-pressed to cover costs. Nonprofit organizations play a major role in resolving the problem. Part V focuses on *ars longa*, the survival of creative goods or their templates to pursue rents over time and through space.

The Theory of Contracts

Most of the economic tools and terms used in this book are quite simple and presume previous exposure to the subject only at the level of the introduc-

tory course. Some mildly technical points are buried in notes so the squeamish can pass over them without courting confusion. One group of issues important for this book, however, lies in a relatively new subfield of economics known as contract theory. Much of this book is about why contracts and deals are structured the way they are, and so simple contract theory plays a considerable role. Since this subject has not yet percolated down to the introductory texts, this section provides a rudimentary account of the field's key insights.¹⁵

A contract is the agreement that governs the conduct of any economic transaction, defining what each party brings into (you supply the lettuce, I bring 89 cents) and takes from (I get the lettuce, you keep the money) the transaction. The problem gets interesting as the exchange grows more complex, perhaps with numerous parties taking actions and receiving rewards at different times, and with the actions and rewards depending on developments that occur after the contract is signed. The decisions and actions controlled by a contract proliferate rapidly as the number of parties to it increases; the problem is simpler for artist and art dealer than it is for the producer assembling a motion picture. Contracts also get more complex when the actions they govern occur over a long period of time, during which more new developments can intrude; the motion picture will be finished in a year, while the artist and dealer might hope to collaborate for the artist's whole career.

Parties enter into a contract, each hoping to get maximum benefit from it. More benefit to one seems to mean less benefit for the other, and indeed the division of benefits between contracting parties ultimately depends on their bargaining power and cannot otherwise be explained. Sometimes bargaining power is a trivial factor, as in the lettuce purchase: many grocers sell lettuce, and many households buy it, so the going price of 89 cents is the only one agreeable to both parties. Contract terms are not so tightly determined, however, when at least some parties lack close alternatives and enjoy an expected net benefit to divide among them. A fundamental principle, then, is that they should mutually agree to arrange their individual tasks and contributions so as to generate maximum value from the project as a whole, then bargain over the division of the net benefit. If defining a participant's contribution and setting its benefit share get confounded, the project will generally yield less than maximum value to all participants.

When the venture governed by a contract has an intrinsically uncertain outcome, or an outcome depending on uncertain future events, the parties face a formidable problem in writing a *complete contract* that pins down every action that each will take and the reward received. These things must be specified for every possible "state of nature"—that is, every set of circum-

stances that might emerge over the venture's duration. We can imagine and often prepare for many contingencies that will affect a contract's outcome. Still, nothing resembling a complete contract ever emerges in practice. Nature is always springing surprises that we failed to anticipate. Even "states" that we can imagine often defy formal description in a contract. Moreover, we have limited patience and powers of concentration for reaching a complete and mutually satisfactory contract—*bounded rationality*, in the language of contract theory. Once the ink is dry, problems of incompleteness likely emerge. The parties can fall into haggling over their response to some unanticipated state of nature, each hoping to slant the response to its own benefit. Ensuring that one's partners honor their commitments entails monitoring costs. All contracts that emerge in practice hence are highly incomplete. Still, they often work quite well for us, thanks to devices that attain much of a complete contract's objective without all of its apparatus.

A good alternative to the elusive complete contract is the *incentive contract*. Rather than specify a party's exact contribution (say, hours worked at some task), the contract can link the party's reward to the value it adds. This is not easy, for two reasons. First, each party's cost of taking part in a project is the effort or other input that is supplied. This cost often is not directly observed by the project's manager or other participants, and so cannot serve as the base for the party's reward. That payoff must turn on the project's overall value, or the output that results. The observed total payoff depends on each party's efforts, but also on other unforeseen and unobservable random factors. The party paid on the basis of the project's outcome has an incentive to work hard, but also gets stuck with the risk that its earnest effort appears, because of some random misfortune, to create little value. To absorb this hazard willingly, a risk-averse participant must be offered higher expected pay at the outset. This premium for risk-bearing provides a reason for weakening the incentive compensation. The weaker the incentive, however, the less the party's input of effort. Parties to a contract seek out some best compromise between strengthening incentives and keeping risk premia down, but the problem still makes the incentive contract underperform a hypothetical complete contract. Contracts in the creative industries encounter this problem, although it is mitigated by the fact that creative participants often behave as if they were risk loving rather than risk averse.

