Creative Industries

Contracts between
Art and Commerce

Richard E. Caves

Harvard University Press
Cambridge, Massachusetts, and London, England

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 preexisting 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 published 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. 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 nonprofit 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 introductory 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 contribution 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.

3

Artist and Gatekeeper: Trade Books, Popular Records, and Classical Music

The filtering of artists by gatekeepers and the promotion and distribution of their works goes on in all the creative sectors, leading to trouble-prone relationships between artists and the firms that link them with humdrum inputs. Authors and trade-book publishers offer a useful comparison to the visual artist and gallery. Whereas the gallery seeks to promote an unpredictable career, the publisher addresses a problem that is usually simpler. The manuscript of the book is in hand; even if not, as with celebrity authors and blockbuster novelists, the prospective development of the artist's career is not such a central concern. While that smooths the contracting process, it raises the stakes for the gatekeeper, who has more information to work with and hence a more complex task of processing it. The record album is the unit of transaction between the pop performing group and the record label, just as is the gallery show for the artist and dealer. Each team, however, hopes to launch and sustain a career for the artist. The deal again runs into problems already seen in the visual arts: hidden actions and information, incomplete contracts, the uncommitable nature of artistic inspiration, and moral hazard. The tweedy author and the grunge musician face very similar issues of contracting with gatekeepers. It proves convenient to explore some issues in the author's context, others in the musician's.

These problems of gatekeeping and contracting open several distinctive issues of industrial organization. Agents commonly serve as business representatives for artists, but their role in allocating creative resources goes much farther. They perform their own gatekeeping function in deciding which artists to represent, and along with this undertake the matching of artists to jobs or humdrum partners. Some of these tasks have natural-monopoly elements, others need one-on-one relationships to artists, so the organization patterns of agents vary greatly among art realms. Another issue of industrial organization, demonstrated in trade-book publishing, is the informal but extensive in-

formation networks that link editors in different publishing houses. The creative-industry firm is no fortress of proprietary information.

Author, Agent, and Publisher

Like classical musicians and cinema and TV film artists, trade-book authors employ agents as intermediaries, to match author with publisher and bargain the best terms. The agent deals proximately with an editor, who becomes the manuscript's champion or manager within the publishing firm. Because all functions that a publisher performs—copy-editing, printing, distribution, and so forth—can be subcontracted, the tasks performed by the publishing-house editor and the small-scale entrepreneurial publisher differ only in governance relationships among the parties.

Agent as Intermediary

The creative act of authorship has its sociable side, but compared to visual art and pop musicianship it is a solitary and individual calling. The typical beginning author cannot follow the visual artist and build a reputation outside the realm of commercial gatekeepers. There does exist a marketplace for literary reputations in the high-minded circles of the university campus and the little magazine, but it has separated itself from the mainstream of commercial writing in books and magazines.[1] There is, then, no marketplace for early reputations of emerging authors ("celebrity" books aside). The author simply prepares a manuscript and presses it into the hands of a gatekeeper, either the publishing-house editor or an agent. The chances that the gate will open are small. The president of Doubleday once stated that three or four of the 10,000 submissions received "over the transom" each year were accepted. Novels face even worse odds (one in 15,000), while scholarly and textbook publishers net many more fish from the stream of submissions.[2] These odds understate the author's chances, since they pertain to any single publisher's acceptance rate, whereas many manuscripts find homes after multiple submissions.

The publisher might decline a particular manuscript for many reasons: it fundamentally lacks quality and originality; it shows promise that could be realized only with substantial further work by the author; it is too specialized to be profitable; it fails to fit well into the publisher's portfolio of books; or it arrives at the wrong time, when the publisher's pipeline is full. Because any novice's manuscript has a tiny chance of acceptance, the publisher either incurs a high cost (per acceptance) of filtering out the losers, or chances many errors of omission (rejecting manuscripts that are worth publishing). There

is room for an intermediary who can efficiently match up publisher and author. The agent does that job. Publishers once shunned agents as collectors of rents that an inexperienced author might leave on the bargaining table. Agents are now welcomed for reducing the publisher's recruitment costs. Just when and how this attitude shift occurred is not clear. But a logically sufficient explanation lies in the growth of the literary marketplace and the total supply of manuscripts available for screening by any one publisher.

The agent's skill as intermediary lies in sensing the threshold quality a manuscript needs in order to interest a publisher and knowing the sorts of manuscripts sought by various publishers. The agent who selects from the "slush pile" of unsolicited submissions a would-be author to represent thus serves as a first-stage gatekeeper. The agent may press the author to change or improve the manuscript enough that an editor can see its potential. The agent then submits the manuscript to appropriate publishers, not so much to foment an auction as to identify the most fruitful author-publisher match. To this task the agent brings both superior knowledge and a reputation asset grown from repeated interactions with publishers. The experienced agent seeks a reputation for truthful assessments—not overclaiming either the quality of submitted manuscripts or the collaborative capability of the author. The quest for this reputation is propelled by the agent's own compensation structure. The agent's conventional 10 percent of the author's royalty does not provide a compensatory total income (that is, one matching other job opportunities) unless he "hits" with a reasonably large proportion of submitted projects.[3] Agents nurture relationships with editors rather than with the publishing houses that employ them, and the editor's claim on the publisher's resources may be more important than the publisher's identity. Editors, who frequently change jobs between publishing houses, function as in-house entrepreneurs with concentrated incentives to produce successful books. They hence value accurate assessments and astute matchmaking by agents, including theory-of-the-case ammunition for getting promising projects accepted by the publishing house. Editors with projects in mind may turn to agents representing authors who have done similar books. Indeed, switches in occupation between agents and editors appear common.[4]

