1991 VALUATION ACTUARY
SYMPOSIUM PROCEEDINGS

SESSION 3

CMO Boot Camp: In the Tranches

Randall L. Boushek

David A. Hall

CMO BOOT CAMP: IN THE TRANCHES

MR. RANDALL L. BOUSHEK: By way of introduction, I am an actuary in a somewhat nontraditional role as an investment officer and bond portfolio manager for Lutheran Brotherhood, a $9 billion fraternal benefit society headquartered in Minneapolis. My distinguished colleague on the panel is David Hall. Dave is also an actuary in an investment capacity. He is vice president and director of portfolio management for the Hartford Life Insurance Companies.

Our subject for this session is "Collateralized Mortgage Obligations (CMOs): What are they, why are they, and what makes them tick?"

Our agenda for this session will include alternating presentations from Dave and myself. I'll begin by giving you a brief introduction to the subject of CMOs. After that, Dave's going to take over and spend some time talking to you about the prepayment function, which is the engine that drives the mortgage-backed market, and is the primary source of risk in that market. After he's through, I'm going to spend a little time talking about basic CMO design and structure. Dave will then review more complex CMO structures, and I'll finish off with a discussion of what are known as "whole-loan CMOs." Our objective for this session is not to turn you all into CMO gurus. Our objective, rather, is to provide a basic level of understanding.

For this session, we are assuming a zero level of prior knowledge. I realize that for some of you this may be a bit elementary. We ask that you'll bear with us. We're trying to lay a groundwork for everyone.

I'd like to make one preliminary comment on terminology. We will be using two terms quite frequently: Mortgage-Backed Securities (MBSs), and CMOs. When we use the term MBS, we're talking about the broad class of investments that, at the core, are supported by residential mortgage loans. CMOs are a subset of the MBS universe.

I feel that the topic we're going to be discussing is extremely relevant. In the last few years we've seen a tremendous explosion in the volume and the use of CMOs, not only among insurance companies, but also through several financial institutions.

The CMO market was born in 1983. Both Lutheran Brotherhood and Hartford have been active in the market since its inception. At my company, we have approximately 35% of our total assets in MBSs, principally CMOs, including over 200 different CMO tranches. This year, we will commit about one half of our investment cash flow to the purchase of CMOs. I believe at Hartford the percentages are both somewhat larger.

The insurance industry itself has become a very active user of MBSs. Chart 1 shows the distributions of net new investments by the life insurance industry as an aggregate in each of the past five years, based on data from the ACLI.

In 1987 and 1988, MBSs, including CMOs but not exclusively CMOs, accounted for about 14% of net investments. Based on results through the first three quarters of 1991, it appears that they will account for about 27% of net investments in 1991. That's a doubling of the role of MBSs in the last three to four years. In fact it appears that, for 1991, MBSs will represent the single largest category of new investments for life insurance companies. Some of that growth has come at the expense of the direct mortgage origination market. But it has also come as other categories have declined as well.

The first question that I'm often asked by actuaries, accountants and others is, "Where do CMOs come from?"

In the beginning (see Chart 2), there are homeowners who need to take out mortgages to pay for their homes. These mortgages are provided, variously, by banks, thrifts, mortgage bankers, and even by insurance companies that do direct origination.

