1 Introduction

Despite their briefcase reputation, economists have shown a remarkable fascination with farming and its various economic details. This might be expected of the agricultural economists in the profession—after all, that's their job—but it also has been true of general economists. The economics literature is filled with discussions of farming, especially in the context of share contracting, specialization, and the division of labor on farms. This literature includes the classical economists Adam Smith, who noted the moral hazard incentives inherent in some farmland contracts, and John Stuart Mill, who identified the effect nature's seasons had on the ability to specialize in farm production. It also includes modern scholars such as recent Nobel Prize winner Joseph Stiglitz, who formally introduced the profession to the principal-agent paradigm in the context of farming, and Yoram Barzel, who first noted the multiple contracting problems that arise on both the farmer and landowner sides of the market. Economists have been especially enamored of share contracts, and the inquiry into their existence and efficiency has led, almost directly, to the modern theory of contracts.[1] Share contracts have been common worldwide for centuries, but perhaps more surprising, these "cropshare" contracts—as they are called by American farmers—remain common in modern developed agriculture as well.[2] Despite numerous theoretical inquiries into agricultural share contracts, however, and despite their common occurrence, surprisingly little is widely known about their details.

Although there is more to the "nature of the farm" than just share contracting for land, it is fair to say that our economic understanding of farm organization beyond share contracting is limited. For example, prior to the decision about the type of land lease, a landowner must decide whether to rent the land or farm the land for himself. What determines this choice? Furthermore, this decision of ownership versus contracting applies to the other major assets on the farm as well as to the land, and the patterns of owning are vastly different for equipment than for land. But what explains these different patterns of ownership? Although farming is one of the last industries in which the majority of firms are owned by families, over time the scope and depth of family ownership and production has eroded. What explains this dominance and partial erosion? Although it is impossible to analyze all organizational issues on the farm, we examine these questions in particular.

This book has several objectives. First, we wish to demonstrate the power of the transaction cost approach in understanding many organizational features of agriculture. Though we devote a great deal of attention to the issue of contract choice, we also examine the ownership structure of the farm and the question of vertical control. Though our specific models vary from chapter to chapter, the overriding theme is that contracts and other patterns of ownership are chosen to mitigate transaction costs (to be defined momentarily). In agriculture, transaction costs are heavily influenced by Mother Nature. Nature's uncertainty, via weather and pests, allows for suboptimal asset use, and through seasonal forces nature imposes constraints on production cycles that are not often found in the production of other

commodities. We take pains to explore these constraints placed by nature and other farming details in order to understand the transaction costs that arise under different forms of organization. We then use this to derive testable propositions about contract and organization choice.

Second, we wish to contrast our transaction cost approach with theories based on, or including, risk-sharing motives. Virtually all economists who study the economics of organization recognize that incentives are important. The oldest, and most common, model of share tenancy is one in which there is a trade-off between the incentives of the farmer and his aversion to bear risk. Throughout the book we will refer to this model as the "principal-agent" model. In this model the contracts that provide the best incentives also generate the most risk. Risk is, in effect, the cost of incentives. Our model is also based on incentives, but incentives spread over many decisions made by the farmer, the landowner, or equipment owner depending on the problem at hand. As it turns out, many predictions from these two models are at odds with each other. We devote considerable space to empirically contrasting the two.

Third, we want to study organization (the ownership and the contracting) of modern agriculture in North America. That is, our book is a detailed study of the organization of a single industry—in both a historical and a contemporary context. Until quite recently, the economic analysis of farm contracts and organization has focused on historic and developing country cases.[3] As we mentioned in the preface, this has often led many to conclude that cropsharing does not exist in modern farming communities. Among modern agricultural economists who study North American farming, the focus has not been on contracts but on neoclassical analysis of costs, production, and commodity markets. A study of North American farming, where technology is advanced and where capital markets are well developed, provides an opportunity to test theories often applied to only Third World settings, and to explain a series of farming puzzles that have generally been ignored. Although our book relates to the literature that spans the fields of development, economic history, and agricultural economics, our book is not a literature survey. We make frequent references to such literature in order to provide context for our models and results; however, many excellent surveys already exist (for example, Otsuka and Hayami 1992), and we feel it unnecessary to repeat them.