The importance of agents' matchmaking function is affirmed by evidence on the origins of literary agents in the nineteenth-century Britain. At the outset the services were informal and unpaid; persons with *art for art's sake* tastes often performed a good-offices function solely for the pleasure of seeing a fruitful deal completed. The growth of the literary marketplace (provincial newspapers, for example) encouraged the emergence of author-paid agents such as A. P. Watt, who served as not only a bargaining agent but also a counselor and legal advisor to authors. Watt began by charging fees for services, but authorial poverty and the power of well-tuned incentives soon led him to contract on a revenue share of 10 percent, a figure that (with some upward perturbations) persists to this day.[5] The origins of the agent's role underline a reason why the agent is efficient for performing the gatekeeping function. As the author's employee, the agent can elicit information from the author on components of her "reservation price"—willingness to cooperate, make manuscript changes, and the like—that the author bargaining directly with the publisher might strategically conceal. The presumed nonnegotiability of the conventionalized agent's fee keeps this strategic consideration from inhibiting the transfer of information from author to agent. The elicited information might tell the agent that no deal is likely, but the publisher is spared wasted haggling costs.

The author-agent relationship does tend to destruct when the author succeeds and becomes well-known to publishers. The agent's talents for shaping manuscripts and interceding with publishers become less valuable. The 10 to 15 percent commission implies a high charge for negotiating and monitoring contracts, services that a lawyer could perform for a fee.[6] Agents, like others who help develop an artist's career, find it difficult to write contracts that award them the expected value of their contribution to the artist's lifetime income without creating a strong incentive for the artist once successful to demand renegotiation. The committed duration of artist-agent contracts is commonly short and often open-ended. The downside hazards of long-run commitments must be too great: for the agent, being stuck promoting a burned-out author; for the author who achieves major success, overpaying for the agent's current services.

A good test of agents' function comes in the traits of publishing segments where they are not used, or just falling from or coming into use. The last is illustrated by Christian publishing, which has recently seen much turnover of firms, cutbacks of lists, movement of authors from house to house, and increased prizes for the most successful authors.[7] Each factor favors or reflects agents' services. The turnover of firms and cutback of lists both increase the magnitude of the publisher's filtering problem and make the agent's gatekeeping more valuable. Meanwhile, enlarged top prizes raise the value to the author of the agent's bargaining skills, and mobility of authors likely reflects the agent's matching function. The agent plays little or no role in book publishing outside of trade books, because his advantages in lubricating the process of selecting authors do not carry over. The scholarly author holds professional qualifications and a reputation that have no counterpart for writers of sensitive first novels. The gatekeeping publisher has easier access to information beyond the manuscript itself. Textbook publishing, for example, with high fixed costs for both author and publisher, requires at least informal

precommitment between author and publisher at the outset, and peer review supplies a cost-effective evaluation. For scholarly books the publisher is apt to use other gatekeepers than the agent, such as the well-known academic who, as series editor, takes on the gatekeeping function for a 2 to 3 percent royalty. The academic author desires the quality signal broadcast by a distinguished series imprint, and the series editor's reputation certifies this. The series editor's own reputation suffers if weak or partisan manuscripts get published. Most scholarly authors rationally expect little pecuniary gain, so that the agent's bargaining prowess has little payout. The scholarly publisher runs little risk from accepting manuscripts submitted through a series editor, and need not commission costly evaluations. The selection of the series editor, however, is a serious investment for the publisher.[8]

Author and Publisher

Publisher and author of a finished manuscript face an easy contracting task. The manuscript removes most issues of opportunism and hidden actions on the author's part and makes the publisher's actions easy to specify. For established authors or celebrity books, however, publishers may participate in an auction based on just a twelve-page prospectus; if the remaining 500 pages on arrival lack the lapidary quality of the prospectus, the publisher generally can back out.[9] That is, the deal takes the form of an option contract obligating the author to deliver the manuscript to the publisher, but the publisher need accept it only if it is "satisfactory." Many disputes arise over decisions to decline manuscripts. Insufficient quality or a missed deadline are contractually valid grounds, though they can be abused as makeweights for opportunistic factors. Contracts frequently entitle the publisher to retrieve an advance for a declined manuscript. Exercise of the privilege leads to bitter disputes and often proves infeasible, leading to such compromises as partial recoupment from the author's royalty income when the manuscript is placed elsewhere.[10] The general implications of option contracts are discussed in the next section.

The royalty contract gives the publisher only an attenuated incentive to promote the book. For a trade book correctly expected to enjoy moderate sales, the typical royalty contract, which offers the author 10 percent of the book's retail price, splits the gross profit (that is, the difference between the publisher's marginal cost and wholesale unit revenue, before deducting the publisher's fixed costs) about 58–42 between publisher and author.[11] As was shown in Chapter 2's discussion of profit- and revenue-sharing contracts, the author's royalty based on sales leaves the publisher with an underinducement to promote the book (the author, correspondingly, wants more promotion

than would maximize author's and publisher's joint profits). The author's advance, though, strengthens greatly the publisher's marginal incentive to promote—he gets roughly the full resulting gross profit—until the advance is earned back, which increases the efficiency of the contract.[12] Another conflict arises between publisher's and author's interests in keeping a book in print once its sales fall to a low level. Keeping a book in print and in stock imposes a fixed cost on the publisher, while the author gets royalty income without sharing that cost.

The publisher's standard contract enjoys widespread use, with little variation in its terms. Why does it so dominate the alternatives, especially since many rejected authors would happily sign if they could? Why does the author receive nearly half of the gross profits? In fact some publishers require only that the author cough up the cost of publishing, distributing, and promoting the book. Completed book manuscripts in a sense go unpublished because the author is unwilling (unable) to throw good cash after time that was badly invested (aside from the pleasure of literary toil). The "vanity press" sector does not appear to do a large business.