CHART 1

CMO Boot Camp
New Investments -- All Life Companies

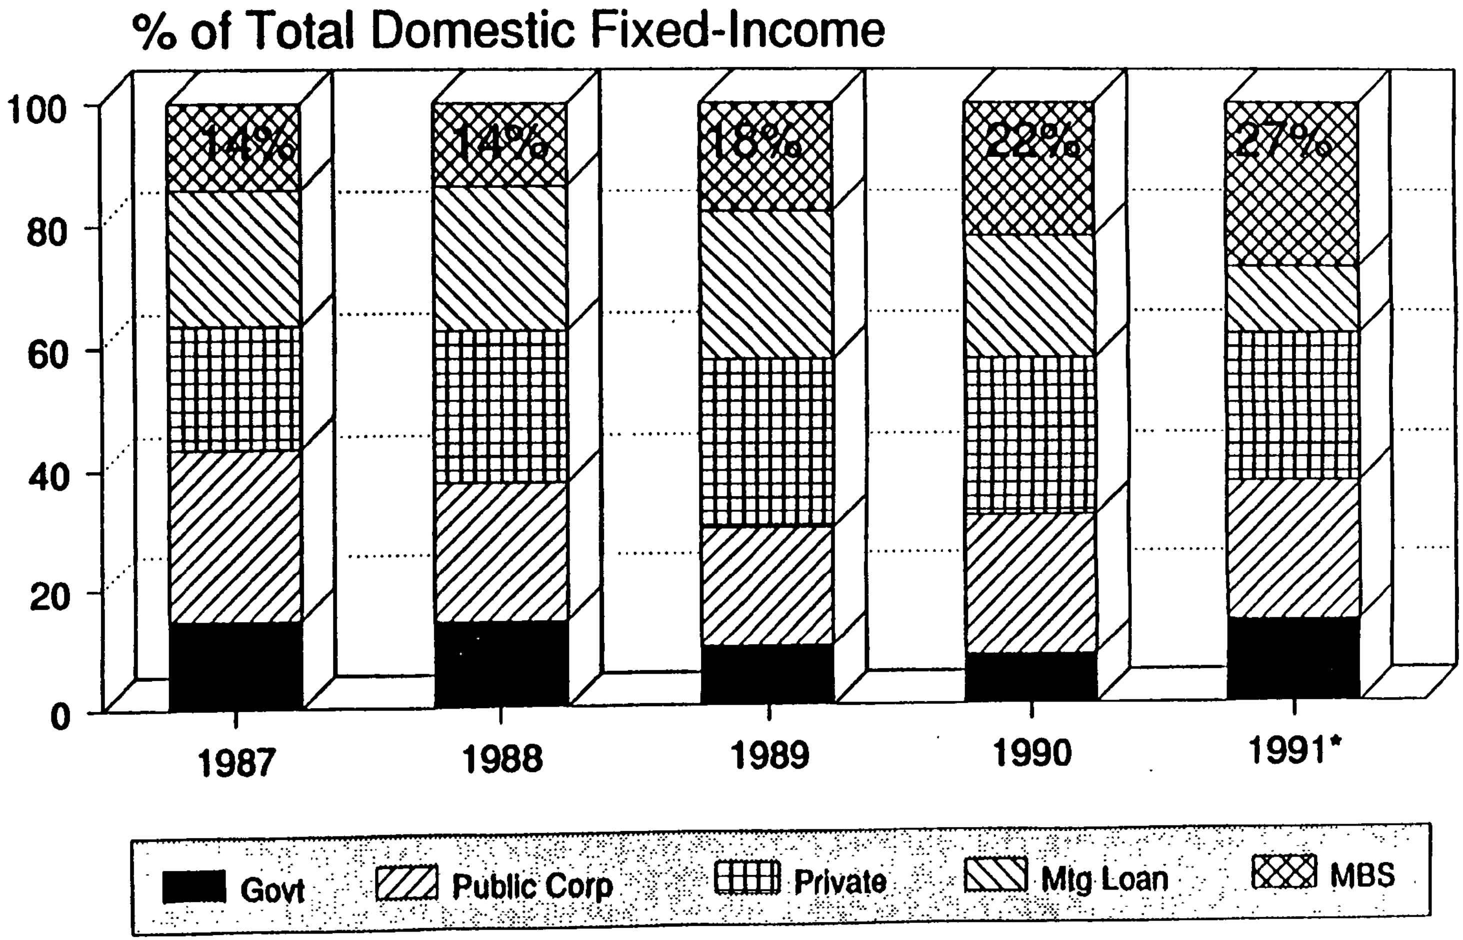

CHART 2

Evolution of a CMO

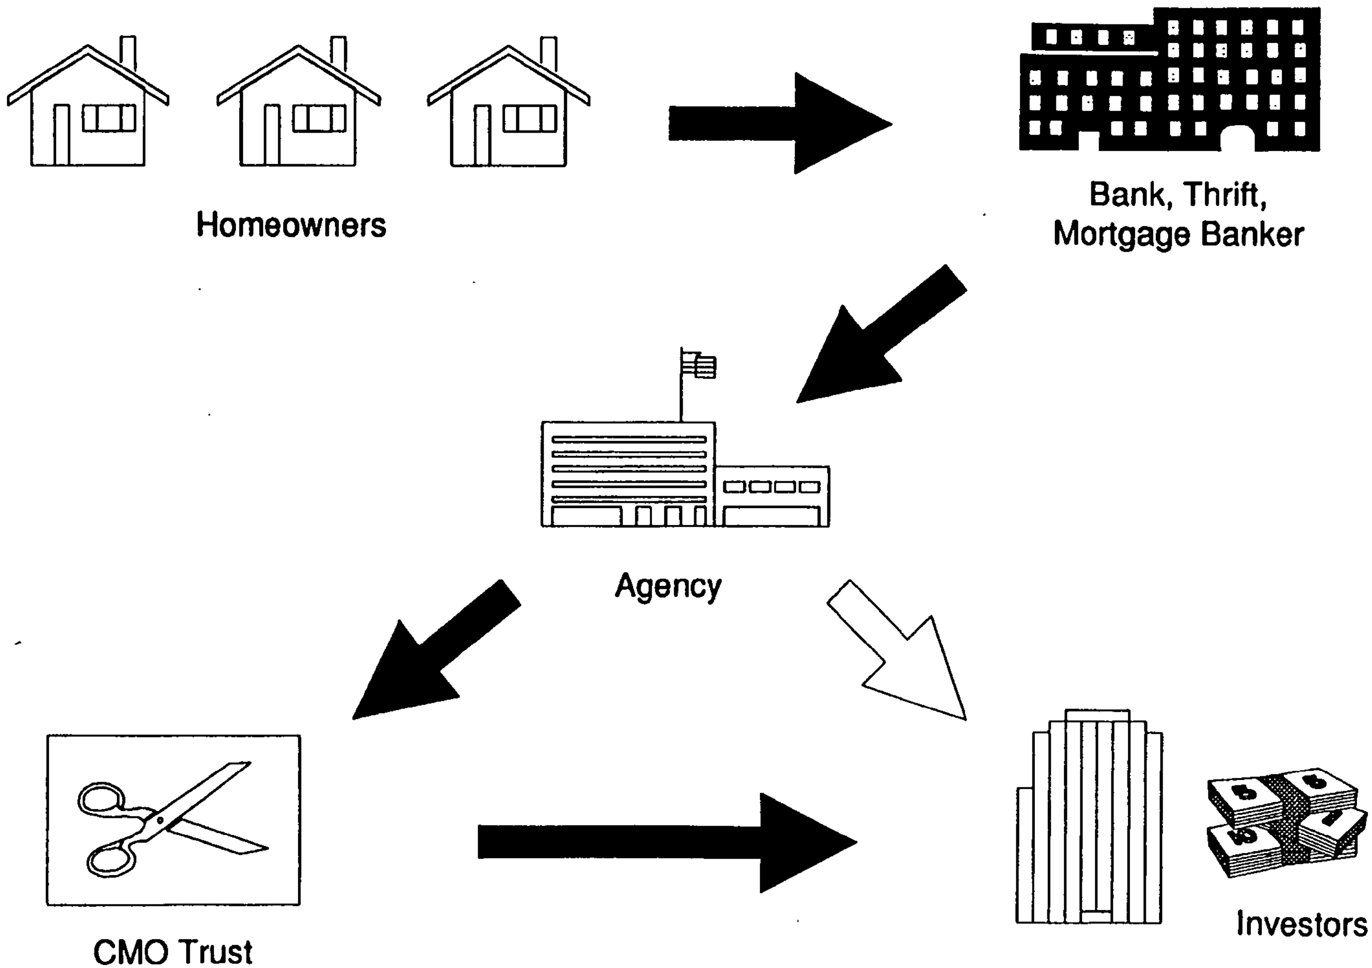

These originators retain some of these mortgage loans. Many, however, are sold as packages of loans to one of three governmental or quasi-governmental agencies: the Government National Mortgage Association (GNMA, or "Ginnie Mae"); the Federal Home Loan Mortgage Corporation (FHLMC or "Freddie Mac"); and the Federal National Mortgage Association (FNMA or "Fannie Mae"). These three agencies exist specifically to provide a secondary market to the mortgage originators, to buy packages of loans from them in order to inject additional funds into the mortgage financing market.

The agencies take packages of loans acquired from different institutions and combine those with similar characteristics (principally maturity and interest rate) into "pools." Each agency has its own name for the pools. We will use the generic term *pass-through securities*.