In creative industries, complex projects require the collaboration of several parties, each providing different but complementary inputs or resources. Such joint ventures have their own problems of effective incentive contracting—the *motley crew* problem. The project's net benefit or surplus is the fruit of these complementary inputs. Each might be necessary for the project to create any value, which makes it impossible to identify and reward the contri-

bution of each one individually. Subtract one input, and there is no project: each participant “makes” the whole project. In the optimal incentive contract, each party’s reward needs to be tied to what it contributes to the combined value created by the project, to induce it to exert the best effort for the project overall. When shares of the project’s benefit are divided among the participants, however, the incentive for each is generally too weak: if I get the whole benefit from the project, I exert effort to the point where the last dollar’s worth of effort I contribute brings in just a dollar of gross revenue; if I get only a 50 percent share, I quit short of that effort level.¹⁶ Most contracts that we find in the creative industries have strong though incomplete incentive provisions. The contracts are commonly simple, and they evade a complete contract’s need to describe each input or action *ex ante* and monitor it *ex post*. They accommodate pervasive uncertainty about the prospective buyers’ valuation of the project’s output, which mocks the idea of enumerating future states of nature.

Simple incentive contracts also benefit from heavy reliance on commonplace but effective enforcement mechanisms. Enforcement indeed is conceptually difficult in incomplete contracts. Not only do parties tend to slack on their contributions in the absence of monitoring or strong incentives; they also may try to hold up their partners in the midst of the project, refusing to supply their input (necessary for any valuable output to emerge) unless dealt a larger share of the benefit pie. People writing contracts anticipate hold-ups, of course, but without a good enforcement mechanism they can only avoid them at a cost. For example, they might forgo project technologies involving sunk and specific inputs—exactly the venture possibilities that make complex contracts attractive. What saves the day is the force of *reputation*. Suppose that we belong to a community of creative workers whose members regularly enter into contract-governed projects that are expected to be profitable but are subject to potential hold-ups. Suppose that some of us (not parties to a given contract) can observe when party Z attempts a hold-up, that we can costlessly pass on this news to all our fellows, and that nobody ever forgets Z’s transgression.¹⁷ Z will be excluded from all future deals. Any potential Z, knowing that this fate awaits, will face future losses that will more than offset the gains reaped by one successful hold-up.

Many complex creative deals occur in communities that are very efficient at maintaining and adjusting reputations for contract fealty. This efficiency in turn reduces the costs associated with writing and enforcing contracts, because it makes feasible highly incomplete contracts that are renegotiated on a “good faith” basis when unanticipated contingencies turn up. The theory of contracts employs the notion of an *implicit contract* that involves no written terms at all, only an informal understanding that the project will be governed

by practices that are common knowledge in the community. Related to it is the concept of a *relational contract*, in which the enforcement mechanism of reputation does not even depend (explicitly) on community-wide reputations. Suppose that Z and Y have complementary endowments of skills or assets that can support a series of expected profitable projects. No other parties fit together with Z or Y so well as the two do with each other, or a comparable fit can be achieved only after the parties incur *switching costs*. Those switching costs create room for a hold-up in any particular Z–Y deal, because any party will pay the switching cost in order to avoid it. One party’s hold-up, however, will likely induce the partner to incur those switching costs, rather than risk still more hold-ups, which would deprive the transgressor of a future series of expected profitable cooperations. The transgressor’s reputation for good behavior with its natural partner may suffice to sustain cooperation in a series of profitable projects, even if the surrounding economic community cannot observe or costlessly verify the behavior of Z and Y under their deals. Such relational contracts are said to create value because they let the partners commit (that is, sink irrevocably) resources into *transaction-specific assets* that hold their full value only when used in Z–Y projects.

Contract theory emphasizes problems that arise from asymmetrical positions of the parties, most notably *asymmetrical information*, leading to the “lemons” problem. If one party (the current owner of a used car) has better information about its quality than does the prospective buyer, the buyer naturally suspects that the seller will overclaim that quality or fail to mention known defects. Two sorts of contract failure can ensue. The unwary buyer who has undervalued the seller’s exploitation of information asymmetry can wind up owning a lemon, or the suspicious buyer can forgo the purchase of what is in fact a good used car. Remedies to the lemons problem exist in guarantees and independent appraisers (as well as reputations, if the seller deals regularly), but the problem seldom goes away. In fact asymmetrical information plays little role in creative industries’ contracts, exactly because of the property that *nobody knows* the value of a project until most or all resources have already been committed to it. Rather than asymmetrical information, there is symmetrical ignorance.

Another problem that does pervasively affect creative industries’ contracts, however, is that of efficiently allocating *decision rights* under an ongoing contract. When a contract is left incomplete, the parties must agree on how to decide their common course of action, after some unanticipated development occurs or an uncertain outcome is resolved. Given whatever deal they initially struck to share the project’s benefits, they will commonly have conflicting preferences about how to revise or fill out the terms of their deal. At this point, it may matter greatly who is better positioned to observe, diagnose,

and act on the resolved outcome. Maximizing the value of the venture calls for assigning the right to decide on the responsive adjustment of the project to the party better suited to act (Z). Giving Z the decision right at some interim stage, however, poses for Y the problem that Z's decision will be bent toward maximizing value for Z and not necessarily for Y and Z together. Y can preserve an acceptable contract, however, by demanding compensation up front for Z's expected opportunistic behavior in making such interim decisions.