Suppose that the standard contract were less generous to the author—say, 5 percent rather than 10 percent royalties. Publishers, gaining a larger share of the profits, would choose to publish more books, and the threshold of commercial prospects would fall. Weaker or more specialized manuscripts would get published, but also some potential authors capable of writing good books would divert their energies to more rewarding activities. Here we find for the first time a pattern that appears throughout the creative industries. Creative inputs are vertically differentiated *(A list/B list)*, but also subject to pervasive horizontal differentiation *(infinite variety)* and uncertain reception in the market *(nobody knows)*. There is a collective interest in keeping enough artists in the game to ensure ample candidates for random success in book publishing (or, in motion-picture production, availability of just the right bit player for a role). That interest is advanced when the author allowed through the gate receives a large prize. We call this the *lottery prize* phenomenon: many keep buying tickets despite the rational expectation that their chances of winning are tiny. Readers benefit if more good books get written, and publishers who gain utility from publishing putatively good books also will not regret liberally compensating the author who is allowed through the gate.

The lottery-prize arrangement may serve the common good, but why does each party individually stick with it? Why does no trade publisher prosper by fishing from the "slush pile" the few promising manuscripts that have slipped past the 10 percent publishers, and offering these rejected authors smaller prizes? Authors have no minimum-wage contract, unlike many performing artists (Chapter 7). The ample supply of parties willing to become 10 percent

publishers by itself tends to make this "deviation" unprofitable.[13] Perhaps an informal standard evolves, widely shared among readers, publishers, and successful authors, as to how good is a "good book," and how deserving of reward is an author capable of passing through a good publisher's gate. This standard discourages bottom-fishing in the manuscript pile, even if it proves (marginally) profitable. That is, the 10 percent royalty publisher who cuts to 5 percent certainly loses the good authors. He also loses some marginal authors who would rather continue searching for a 10 percent publisher rather than "slum" for 5 percent. He loses respect among publishers and potential employees. These factors may deter trade publishers from cutting the royalty below a conventionalized rate, although the argument does not explain exactly why that convention took root in the first place—beyond the focal power of round numbers.[14]

Because most authors complete books only infrequently, career-building investments dwindle in importance. Nonetheless, a successful first book raises readers' interest in the next, just as a successful second book increases demand for the first. The publisher can internalize this benefit only with contractual first refusal of the author's next manuscript. The author has an interest in committing subsequent books to the publisher in order to lift his incentive to promote the first. Such a lock-in has pitfalls for the author, however. On what terms will the next book be published? A right of first refusal leaves the author free to accept any better offer for her next book. When the next manuscript comes around, the original publisher's bid anticipates some of the second book's expected spillover profits from promoting backlist sales of the first, in order to meet the highest outside bid. A firm commitment of the author's subsequent works denies her this leverage. In the first contract she can try to bargain favorable terms for subsequent books, but at that time *nobody knows,* and her bargaining power hence is weak. With the contract terms on future books left incomplete, the locked-in author assumes that the publisher will exploit any opportunity that arises. She either demands upfront compensation for being locked in or underinvests in effort on the next manuscript. In practice, lock-ins are not seen, and first refusals prevail but apparently cause few problems. One survey of authors found that not only did they typically take the first contract offered to them, but that over 60 percent repeated with the same publisher by default.[15]

The history of contracting practice holds interest, because the changes reflect changes in the technology and organization of publishing, changes that have not always improved the incentive structures of contracts. In England in the nineteenth century and before, either the author received a fixed payment from the publisher, or the two parties shared the net profits from the publishing venture. Up to the 1830s and 1840s, the deal covered only a

single printing of the work. Printing type was too costly to be set and retained for possible reprintings, so that any reprint was technically a fresh start for the publisher. The development of stereotype and later electrotype plates eased reprinting and allowed an open-ended deal dependent on the book's sales. Open-endedness increased the uncertainty of the venture, however, encouraging the publisher to share the risk with the author (especially the novice).

The royalty contract with an advance improves efficiency until the advance is recouped, but then generates discord as the author benefits more from the publisher's actions that enhance gross revenue, while the publisher with a revenue-sharing contract underspends. In any case, Victorian novelist George Eliot was the first major writer to receive a royalty contract (in 1860 for *The Mill on the Floss*), and by 1900 it was standard in both Britain and America. At first the royalty kicked in only after the publisher recovered printing expenses, but that feature disappeared.[16] The royalty contract was probably associated with the rise of the large publishing house and the practice of successive differentiated editions of successful novels. A profit-sharing contract needs the account-keeping process (whoever performs it) to be transparent to both parties and free from "judgment calls" that allow opportunism. As publishing firms became larger, with overheads spread among many projects, and as the publisher's costs became a stream incurred over time, accounting opportunism arose of the sort that Hollywood later made famous (Chapter 6). The practice of temporary purchase of copyright by the publisher had strong efficiency advantages (the publisher faced the full cash-flow consequences of his decisions), with one exception. As the time of reversion of copyright to the author approached, the publisher's incentive was to run the presses nonstop and flood the market.[17]

Editor and Publisher

The gatekeeping process in book publishing affects the relationship between the editor and the publisher who employs her. The editor plays an entrepreneurial role inside the firm, starting with the decision to accept a manuscript. Editors in scholarly publishing have substantial authority to sign manuscripts on their own, as quality and breadth of interest are relatively easy to determine, and the firm's risks are correspondingly small. In trade publishing uncertainty is greater, and success is more dependent on the efforts exerted by other specialists within the firm (design, advertising, marketing). The editor assumes the role of advocate, supplying impetus and staking her own reputation on her conviction of a manuscript's commercial worth.[18] One asset that the editor deploys in this contest for resources is her status as the key link be-