Finally, our objective is to explain a variety of economic organizational puzzles in farming. Consider the following. Grain farmers use a large machine called a combine to harvest grain, and depending on the crops grown and the size of farm, this machine may only be used as little as two to three weeks per year. Combines are also one of the most expensive pieces of equipment a farmer might own, with larger models costing $150,000–$200,000. This is a classic case where economists would predict, based on the high cost

and low utilization rate, that farmers should rent combines. Yet most farmers still own their combines and leave them idle for most of the year.

Our book examines many such puzzling observations, and while we develop numerous formal models within our basic transaction cost framework, the book is also an empirical analysis of testable predictions using contract and organization data. We use five separate data sources, supplemented with census data, to provide the bulk of our information. These data allow us to use standard econometric methods to test our predictions. In addition, we rely on historical case studies, on such topics as Bonanza farms and custom combining, to supplement our statistical analyses.

1.1 The Transaction Cost Approach to Contracting and Organization

Transaction Costs and Property Rights

At several places in this introduction we have mentioned the "transaction cost" approach, and the time has come to explain what we mean by this phrase and how we believe it differs from other economic approaches to organization. The transaction cost approach begins, of course, with Coase's classic works on the firm (1937) and social cost (1960).[4] In the latter paper Coase pointed out that when transaction costs are zero, the allocation of resources is independent of the distribution of property rights. Ironically, his most famous example is an agricultural application: the cattleman dealing with his crop farming neighbor over tresspassing cattle. When transaction costs are zero, the number of cattle tresspassing does not depend on whether the cattleman possesses the right to trespass or not. The outcome is determined by the joint wealth maximizing level of output on the two farms.

It remained for Cheung, in his pathbreaking book, to recognize the general implications of Coase's work to contracts. Cheung (1969) showed how, under the conditions of zero transaction costs, a cropshare contract could achieve the same outcome in terms of crop output as a cash rent contract could.[5] The result is completely general. When transaction costs are zero, it does not matter how the ownership of the inputs and outputs is distributed by the terms of a contract. Farmers can control land through cash leases, share contracts, or ownership; farms can be family run, sole proprietorships, or they can be large-scale corporations; and farms can be integrated completely from breaking ground to baking the bread, or disintegrated to the point of owning a wheat field for one day—it matters not one iota.

At this point many economists, and others, are ready to abandon Coase's idea. It cannot be stressed enough, however, that Coase's point was not that a model based on zero transaction costs had any relevance for understanding economic organization. Just the opposite. He

argued that any analysis of economic organization must hinge on an examination of the transaction costs involved. His argument is as follows: If transaction costs equal zero, then property rights are perfect and organization does not matter; if these costs are not zero, then the explanation of organization lies in transaction costs. The grand hypothesis of the transaction cost approach is that contracts and organization are organized to maximize joint wealth net of transaction costs.

All of this, of course, begs the question "What are transaction costs?"—a question that is made more pressing given that Coase himself has never defined the term but instead just provided examples. Indeed, the transaction cost approach has been hindered at times by ambiguities in language and a general reluctance to define terms—especially the terms "transaction costs" and "property rights." In fact, there are two well-developed concepts of transaction costs in the economics literature. The first, developed by Demsetz (1968), defines transaction costs as the costs of transferring property rights in a market exchange. This is the definition found in the *The New Palgrave Dictionary of Economics*.[6] This approach typically posits some type of "transaction technology" that taxes the transaction and acts in many ways just like a tax. Because this notion of transaction costs was developed to analyze the volume of trade, its major drawback is that it is not useful for examining questions of contract and organizational choice. In another survey article, Allen (2000) calls this the "neoclassical" definition of transaction costs because of its emphasis on the volume of trade.[7]

We do not use the neoclassical concept of transaction cost. Instead, we use what has been called the "property rights approach" to transaction costs, where these costs are defined as *the costs of enforcing and maintaining property rights*—regardless of whether a market exchange takes place or not. Property rights, in turn, are defined as the ability to freely exercise choices over the asset in question. Transaction costs include the deadweight losses that result from enforcing property rights as well (Allen 1991, 2000; Barzel 1997).[8] As a result, transaction costs are more than the costs of a market exchange. That is, property rights may be required to be enforced in a private contract, through courts or other third party agencies, against thieves, or across market transactions. We employ this concept of transaction costs throughout our book because it is complete enough to explain organizational choices and because it more closely aligns with the modern literature on contracts and organization.