This problem of allocating decision rights turns out to be pervasively important in the creative industries, because complex creative projects regularly involve sequences of actions, each one a sunk and irrevocable cost for the party incurring it. The writer prepares the movie's screenplay; the producer mobilizes the actors and other inputs to make the film; the editor edits it and the composer prepares the soundtrack music; the studio plans the film's advertising and distribution. At each stage one party makes a decision about sinking more resources into the project, while the previous contributors' investments are wholly sunk. Decision rights efficiently lie with the party about to contribute still-fungible resources. That is because the project should go ahead if and only if its currently expected value exceeds the still-avoidable cost of these fungible resources (even if not the costs previously incurred). This allocation of decision rights is achieved by an *option contract*. Party Y, in a position to decide the next move, negotiates a contract with predecessor X that specifies how X will be rewarded if Y decides to go ahead with the project. Y also purchases from X the opportunity for a specific time interval in which to investigate the current omens and resolve whether and how to proceed with the project. Once Y exercises its option to purchase and sinks its own input into the project, decision rights are handed on to Z, who takes an option to consider the next step.

This book relies on contract theory, but also on the field of industrial organization, which seeks to explain the sizes, activity patterns, and numbers of firms that we find competing in particular markets. For this task contract theory and industrial organization come together. From the viewpoint of contract theory the firm can be regarded as a *nexus of contracts*—between the owning shareholders and the chief executive, the chief executive and the vice presidents, and on down to the factory manager and the workers. Contracts in force within the firm may not differ fundamentally from those between independent parties. Still, there is a difference between the firm's hiring workers on a presumed long-term basis and the workers coming together for circumscribed projects. Long-term employees get their marching orders from a supervisor each day, under a flexible continuing contractual arrangement and usually without any fresh or recurring contract costs. Every project outside

the firm involves a separate contract, although it may benefit from relational practice. Which mode of organization works best varies from situation to situation; the motion picture industry demonstrates this dramatically in the sharp transition that occurred in the 1950s from one mode to the other (see Chapter 5).

Drawing on contract theory's contribution, the field of industrial organization brings its own techniques to explaining why the production of creative goods is organized as it is. Why are some markets served by a single seller (for example, a city's major symphony orchestra), others by many (commercial art galleries in New York)? The answer is that each market's technology and activity patterns dictate that participating firms incur some fixed cost, or produce subject to some economies of scale. Relative to the amount of market demand that the firm faces (and the aggressiveness with which the firm competes with its rivals), these fixed costs determine the maximum number of competitors that can "fit" into the market and still earn normal profits. The number may be shriveled further in some industries by first-mover advantages that incumbents enjoy simply from having already arrived. To explain why an industry contains only a few firms, we look for these scale and first-mover factors. The high concentration of cinema film studios, for example, turns out to be due to the fixed cost of maintaining an apparatus for distributing and promoting films. Movie production takes place as one-shot deals on a much smaller scale.

The number of firms occupying a market is important for industrial organization, because it strongly influences how competitive and efficient they are, but it is not the only important question about how the marketplace is organized. Firms in a market, like the products they offer, are rarely homogeneous. Directly competing firms fall into groups differing in the product varieties that they turn out or the way they organize those goods' production. In the creative industries these diverse groups of market competitors turn out to reflect the basic properties of creative goods. For example, the inner and individualistic aspect of creative production mixes badly with the orderly, rule-driven routines of the large, bureaucratic firm. Creative industries tend to be organized so that the most delicate dealings with artists concentrate in small firms specialized to this task, while other firms (usually much larger ones) undertake to combine creative work with large teams of humdrum inputs.

Another question about industrial organization often carries the tag "firm versus market": explaining why some transactions take place at arm's length between independent parties while others get "internalized" within the firm. The question most commonly takes two specific forms. Why are firms sometimes vertically integrated, supplying their inputs internally and/or distributing their product through their own sales organizations, while elsewhere

these transactions take place at arm's length? Why does the firm operate in several different product markets, perhaps using some inputs or facilities in common, when those markets could each be served by independent, single-market firms? The answers to these questions depend heavily on the theory of contracts, already outlined. Where a firm and its major supplier can build up a relational long-term contract that permits efficient transaction-specific investments, they may get along perfectly well as independent organizations. Where such contractual dealings are soured by opportunism, haggling, and monitoring costs, vertical integration comes to the rescue. These choices of organization are discrete. Contracting problems certainly exist within the integrated firm as well, but it has a trump card in the flexible, long-term employment contract, which lets it retain an experienced team of inputs and deploy them flexibly in light of the day's fresh news. Hence, especially where different organizational arrangements are about equally effective, we may find quite different organizational forms coexisting, or observe dramatic shifts from one to another.