[illegible]he firm and the author.[19] The editor's entrepreneurial role diverges [illegible] formerly central copyediting function, now spun off to specialists. [illegible]their entrepreneurial roles, editors' compensation apparently contains little pecuniary incentive.[20] Bonuses may be growing more common, however. Highly successful trade editors are commonly rewarded with their own publishing imprint, which at least enlarges the reputation capital that they can glean from producing successful books. As head of an imprint, the editor also has the discretion of offering larger advances and taking a more active role in the book's design, marketing, and publicity.[21] The mobility of editors among publishing houses is high and increased by the goodwill asset built up between editor and author, which allows a publisher to bid for an editor's services with the hope that her authorial ducklings will waddle along behind. The *nobody knows* outcome of editors' efforts allows reputations to swing unpredictably between dud and star status, which also increases mobility.[22]

Editors' circulation among jobs flags a pattern of behavior that appears throughout the creative industries. Regular employees who play a creative role (for example, engineers in semiconductor firms), or who hold goodwill assets with independent artists (such as trade-book editors), maintain extensive contacts with their counterparts in other firms. The contacts involve the exchange of copious information and deal possibilities in the form of favors given and returned. Economists normally assume that proprietary information is an asset used intensively within the firm and guarded vigilantly against leakage to the outside world, so this interfirm "networking" needs an economic explanation. Through these industries surge large, amorphous flows of new information. With constantly changing data, new information is costly to obtain, nobody knows it all, and employees with similar professional interests can benefit by sharing on a reciprocal basis. An employee of firm A may have some bit of information about a failed project of A's that is useful to firm B, yet its leakage imposes no cost on A. When an employee of A passes it to B's employee, the two firms taken together are better off. Furthermore, A's employee (and therefore the firm itself) acquires a call on some future reciprocation. One-third of editors reported having had projects referred to them by other publishers.[23]

This process of exchanging information complements the mobility of personnel between firms. The network not only facilitates job switching; it also widens individuals' circles of contacts and speeds the dispatch of information toward its points of greatest usefulness. The role of informal reciprocity and reputation is underlined: the better connected the network, the faster the news of defection travels, and the longer it is remembered.[24]

Dealings between Artist and Record Company

Young pop musicians (introduced in Chapter 1) seek work at local clubs, schools, and the like, honing their performing skills and (likely) writing a body of songs that they hope will eventually attract a record company ("label"). A band recruits a personal manager to help with both its creative development and business management. This role is both complex (in terms of the mixture of skills and the personal congeniality required) and risky (because of the large proportion of performing groups who break up or fail ever to earn substantial revenue). Managers accordingly demand a large share of a successful group's income, 15 to 20 percent of gross earnings (and thus a large bite of net earnings), excluding advances from labels for recording costs, support of tours, and the like. The contract's duration is clearly a sensitive issue. It generally runs from three to five years, with the artists holding the right to cancel if their earnings fall below some threshold. The manager receives a continuing share of royalties from any records made under the contract.[25]

Performers and labels pursue each other in various ways. The labels employ talent scouts ("strange noctavagants," *Fortune* once called them) to seek out talent, and in-house producers become advocates for groups that they record in the manner of publishing-house editors.[26] Musicians employ attorneys in the manner of literary agents to shop demonstration ("demo") tapes around to labels. A label may receive three to four hundred tapes a week, so only those backed by a competent certifier are likely to get attention.[27]

Terms of Recording Contracts

The principles behind the typical contract for a present-day popular music composer-performer resemble those in the visual arts. The artist (group) provides the creative inputs through songwriting and performing to create master tapes of songs suitable in quality and quantity for release as an album. The label produces and distributes the album on compact discs and cassettes and advertises the artist and album through videos, by supporting performance tours, and by promoting the album to broadcasters and record buyers. The artist collaborates not just in producing the master tape and video but also in touring and associated direct promotion.

Like the visual artist (and less like most trade-book authors), the musician and label seek to promote an extended career and gain a long-lived earnings stream. The payout is highly uncertain, however. *Nobody knows:* casual estimates suggest that roughly 80 percent of albums and 85 percent of single records released fail to cover their costs (the "stiff ratio").[28] The chances for a

successful second album are not high even given a successful debut album. The artist who had her whole life to ready the first album now has six months for the second. Apart from that, among all the world's possible record albums, randomly excellent, okay, or terrible, the musician's first released album might be a lucky draw of excellence that is unlikely to be repeated.[29]

For artist and label the issues of contracting essentially duplicate those faced by visual artists and authors. Royalty contracts based on sales revenue again dominate, with a more efficient profit-sharing deal likely ruled out by the difficulty of monitoring the label's bookkeeping. Contracts also reflect both the uncertainty of success and the inability of young performers to supply inputs other than their time and talents. The artist can commit to produce an album that the label might find acceptable to release, but of course cannot predict or guarantee the public's response. The label cannot efficiently guarantee in advance what promotional effort it will undertake: this depends on the album's character and how much (if any) public interest it initially elicits.

In the standard recording contract, the artist commits to deliver exclusively to the label master tapes for a series of albums. The label holds an option to distribute these records for a period of time, and commits to pay royalties based on revenues generated by each recording. For each album the label provides an advance to the artist that at least covers the cost of recording the album (including union-scale wages to the session musicians) and may provide income that anticipates earnings from royalties. The advance against royalties is an absorption of risk by the label and a guarantee of minimum royalty income for the artist. Royalty rates fall in the range of 11 to 13 percent for a new artist (as low as 9 percent for a small, independent label), 14 to 16 percent for artists in the middle range of success, and 16 to 20 percent for superstars.[30] The rate on a given album may escalate with the number of copies sold.