In order for transaction costs to exist, two conditions must be met. Information must be costly to obtain, and assets must be variable in their quality or characteristics, and alterable by man. That information must be costly is rather obvious. If everything is known, then enforcing and maintaining one's claim to property is redundant.[9] That assets must be both variable and alterable is perhaps less obvious. Essentially the only way someone can systematically infringe on another's property rights, and therefore make necessary efforts

to enforce or maintain them, is for a confusion to exist over the effects of nature and the actions of people. The more uncertainty there is in nature and the more individuals are able to influence final outcomes, the larger the transaction costs. What makes farming such a rich field for a transaction cost approach is the obvious impact of Mother Nature, and the equally important impact of farming decisions on crop output.

Our approach to farming contracts and organization is a transaction cost approach because we develop a set of specific models that depend on the ability of contracting parties to police their interactions with each other. Although farmers enter into contracts with various parties (for example, custom combiners, laborers, landowners, pesticide applicators, storage firms), these contracts are never complete and problems arise in their enforcement due to nature's uncertainty and the complexity of the assets involved in production. Farmers can hide bales of hay that were intended to be shared with landowners, harvest crews can arrive late causing a reduction in crop value, and, of course, hired workers can generally shirk their duties. Transaction costs are the costs of engaging in and preventing these activities, along with any lost gains from trade that result. Both landowners and farmers seek to mitigate these costs. A theme throughout the book is that contracts have incentives that often substitute for direct monitoring. As a result, contracting problems are often solved by altering incentives given the constraints imposed by the particular farming technology, the role of nature, and the potential gains from specialization.[10]

Our transaction cost approach is in the tradition begun by Coase and Cheung, Alchion and Demsetz (1972), and most recently exposited by Barzel (1997). It is similar to Williamson's (2000) discussion of the New Institutional Economics, but differs from his (1979) view of transaction costs that emphasize the role of specific assets in determining organizational forms. Recently Hart (1995, chap. 2) developed what he calls a "property rights approach" to firm ownership. Our book has a similar spirit to that of Hart, but its method is broader. Hart's framework stresses the investments individuals make under different ownership structures. He notes that investments may become sunk, raising the costs of negotiating over the gains from trade in future periods, and that different asset ownership structures will influence investment and total value. Our model is more general and more relevant to farming where investment in such assets tends to play a minor role.

Five Important Ideas

Five important ideas define our framework and require discussion. First, we assume that all parties (farmers, landowners, other input owners) choose contracts and organizational forms because they *maximize the expected value of the relationship*, given the characteristics of all parties, the desired output, and the attributes of assets such as land and equipment. By focusing on joint wealth maximizing allocations, we ignore issues of bargaining and surplus division. In addition, the empirical implementation of bargaining strength seems impractical

given the data available for our study. In the context of competitive farming where specific assets are minimal, this issue is relatively unimportant. Competition among farmers for land, and among landowners for renters, and competition between on- and off-farm opportunities generally determine the returns to individual factors of production within narrow bounds. This method assumes that "natural selection" has resulted in the most valuable contract or organization being chosen, and is based on the idea first proposed by Alchian (1950).[11] Farmers and landowners, like everyone else, are keenly aware of their incomes and just as aware of the effect of one type of contract over another on their bottom line. Given the general stability of farming communities, it seems only reasonable to assume that contracts and organization are fundamentally driven to maximize wealth.[12]

Second, while we abandon some aspects of typical contracting models, *uncertainty remains a crucial component*. Uncertainty allows individuals to exploit an exchange at the expense of the other party because it masks their actual effort. This factor is important in agriculture because weather, pests, and other natural phenomenon contribute so much to the final output. In a land lease, for example, uncertainty from weather and other natural forces means that the farmer has the opportunity to "exploit" the landowner in several ways: undersupplying effort, overusing soil quality attributes, and underreporting the shared crop, to name a few. The type of behavior we often focus on is moral hazard (or hidden action) where the farmer, landowner, or other asset user does not bear all of the costs of his actions. Moral hazard is just one type of transaction cost phenomenon, and like transaction costs in general, uncertainty is necessary for it to exist. Still moral hazard is not the only incentive effect we study; there are also measurement and enforcement costs arising from uncertainty.