Each agency has its own parameters for what can go into a pool, and each buys different types of mortgages. GNMA, for example, buys Veterans Administration (VA) and Federal Housing Association (FHA) mortgages. FNMA and FHLMC buy conventional, conforming mortgages. Some people feel that there is a slight difference, also, in the credit-worthiness of these three agencies. GNMA obligations are a direct obligation of the United States government. FHLMC and FNMA have only implied guarantees from the government. These are actually privately owned corporations with direct borrowing rights at the Treasury.

Investors who buy shares in a pass-through own a proportional and equal share of each cash flow generated by that pool. One of the principal drawbacks of pass-through securities is that they pay down continuously for long periods of time. Unfortunately, there are some investors who do not want cash flows early on, and others who do not want cash flows later on. The CMO market was developed, in part, to try to satisfy a number of these different investors and make the mortgage market more palatable to a broader range of participants.

In a CMO, we interject one more element into the diagram (Chart 2), that being a special-purpose corporation or trust. After the agencies (GNMA, FHLMC, and FNMA) buy packages of loans and put them into pools, the trust buys a number of similar pools and combines them into a melting pot. Then the trust slices and dices the combined cash flows in any number of ways for distribution to investors. When you buy an ownership interest in a CMO, you are not buying an equal and proportionate share in every cash flow that the underlying mortgages produce. Rather, there are patterns and priorities, depending upon the particular tranche that you buy.

There are two Congressional Acts that have a lot to do with the growth of the CMO market. The first of these is the Secondary Mortgage Market Enhancement Act, or SMMEA, which was enacted in 1984. Basically, that Act overrode statutory limitations that had prevented a number of insurance companies from owning CMOs. Under SMMEA, the only investment limitations that can apply to agency or high-quality mortgage securities are the same limitations that apply to Treasuries.

Interestingly, SMMEA has come up recently in a number of news articles. There was a provision in SMMEA that states could explicitly override it, prior to October of 1991. A few states have done so, most recently and notably New York, and legislation has been pending in a few others. Effectively, however, those overrides have had little practical impact. For example, New York legislation (at least in its preliminary form) will allow an insurance company to hold up to 80% of assets in high-grade MBSs.

The Tax Reform Act of 1986 was perhaps even more significant than SMMEA. This Act created the Real Estate Mortgage Investment Conduit (REMIC). It basically simplified the tax and accounting issues involved with MBSs, CMOs in particular, and led to a boom in the market.

Chart 3 shows the growth of the CMO market, beginning with 1984. You can see the jump in 1986 (following the REMIC legislation), the constant and continual growth up through 1990, and the literal explosion of growth in 1991. We will be approaching $175 billion of new