The effective value of these nominal royalty rates is reduced by what Donald Passman called a series of "cheats" by the labels that have crawled into the standard contract. The royalty is based on the album's suggested retail list price, but this is reduced by an arbitrary "packaging charge" for the record's container, 20 percent for cassettes and 25 percent for compact discs. The count of records shipped (for determining the royalty payment) is reduced by genuinely donated copies used for promotion, but also by phony donated copies (15 percent) conventionally generated by boosting the wholesale price quoted to the retailer but shipping an offsetting number of free copies. Until recently the count of records shipped was reduced 10 percent to allow for breakage, although fragile shellac records disappeared from the market 50 years ago. The breakage and free-good allowances together knock the artist's royalty base down to 76.5 percent of actual shipments.[31]

A vital feature of the recording contract—one that appears throughout the creative industries—is its option structure. An option in this context gives the gatekeeper an exclusive right (for some period of time) to consider acquiring the artist's creative product on pre-agreed terms; if the gatekeeper declines, he has no further obligation to the artist, and she keeps any payments already made to her (in some art realms that includes payment for the option itself). The label, having advanced to the artist the cost of recording an album, is not required to release it, if the tape disappoints sufficiently to kill expected profit.[32] That means gross profit (without deducting the advanced recording cost) from the record at hand plus net profit expected from future albums. If the label continues to release tapes delivered to it, the contract then obligates the group to supply further albums. Nowadays the terms commonly require the artist to deliver a second album within some number of months after the first. The label's acceptance carries an increased advance and royalty rate, but it also commits the artist to ready yet another album, the cycle continuing for a total of eight to ten records.[33]

The artist receives royalties on a record only after the advance is recouped. No necessary relation exists between profit to the label and the recoupment of the advance. Likewise, an artist might issue a series of apparently successful records and yet receive no royalties except for any excess of the advance over recording costs. These outcomes depend simply on the size of the negotiated advance relative to the maximum joint profit actually available to artist and record company together. The label insists on "cross-collateralizing" the artist's albums issued under a contract. Cross-collateralizing pools revenues from independent projects (or markets) for the purpose of recouping advances for any of them: if the first record fails to earn back its advance, the deficit becomes a claim on the royalty stream of the second record, and so forth.[34]

The option contract looks unfair to guileless young musicians, yet its form has a compelling logic. Many creative goods proceed through stages of production with costs completely sunk at each stage, yet with the final payout remaining highly uncertain. At each stage, fresh news regarding the ultimate payout likely emerges, and additional costs are incurred. The efficient contract allocates decision rights at each step to the party about to sink liquid purchasing power and able to test its expected payout against the project's current prospects. If would-be superstar musicians became much scarcer relative to the public's demand, the new equilibrium contract would likely provide the musicians with more favorable advances and/or royalty rates, but

preserve the option structure. This proposition does depend on the two parties' attitudes toward risk; if artists are risk averse, the label might give a firmer commitment to release even in the presence of fresh bad news. In fact highly successful artists bargain for somewhat stronger obligations for the company to release; for example, the master tape need only be "technically satisfactory" or "in the artist's previous style" rather than "commercially satisfactory."[35]

Implications for Incentives

This contract has a number of salient incentive features. By guaranteeing terms of access to the artist's future albums, it bolsters the record company's incentive to promote the artist's first album despite the high expected stiff ratio. A successful group is under exclusive contract long enough for the label to recoup on later albums any losses run on earlier ones. Promotion of one album generally has positive spillover effects for the artist's previous albums, and this is also internalized. If the artist is successful enough, the contract's built-in schedule of royalties may not escalate fast enough to deny substantial rents to the label. But the contract is then prone to renegotiation: injustice perceived by the performer in her current contract does not favor the timeliness and quality of the next album delivered. The label has only limited ability to collect its biggest contractual jackpots in full. The label's incentive to promote a successful artist withers as the contract's expiration date approaches, shrinking the rents expected to fall into the label's hands from subsequent albums. For the artist the multiyear option contract does not provide much downside protection. Even if the label continues to pick up its options, unsuccessful albums may cause it palpably to "lose interest" in the artist. While it might be obligated by contract to issue the artist's next album, it is not obligated to promote it. The label's ability to terminate a contractual arrangement is quite complete, and the artist's ability to sustain the forward motion of a best-efforts commitment quite limited.[36]

The royalty contract based on the label's sales revenue diverges in some ways from the terms under which authors and visual artists get paid. While the label funds the physical production and distribution of the album, large components of promotion costs (as well as all recording costs) are recoupable from the artist's royalties. Because each party to a royalty contract benefits from the other's promotional outlays, the label might even overpromote an album, relative to the outlay that would maximize parties' combined profits. The incentive's strength increases with the company's confidence that the label will recoup all costs that it has advanced. The incentive decreases with the

artist's effective royalty rate. The higher the rate, the less of a promotion-induced extra dollar of revenue lands in the company's pocket. There is no easy way to test empirically for overpromotion, beyond noting the industry lore on promotional expenditures charged to artists' royalties that have little function save gratifying the artist: vanity billboards on Sunset Boulevard and "limousines to take you to the bathroom" might attest to either an overinducement to promote or the strong appeal of ego gratification to the immature.[37]