Third, *all assets are complex in the sense that they are comprised of many attributes*. When assets are complex they create an opportunity for transaction costs to arise for every attribute, which subsequently allows for divided ownership over the various attributes because multidimensional assets are nontrivial to measure. A plot of land, for example, is characterized by its size, terrain, nutrients, moisture, soil type, and so on. Different ownership and contract types affect the various attributes in different ways, creating trade-offs. These trade-offs allow us to explain the choice of organization based on different transaction costs.[13]

Fourth, though nature has a random component in uncertainty, she has a *systematic component we call seasonality*. For contract choice we focus on the random aspect of nature. Poor harvests, soil erosion, and nutrient and moisture depletion can be blamed on acts of nature, even though land overuse may arise from improper tilling and pest control or other practices. Although random acts by nature are a common element in modern contract theory, our idea that nature also plays a systematic role is not found in the contracting literature.[14] Seasonality, instead, refers to crop cycles, the number and length of stages, and timeliness.

In part III of the book the predictable aspects of nature—its seasonality—become very important because they limit the degree to which farmers can specialize in production. Most types of farming are greatly restricted by nature. Both plant and animal crops have "growing seasons" that restrict the nature of farm production. As a result, farmers are seldom able to exploit many forms of economies of size and tend not to develop into large corporate farms.

Finally, throughout our book we attempt to explain farming contracts and organization in the context of *risk neutrality*. As we note below, this is a significant departure from most attempts to explain such matters. Risk aversion is an assumption about preferences that we do not make. There are several reasons for this. First, in modern agriculture, where the ability to avoid risk though insurance and asset markets is so well developed, it seems implausible that farmers and landowners would use their land contracts to further avoid risk. Second, as we show in chapter 6, empirical implementation of even the simplest risk-sharing hypotheses is difficult, and often impossible, because of the stringent data requirements. Third, by avoiding the complexity of preference and uncertainty modeling, we are able to develop models that yield clear and testable predictions. Ultimately, the importance of risk sharing is an empirical matter. Though we push our risk-neutral model in all directions, in part II we pause to compare our results with those based on risk sharing, where we find no compelling grounds to abandon our assumption of risk neutrality.

Modeling Transaction Costs

To summarize, we use a series of transaction cost models in which all parties are risk neutral, in which all assets are complex, and where nature is both an uncertain and seasonal force. These conditions make it costly for the contracting parties to identify exactly the input contributions of their counterparts and, similarly, make it costly to identify the quantity and quality of the output. In agriculture, nature's seasonal forces limit the gains from specialization and the ability of parties to monitor each other.

Although the specific models we use vary from chapter to chapter with the details of the questions we address, we outline the basic framework and its characteristics in this first chapter. In our models output takes the following general form:

$$Q = h(\text{land, labor, capital}) + \theta, \tag{1.1}$$

where Q is the observed harvested output that is assumed to have a unit price and $\theta \sim (0, \sigma^2)$ is the randomly distributed composite input of nature. Exactly how each input is defined depends on the question at hand. For example, when analyzing the choice of contract between cash rent and cropshare, the critical inputs are the unpriced land attributes supplied by the landowner (such as fertility and moisture content) and the labor effort of the farmer. When analyzing the decision to rent or buy an asset such as land, we make a distinction between the priced attributes of the land (such as size measured in acres) and the unpriced

attributes (such as soil fertility). In other problems the various inputs may be seed, fertilizer, pesticide, or other nonland nonlabor inputs; labor time as opposed to effort; the number of tasks; and so on. The complexity of the models progresses throughout the book. The chapters on contract choice are simpler because they suppress issues of timing and specialization and focus on moral hazard and enforcement-monitoring costs. Later, these other aspects are introduced to discuss ownership types and vertical integration.

As these complications are added, the specific form of production function in equation (1.1) alters, but regardless of the specific form, our production function contains in it the basic structure necessary for the existence of transaction costs. Notably, output is determined by human action $h(\cdot)$ and nature θ. Inputs are not observable, and although the output Q is observable, it is a complex asset and cannot be perfectly measured. As a result, effort can be altered to suit the private interests of one party at the expense of the other. For example, low levels of output that result from low labor inputs could be blamed on poor weather conditions. Suboptimal level of inputs might also include applications of the inputs at an incorrect time. The simple additive uncertainty component also simplifies the analysis in models that maximize expected values.

In addition to this basic structure, we also make several assumptions regarding the function h. First we assume that h always has positive but diminishing marginal products. Second, we assume that all inputs are independent of one another. Both of these assumptions are intuitively appealing and create models that generate clear predictions. The assumption of independent inputs simplifies the model and increases the number of testable implications. Not only do we have no a priori theoretical grounds to assume which inputs are substitutes or complements, but there is empirical justification for their independence. First, were they not independent, contracts could adjust some input prices upward, others downward, to influence farmer behavior. This, however, is not observed for the cases we study.[15] Second, in chapter 5 we show that input cost sharing in cropshare contracts exhibits an all-or-nothing dichotomy; that is, input costs are either shared in the same proportion as output or are not shared at all. This result is consistent with independent inputs.

Price taking is another common feature of our models. For example, we always assume that the opportunity cost of the farmer's input is the competitive wage rate w per unit of farmer's effort, and the opportunity cost of the unpriced land inputs is r per unit. We also assume that farmers sell their output on world markets and that they cannot influence this output price. These assumptions seem reasonable in the context of modern agriculture with world trade, where individual farmers are small relative to both the input and output markets.

The logic of our models is straightforward. Once the precise production structure is constructed, we begin by deriving the first-best, zero transaction cost outcome for a specific problem. By "first-best" we mean not only that inputs are used in the optimal amounts, but also that inputs are fully specialized and applied at the appropriate time. We use this outcome

as a benchmark to compare the actual contracts and organizations, because the presence of transaction costs makes this outcome unattainable. The second step is to examine the various contracts or organizations and determine the optimal value functions under each and to examine the comparative statics of these functions. Next, we assume in all cases that the joint wealth maximizing contract or organization is chosen. Finally, we test derived predictions using our data from North America.

1.2 The Role of Risk in Contract Economics

The transaction cost approach, with its trade-offs of one incentive against another, can be contrasted with the classic "principal-agent" approach to share contracting where it is assumed that contracts are designed to spread the risk of crop farming away from the farmer and partially on to the less risk averse landowner.[16] The fundamental idea that farmland contracts are designed around a trade-off between risks and incentives is commonplace among economists. Indeed, Stiglitz (1987) writes, "The sharecropping model has served as the basic paradigm for a wider class of relationships known as principal-agent relationships" (321), and Sappington (1991) notes, "The classic example of the principal-agent relationship has a landlord overseeing the activities of a tenant farmer" (46). In their important study Otsuka, Chuma, and Hayami (1992) claim that risk aversion "provides the most consistent explanation for the existence of a share contract" (2012). Relying on the standard risk-sharing framework, they further state: "As in typical agency models, the most obvious factor to be accounted for in considering the optimum contract choice is the presence of uncertainty coupled with the risk aversion of the contracting parties" (1987). The dominance of this approach in modeling the behavior of farmers and landowners is not limited to those studying developing countries or economic history (for example, Otsuka, Chuma, and Hayami 1992; Townsend 1994). It is routine among agricultural economists studying farm behavior—including acreage and crop choice studies as well as contract studies—to assume that farmers are risk averse and stress the role of risk sharing in determining behavior.[17]

Despite the prominence of the risk-sharing paradigm, the empirical evidence to support it is scarce.[18] In agriculture there has been little empirical work at the contract level and nearly all of this has been in developing economies (Otsuka, Chuma, and Hayami 1992). In one of the early studies to confront risk sharing and contract choice, Rao (1971) found that crops with high yield and profit variability were less likely to be sharecropped, directly refuting the anecdotal evidence originally provided by Cheung (1969). At the same time, studies by Rao and others (for example, Higgs 1973) tend to use rather small samples of highly aggregated data, making clear inferences difficult.[19] In chapters 6 and 7 we examine

Table 1.1
Crop riskiness and share contracting

Region	Yield coefficient of variation (% of share contracts)	
	Corn	Wheat
British Columbia	.27 (20%)	.18 (79%)
Louisiana	.29 (62%)	.21 (76%)
Nebraska	.12 (69%)	.11 (86%)
South Dakota	.14 (64%)	.25 (61%)

Sources: Appendix A and Allen and Lueck (1995, 1999a).

the risk-sharing hypothesis in detail, but even a glance at the facts suggests that this is not likely the case in modern farming. Table 1.1 shows the coefficient of variation for two major crops (corn and wheat) in our four distinct regions as well as the prevalence of share contracting for farming in those same regions. Contrary to the risk-sharing thesis, land used to grow high variance crops is not cropshared more often. Table 1.1 actually suggests the opposite: Crops with less yield variance are more often cropshared. As we show in chapter 6, this finding is consistent with transaction cost models that focus on measurement costs.