The artist may have little control over the promotion outlays chosen by the label, but the costs of recording the album depend directly on the artist's quest for quality (and degree of organization for achieving it). The royalty contract seemingly induces the artist to underinvest in studio time and resources, because the company pockets a large share of the additional revenue once the advanced recording costs are earned back. Nonetheless, record executives commonly complain of the artist who out of indiscipline or perfectionism takes up far more studio time than necessary.[38] Such a concern by the artist with creative success over income can be written off to *art for art's sake.* The uncertainty of the album's success, however, supplies an economic reason why the performers might rationally overinvest in studio costs. If the album fails to recover its costs, the artist receives no royalties while the loss—whether large or small—falls entirely on the label. The artist's royalties increase dramatically, however, when album is a big success—indeed, the royalty *rate* commonly escalates. If inflating the studio costs raises the chances of a big win, while the downside loss falls entirely or mostly on the label, the artist could select a heavier investment in album production costs than the label would prefer.[39]

Governance of Contracts

Recording contracts encounter various problems of opportunism as they run their course. Previously mentioned was the erosion of label's incentives to invest in the artist's career as the contract's expiration date approaches. From the artist's viewpoint, a problem of moral hazard arises because the label keeps the books that determine the earnings remitted to the artist. This problem festers at some level in all of the creative industries, reaching its peak in the motion-picture industry. Recording artists face less commonly the problem that afflicts visual artists (Chapter 2) of bankruptcy and disappearance of the commercial partner, but a greater hazard of opportunistic transactions. In a long dispute between the Beatles and EMI and its U.S. subsidiary Capitol, undercounts of sales for royalty calculation were alleged, as were

transfers of "free" promotional records to subsidiaries that released them for commercial sale.[40] Music videos are a valuable tool for promoting albums, although their high fixed cost warrants their preparation for a new artist only if an album has already shown substantial public appeal. A video therefore is not committed in the contract initially and involves renegotiation. The video itself generates a royalty stream for both artist and label, apart from its promotional benefit. Usual practice is apparently for the label to recoup 50 to 75 percent of the video's cost from the artist's royalties—from the video, from the record, or from the two cross-collateralized.[41] Overall, the recording contract seems peerless in the scope that it offers for governance disputes. With the recent consolidation of several record companies these have reached fever pitch, causing many artists to consider managing their own promotion and record distribution over the internet. But one website among the endless horde, with no mediating gatekeeper, is probably an inadequate substitute for an established label.[42]

Artists' concert tours raise their own set of complications. Concerts on the road bring in substantial revenue streams while also incurring heavy direct costs as well as personal stress for the performers. They promote record sales, just as records increase the demand for concert appearances. This interdependence provides an example of an economic relationship that appears here and there throughout the creative industries. Assume that an artist's recordings and concert tours are each profitable, taking just their own respective costs and revenues into account. Because each increases demand for the other, each profit stream's recipients would willingly spend to induce an increase in the scale of the other activity. When a deal is worked out, it is not clear whether (in this case) the label will contribute to touring expenses, or the tour's beneficiaries (artist and promoter) will sacrifice to increase the stream of recordings. A mutually acceptable deal could send payments flowing in either direction. Labels have supported tours to a degree varying over time, cutting back when music videos emerged as an effective competing promotion technology. New and unknown groups can benefit greatly from opening for established groups at concerts; indeed they might willingly pay for the privilege and have been known to do so.[43] In the 1960s touring pop groups were paid very little, but their recordings received extensive local airplay. British artists found that to build a base of fans large enough for a high level of success with recordings they had to make several money-losing tours of the United States—clearly a high cost for establishing a group.[44] Gradually the concert circuit evolved into a spectacular "show" and a way to "break" (that is, promote into widespread recognition) new groups and advertise new albums of well-established ones. The revenues from concert tours have increased greatly, but so have their costs and promotional benefits, and taken by themselves they apparently are profitable only for the most popular groups.[45]

Contracts between performers and local promoters and venues suffer their own epidemic of failures. A profit-sharing arrangement is common: the performers get a guaranteed minimum payment; before any further distribution to the performers, the promoter recovers expenses and a "guaranteed" profit; and the remaining net revenue is divided between performer and promoter in a ratio between 60/40 and 90/10. Large venues (stadiums) with local spatial monopolies capture some rents (often partly absorbed in union wages and feather-bedding agreements); local promoters, however, lack any rent-yielding unique assets. Scope for opportunism arises as performers move rapidly from one venue to the next, while accounts arrive only months later, when the expenses and gross revenues are costly or impossible to audit. One booking agency, Frank Barsalona's Premier Talent, rose rapidly in the late 1960s to dominate the scheduling of rock concerts partly through forming continuing informal partnerships with local presenters. The agency would guarantee that the local presenter broke even if a concert unexpectedly lost money; the presenter presumably sustained the relationship by rendering honest accounts. Enforcement of contracts in the creative industries depends heavily on the power of repeated interactions among parties who value their reputations for cooperative behavior.[46]

Agents and Job-Matching

Agents play several roles in the creative industries. Literary and concert-booking agents perform a job-matching function—lining up trade-book authors with publishers and popular musicians with performing venues. The gatekeeping role assumed by the literary agent results from excess supply due to *art for art's sake,* but the job-matching task would remain even if every manuscript found a publisher and every publishing house issued its preferred number of titles. Creative outputs such as book manuscripts differ pervasively *(infinite variety);* in the short run that matters for the matching process, publishers are likewise heterogeneous—in their fields of interest, backlists, current congestion of the publication pipeline, promotion methods, and so on. Indeed, most creative activities have to solve some job-matching problem—allocating vaudeville acts among theatres, big bands among ballrooms, classical pianists among concert series, actors among movie projects. The organization of matchmaking agents varies greatly. In classical music (this section's empirical focus), the agency function at times has been nearly monopolized.

Hence, we need to pin down the basic functions that agents perform—which side of the market they represent, and what forces determine their number relative to the parties whom they represent.

Whose Side Are You On?