Our approach contrasts with the risk-sharing model.[20] In the classic risk-sharing model, a typical model assumes that a principal maximizes some objective function subject to an agent's incentive and individual rationality constraints. For sharecropping, most of these models postulate a risk-neutral landowner (principal) leasing land to a risk-averse farmer (agent). These models generate a trade-off between risk avoidance and imperfect incentives. A principal who cash rents to an agent has no incentive problem, but the agent "bears all the risk." By sharing with an agent, the principal suffers from agent moral hazard, but the agent no longer bears the full risk of the project and the payments can adjust accordingly. By using the transaction cost approach, we avoid the empirical difficulties of risk-sharing models while retaining other aspects (for example, uncertainty and moral hazard) of these models. We abandon both the principal-agent distinction and the assumption of risk aversion, and instead, we assume all parties are risk neutral.[21]

By treating both parties as risk neutral, we avoid the problem of defining which party is the principal and which is the agent, and also which party is more or less risk averse.[22] In modern farming it is especially difficult to establish such a dichotomy because farmers and

Table 1.2
Characteristics of farmers and landowners

Variable	British Columbia 1979	British Columbia 1992	Louisiana 1992	Nebraska–South Dakota 1986
Average age				
Landowners	52.8	57.0	63.9	≈50
Farmers	40.9	47.2	46.5	≈40
Average years of education				
Landowners	8.3	NA	NA	NA
Farmers	11.0	NA	NA	NA
Average acres of owned land				
Landowners	NA	499.5	748.5	661.2
Farmers	NA	439.4	122.7	435.5
Average acres of owned land				
Farmers with no leased land	NA	147.4	418.4	NA
Farmers with only share leases	NA	412.1	116.8	NA
Farmers with only cash leases	NA	241.3	185.4	NA
Percent of women				
Landowners	NA	NA	NA	34
Farmers	NA	NA	NA	6
Percent of landowners with farm experience	60	69.5	57.2	NA
Percent of farmers				
that rent and own land	NA	93	57	NA
that rent and rent out land	NA	6	6	6
that both share and cash lease	NA	10	24	23

Sources: Appendix A and Allen and Lueck (1995, 1999a).
Note: NA = not available.

landowners have nearly identical demographic characteristics and because farmers make virtually all the decisions, contrary to their oft-designated "agent" status. Table 1.2 points out a number of characteristics of the farmers and landowners that are common across all of our data sets. For example, the table shows that 60 percent of the landowners are or were at one time farmers. Furthermore, table 1.2 shows that renters are often landowners, and in some cases (6%) rent out land simultaneously as well as hold both share and cash rent contracts. The similar social-economic background and demographic features of farmers and landowners along with the coexistence of owning and leasing are inconsistent with a model that posits dichotomous preferences and risk sharing.

Another advantage of our risk-neutral approach is that we do not require data on exogenous risk to test the implications of our models. A significant difficulty in conducting tests of risk sharing lies in finding a reasonable empirical counterpart for the pure random variance

in output caused by nature.[23] Obtaining such data is difficult because output data at the contract level are "contaminated" by inputs from nature and the farmer. Finding such measures in studies of franchising and other areas has proved difficult, and as a result scholars have either ignored them or relied on proxies that may seem reasonable, but are not often clearly linked to the underlying theoretical model and may be highly endogenous to the firm's behavior. To test explicitly for the negative relationship between risk and incentives, it is crucial to have such data, which are notoriously difficult to obtain. Such data are not necessary to test our model based on transaction costs and risk neutrality.

1.3 The Role of Government

Government intervention in agriculture in the United States and Canada is long-lived and prominent.[24] Yet, for the most part, we ignore the role of government in affecting the choice of contracts and organizations in agriculture. We have two primary reasons for doing this. First, and most important, these interventions generally do not affect the kinds of incentives we study and thus do not alter the relative costs and benefits of various contract and organizational choices. Second, to the extent there is an impact, our statistical data cannot readily isolate the effects of government. To illustrate these points, we first describe the basic features of agricultural policies and then link them to our study.

In the United States, large-scale intervention in agricultural production began with the New Deal legislation of the 1930s. This legislation established a set of policies providing for price supports and production controls, crop insurance and disaster payments, export subsidies, subsidized farm credit, land and water conservation, subsidized food distribution, and expanded research and extension activities.[25] These programs have varied over time but can be summarized as follows. Price supports (through target prices and nonrecourse loans) and acreage restrictions have been the basic policy for cotton, rice, wheat, and feed grains (barley, corn, grain sorghum, oats). Producers of these crops were entitled to government "deficiency payments" that cover the difference between the target price and the prevailing market price.[26] Soybean prices have been supported through loans and government purchases. Sugar prices have been supported by import restrictions and some price supports, and thus have impact on producers of sugarcane and sugar beets. Many other products, including milk and certain fruits, vegetables, and specialty crops (for example, almonds, oranges), are governed by marketing orders. These marketing orders typically limit the production of the governed commodities by specifying quality and other details of the product.[27] Milk marketing is different from other marketing orders in that price supports are explicitly used to limit output. The 1996 Farm Bill introduced some important changes in U.S. policy.[28] The target price system was replaced with a series of "transition payments"

made directly to farmers. These payments, based on historical production, are in place for the years 1996–2002 and are thus "decoupled" from current production decisions. The 1996 bill did not alter the basic structure of the programs for dairy, peanuts, and sugar.

Several other policies are worth noting. Since the 1930s limits have existed on the amount of government payments any single farm(er) can receive, although there are methods of avoiding these restrictions (as we note in chapter 9). Farm capital is generally treated more favorably in the federal tax code than is nonfarm capital. And there have been conservation programs (Soil Bank, Conservation Reserve) since the 1950s, which pay farmers to take land out of production. There are other agricultural programs but these do not relate to the topics we study.[29]

In Canada government intervention has been less intrusive and of a slightly different form.[30] There have been no systems of target prices and government payments for grain. Wheat producers, however, must sell their crop through the Canadian Wheat Board (CWB).[31] Originally established in 1935, the CWB was given monopoly control over Canadian wheat in 1943. The rest of Canada's agriculture tends to be governed by what is called "supply management" (Schmitz, Furtans, and Baylis 2002). Products such as eggs, milk, and poultry are influenced by programs that limit imports with quotas and tariffs as well as domestic production quotas.

Agricultural programs in both countries are substantial in size and appear to be politically quite stable. In 2000, more than $29 billion in direct payments were made to farmers in the United States, including $11.6 billion for feed grains, $5.4 billion for wheat, $4 billion for cotton, and $1.5 billion for disaster assistance.[32] Nearly 34 million acres were enrolled in the Conservation Reserve Program, which pays farmers to keep formerly arable land out of production. This is roughly 7 percent of the 430 million acres of cropland (see table 1.1). According to the U.S. Census of Agriculture, net farm income was $46 billion (income = $196, expenses = $150) in 1997, which makes government payments the equivalent of more than half of net farm income. Programs in Canada tend to be less generous to producers than in the United States, yet they are still substantial.[33]

Among agricultural economists, the effects of government policies have been intensively studied with several areas of focus. One important area in this literature has been how farm policies affect crop production, including such issues as crop choice, acreage in production, farm prices, and crop yield. A second focus has been to estimate the deadweight losses of farm programs and the distribution of their benefits and costs. A third focus has been to study the political economy of farm programs in order to explain their form, their survival, and their variation across crops, regions and time. Despite this expansive literature, almost no analysis of the effects of farm programs on contracts and organization in agricultural production exists. For example, in the texts we cite here, there is almost no mention of such possible effects, with the exception of some discussion of the effects of limitations on

farm payments (for example, Knutson, Penn, and Flinchbaugh 1998, 265–66; Pasour 1991, 140). This omission is, perhaps, not too surprising given the relatively limited attention of agricultural economists to contracts and organization.

In this study we focus on incentives arising from the transaction costs created when farmers, landowners, and other asset owners come together to produce agricultural output. Do these farm policies influence these incentives and thus influence the choices farmers make regarding contracts and organization? Generally we would say no: Farm programs do not differentially affect these incentives and thus do not substantively alter economic organization of farms. For example, there is no discernable gain from choosing a cash or a share land lease arrangement in order to increase the benefits derived from government policies.[34] Even if there were predictions, our large data sets all contain cross-section observations that are typically not suited to test potential prediction about the effects of farm policies. Because our cross section data are comprised of observation of farms producing program crops, we cannot capture changes in farm policy. Instead, we would need a panel that covered changes in regimes (for example, before and after the 1996 Farm Bill; before and after changes in crop insurance that might influence risk and thus alter benefits of cropsharing). For example, all of our data come from before passage of the 1996 Farm Bill, so we are unable to test for effects of this new regime.