Why do matchmaking agents represent the side of the market that they do? We usually see them representing sellers. But that arrangement is not universal (think of executive search firms and other such independent purchasing agents), and for some matching problems it is hard to see any case for attaching the matchmaker's services to one side rather than the other. Why should marriage brokers represent would-be brides rather than potential grooms? The gatekeeping role suggests one way to explain which side is represented. With potential authors vertically differentiated *(A list/B list)*, and a cost incurred to assess the quality and prospects of any given author, it would clearly be inefficient for every publisher to read every manuscript that tumbles over the transom. Publishers, we assume, cannot themselves divide up the stream of submissions and share their assessments with each other, while an agent who finds a good manuscript knows which publisher will find it most attractive. Agents, as they decide which authors to represent, "pool" the task of assessment, performing it once for the benefit of the several publishers to whom the agent refers promising manuscripts. All gain from this pooled screening.[47] The gain increases with the number of writers who never find a publisher.

Matching is a two-sided process: publishers need to find out about writers, and writers about publishers. Suppose that publishers' policies are readily revealed by their backlists and reputations with their authors for capability and integrity; authors' qualities are more costly to expose. If the transacting parties on one side of the deal are represented by agents who pool information and economize on its transfer to the other side, pooling the more costly assessments of authors and their manuscripts beats pooling the less costly assessments of publishers' traits on behalf of authors.[48] The same logic applies to the differential importance of the information to the other side. A publisher loses heavily if a celebrity author's book flops, but the celebrity might be nearly indifferent about which of several mainline trade publishers publishes the book.[49]

Other influences also weigh in. Suppose that it costs the same for a publisher to size up a prospective author as for the author to evaluate a prospective publisher. Even after the gatekeeping agents have swept out the losers, the authors remain more numerous than the publishers. (Think of each house serving as exclusive publisher to a number of authors.) If agents are to pool information on one side of the market only, they should pick the more numerous authors, thereby consolidating more information than if they represented the less numerous publishers. The agent's gatekeeping function, which excludes many authors, is really a special case of this "differential numbers" effect.

The agents' assignment might also be determined in part by functions that they perform along with job-matching. The auditorium sits immobile at the corner of Fourth and Main, while someone must ensure that the band finds its way there; after the concert the band rides off into the night in its limousine, and somebody must check that the venue makes its obligated payment to the artists. Agents will tend to represent the artists, because they place higher value on such auxiliary services. In summary, job-matching agents likely represent the side of the market that has the most clients, is subject to more exclusionary gatekeeping, is less able to reveal credible information about itself, has less need for differential information about the other side, and/or has a greater need for the agent's ancillary services.

How Many Agents?

How many parties will each agent represent, and how many agents serve the market? Suppose that what each seller offers and each buyer seeks could be fully described as quantitative, measurable traits. Suppose also that the best match-up of heterogeneous buyers and sellers can be reduced to a computer algorithm.[50] Then matchmaking, like an automated stock market, would be a "natural monopoly" efficiently run by a single agent. Where the agent's role arises from the costliness and incompleteness of information, however, the virtue of a monopolistic job-matcher breaks down. Wherever the agent's role centers on the costly pooling of heterogeneous information, the agent's capacity to absorb and exchange information is limited, and with it the number of clients each agent can represent. If agencies perform auxiliary services subject to scale economies, the agency firm might become a large partnership of individual agents. But the more idiosyncratic the information and the more often and extensively it changes, the more personalized the agent's role becomes, and the fewer parties he can serve.[51]

The organization of agency relationships is also driven by problems in governing the deal between the agent and the party represented. Take the popular dance bands that toured among ballrooms and theatres in the 1920s through 1940s. Its booking agent might schedule a band at one location for Saturday night, then uncover and accept a better offer for that same night without renegotiating the first—resulting in an enraged presenter and a disgraced bandleader. Or the band might fail to show up due to personal disor-

ganization that could be warded off by a suitably hectoring agent. These problems led to the rise of relatively large booking agencies, notably Music Corporation of America (later MCA). The organizations' sizes were partly due to scale economies in managing the logistics of bands' travels, but the valuable fixed assets of these booking organizations also ensured their diligent and consistent performance to parties on both sides of the market.[52] That is, the victim of a breach of contract could locate the perpetrator and bring suit that could (if successful) allow it to recover its losses from those valuable assets. A fly-by-night operator by definition offers the contracting party no such protection.[53] In the creative industries, large-size firms sometimes owe their prevalence not to conventional scale economies, but to the value of large blocks of exposed assets as collateral for proper performance of obligations. The firm with exposed assets has incentive not to cheat on its obligations; the contracting partner, recognizing this, has more incentive to sign.

While the job-matching agent might formally represent either seller or buyer, there is a reason why he must align with one or the other. Linking seller S and buyer B may be the best deal for both parties, one that would be picked by an agent representing either sellers or buyers. The contract price however, divides the benefit between S and B, and the agent negotiating that price can serve the interest only of one party at the expense of the other. The problem of interest conflict when agents get onto both sides of a transaction recurs throughout the creative industries. MCA fell into the role when it diversified its activities from agency to the production of films, and was forced to divest its agency operations (see Chapter 7).[54] Creative Artists Agency under Michael Ovitz became a successful and innovative packager of scripts and creative talent for TV and cinema films, but also found itself compromised by both hiring and representing its own talent.