Our treatment of government policy does not mean that government has no impact on contracts and organization, but that for the issues we examine, the implications are not readily forthcoming. Indeed, if one can discern how such programs have an impact on the incentives we find important, we would expect government to have predictable effects. At several points in the book, we consider some possible impacts of farm policies on contracts and organization. In the chapter 4 summary, we discuss how government programs might influence the choice between a cash contract or cropshare contract. In chapter 8, we consider how tax policies and subsidized credit influence the decision to own or contract for control of assets. And in chapter 9, we consider how limits on government payments and taxes influence farm ownership and organization.[35]

1.4 Other Literature

The agricultural economics literature has historically shown considerable interest in contracts and organization, although the focus has generally been different from what we examine in this book. In the first half of the twentieth century, many economists explored the differences between agriculture and other industries where large-scale production and corporate organization was dominant.[36] Agricultural economists have examined the structure of farm organization and farmland leasing, and those working at the various agricultural

experiment stations at land grant universities have collected data on such issues.[37] Most of this work either predated Coase's work on the firm and social cost or was conducted without explicit recognition of it. Quite often the work is based on transaction cost ideas or has a subtext of costly information, but these issues are never directly mentioned.

Among modern agricultural economists who study North American agriculture, the focus has not generally been on contracts and economic organization, but on neoclassical analyses of costs, production, commodity markets, and the effects of agricultural policies. In the past decade, however, this began to change. For example, there have been analyses of contracts in poultry (Knoeber and Thurman 1994; Tsoulouhas and Vukina 1999), vegetables (Hueth and Ligon 1999), and vertical coordination (Frank and Henderson 1992; Hennessy 1996). This and related literatures are carefully summarized in Knoeber (2000); other less extensive summaries are found in Sexton and Lavoie (2001), Deininger and Feder (2001), and Vercammen and Schmitz (2001). To the extent that modern agricultural economists have examined issues of economic organization, they have relied more on risk-sharing arguments than on transaction costs. As a result this literature is quite distinct from the analysis in this book.[38]

In addition to agricultural economics, there is a considerable literature on farm contracts and organization in economic development and in economic history. In economic development, cropsharing has been a focus of analysis ever since Stiglitz's (1974) paper. This literature, like that in agricultural economics, has been dominated by risk-sharing models and an emphasis on theory over empirical work.[39] In recent years, however, there has been less emphasis on risk sharing and more discussion of multiple incentive margins (for example, Dubois 2002). In economic history, of course, agricultural topics are a mainstay, and discussions of contracts and farm organization are common. Transaction cost models have been much more prevalent in economic history than in other fields, perhaps because of the influence of Nobel Laureates Robert Fogel and Douglass North. Historical issues of slavery and serfdom obviously suggest the importance of property rights, and thus have led scholars to transaction cost economics.[40]

1.5 Organization of the Book

All substantive chapters (3–9) contain both a theoretical model and empirical analysis. Although there is a natural progression in the chapters, the book is analytically divided into three parts. Part I examines contract choice using the transaction cost paradigm, focusing on explaining the prevailing simplicity of contracts and the choice between cropshare and cash rent agreements. We show how land leases exist in a context where reputation and the common law are important. These factors allow farming contracts to be relatively

simple arrangements even though the values of the transactions may be quite large. We also show that the ability of farmers to exploit soil, underreport output, and overreport inputs best explains choices between cash rent and cropshare and explains the details of cropshare contracts. Part II examines the implications of risk-sharing. Our data refute the common prediction of risk-sharing models—namely, that as the risk to farmers increases the incentives to farmers decreases. We also show that there is no evidence for the existence of ratchet contracts in our data, another prediction that is often found in contract theories based on risk. Finally, part III examines ownership and firm organization choices, such as why family farms have dominated farming and what determines the ownership pattern of assets. Here we return to our transaction cost framework for explanations. We generally find that assets are owned when gains from labor specialization are low and timeliness costs are high. We also find that family farming tends to dominate corporate farming for the same reasons. For those readers only interested in the transaction cost approach to farm contracts and organization, part II can be skipped without any loss of continuity.