Agents in Classical Music

The booking of classical musicians—singers, instrumental soloists, and orchestra conductors—provides a case study of agents' organization. In the United States, booking agents arose in the 1880s as travel became efficient enough for large numbers of musicians to take to the road. The same pattern emerged in Europe of individual agents, usually well-connected people with musical training and often substantial family wealth. The representation role focused on the artists rather than the venues, for several reasons suggested previously.[55] Personal integrity was problematic, for many musicians (then as now) would perform for little pay in order to build their careers; some agents took their money for promotion and did nothing, confident that the artist

was going nowhere.[56] Large agencies emerged, such as Ibbs & Tillett in London and Henry Wolfsohn in New York, for the expected reasons. Wolfsohn employed a large staff of road agents to sell his artists to local sponsors in the United States. Ibbs & Tillett carefully built its reputation for integrity in the wake of personal agents who ran into financial trouble and began appropriating funds due to artists; it even took part in bailing out one failing competitor.[57]

The size distribution of agencies' operations has varied over time. Small (personal) scale always yields an advantage in the agent's opportunity to help build the artist's career. For artists of high promise, these services frequently dominate the ones that large agencies can supply efficiently.[58] Large-scale agencies plied different strengths. Arthur Judson's Judson Concert Management arose in the 1920s, becoming Columbia Artists Management Inc. (CAMI). Before blossoming as an agent, Judson had managed both the Philadelphia and New York Philharmonic orchestras. His specialty was representing conductors, and his continuing orchestra-manager job gave him easy access to any promising guest conductor or soloist who appeared with the New York Philharmonic. Conductors are key decisionmakers in recruiting orchestral soloists, so Judson's dominance as an agent for conductors gave leverage for representing soloists. As intermediary in many deals, he could punish any musician who eschewed or dispensed with his services. By 1930 he had bought up the six largest competing agencies in New York. Judson clearly saw himself as the honest broker between music presenters and musicians, and he managed repeatedly to be on both sides of transactions without protest from the contracting parties.[59]

Judson's successor, Ronald Wilford, preserved CAMI's dominance, representing about 100 conductors and 800 musicians overall (mostly singers). Judson's unique advantage of a central position in the information web was sustained, and conductors remain the linchpins for leveraging the placement of other CAMI soloists. CAMI's objective was clearly to maximize benefit for the musicians represented, however. Judson's domain had been the United States, while Wilford became active internationally and began pulling talent away from agents abroad. Wilford pioneered the practice of conductors holding down the job of musical director with several orchestras at once, while topping their incomes with the cream of guest-appearance fees. His success may explain why, from 1960 through 1990, the price of conductors' services was estimated to rise three times as fast as a laborer's average weekly wage.[60] Apparently, either Olympian optimization or outright monopoly over talent can sustain a large-scale agency organization in classical music.

Another success in large-scale organization was the approach to promoting classical music (and its performers) innovated in 1921 by the Redpath-Chi-

cago Lyseum Bureau, particularly its partner Dema Harshbarger. She devised the Civic Music Committee, for which she would identify leading citizens of a town or city and persuade them to organize a concert series, guaranteeing the fees for visiting artists from subscription sales and undertaking local publicity. The Civic Music Committees expanded local demand for concert tickets, reduced the agent's risk of a nonpaying presenter, cut the agent's promotion costs, and usually recruited their soloists from Chicago Lyseum. By 1932 the energetic Ms. Harshbarger had papered 32 states with Civic Music Committees. She was glimpsed on the horizon by Judson, who made haste to defend the east coast with his own Community Concerts Corporation, which mimicked Civic except for requiring that touring artists be taken from his firm's list. Later (1955) this tying arrangement was enjoined under the antitrust laws, but by then declining public interest in the vocal recital, the mainstay of these series, had shriveled its importance.[61]

For ensemble performers, the agent's function in classical music devolves to contractors, who organize orchestras and choral singers for free-lance work. New York's performing groups, except for the New York Philharmonic and Metropolitan Opera Orchestra, are staffed by free-lance musicians assembled for individual concerts by a dozen contractors, who maintain rosters from a pool of about 1,000 free-lancers. The contractor's specialty is knowing exactly the various players' strengths and weaknesses and their compatibility with the conditions of various jobs. They also ensure that the musicians' wages are placed in escrow accounts and payments are subsequently made. The contractors perform a gatekeeping function in admitting new talent to their rosters, but this role was restricted in 1994 by the musicians' union's success in forcing each of the biggest free-lance orchestras to establish its own roster of regular players, whom the contractor is required to contact first. This seniority system obviously impairs the ability of the contractor to assemble the best available group on each occasion.[62]

4

Artists, Starving and Well-Fed

The imperative *art for art's sake* implies that artists will choose low-paid creative work over better-paid humdrum labor. The gatekeeping process further implies that the artist community will include at any one time those who have passed through the gateway and claimed their prizes (great or small), those still clamoring at the gate (earning apprentice wages at most), and those who have relegated their creative urges to hobby status. The mean levels and dispersions of artists' incomes should reflect these processes, as should the split between their earnings from creative and from other work.

Artists' Success: Superstardom

To explain the great dispersion of artists' earnings and the legendary riches of stars, economists have developed a model of the economic position of superstars.[1] Assume that artists (rock performers, say) differ in quality in the eyes of fans, and all who might attend rock concerts agree on who is the best, second best, and so forth (the *A list/B list* property). Assume, too, that we have a way to quantify these quality differences. To attend a rock concert, fans pay the price of the tickets plus the time diverted from other activities. The higher the artist's quality, the more utility fans get, given their time and money costs. In particular, the model assumes that lower-quality performers are poor substitutes for stars. Great and mediocre concerts take up the same time. A strong preference for quality implies that the fan would forgo many concerts by average rock performers in order to hear one by a superstar, even if the superstar's tickets cost more.

The artist incurs a cost of performing. It does not obviously cost the superstar more to perform a concert (she might offer a more glittery light show, but the argument requires only that the cost for artist and light-show together rises less than proportionally to quality). The superstar has an enviable economic choice. If her tickets are priced the same as average artists', she will