CHART 3

CMO Boot Camp
Issuance of Agency CMOs

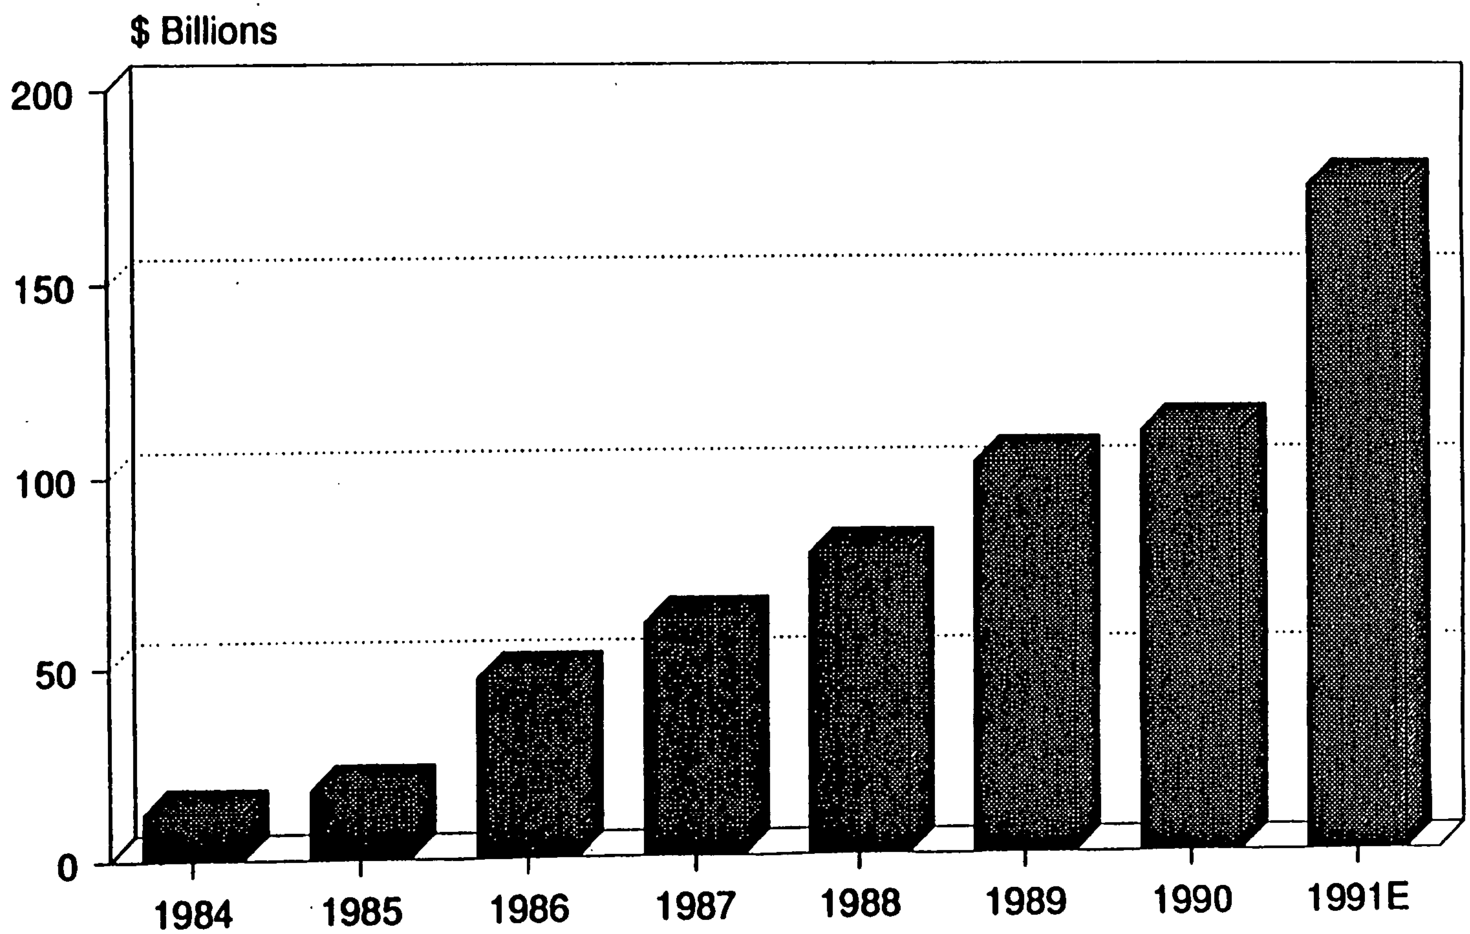

CMOs created in 1991, owing in large part to the tremendous increase in demand from both insurance companies and banks for AAA, high quality paper.

Chart 4 contrasts the growth in the CMO market against the overall issuance of mortgage-backed pass-throughs. Notice, as we move forward from 1984, that CMOs are consuming more and more of the fixed-rate pass-through issuance and outstanding issues. For certain coupons, the situation has become quite extreme. As an example, almost 90% of all outstanding FNMA 9%'s are now in CMO structures.

MR. DAVID A. HALL: This section would be more appropriately titled, "Modeling the Prepayment Function," which could be called "The Impossible Dream." As you'll see in my conclusion, it really can't be done.

There are five principal sources of mortgage prepayments:

1. Sale (net of assumability) -- Clearly, the most natural one is the sale of a home. The caveat is that some mortgages are assumable, and therefore the sale of a home might not cause prepayment of a mortgage.
2. Default -- A mortgage default would generate a prepayment.
3. Casualty -- Casualty of the house: If it gets swallowed by an earthquake, or burns down, the mortgage would generally be prepaid.
4. Accelerated payments -- These are a form of prepayment, sometimes called a partial prepayment. This is the phenomenon of a mortgage holder paying down his mortgage a little bit faster than his schedule calls for. There have been a lot of articles as to the power of compound interest, describing how paying a little more in the beginning dramatically reduces the number of payments which are ultimately required.
5. Refinancing -- This is probably the most problematic source.

CHART 4

CMO Boot Camp
Comparative Issuance

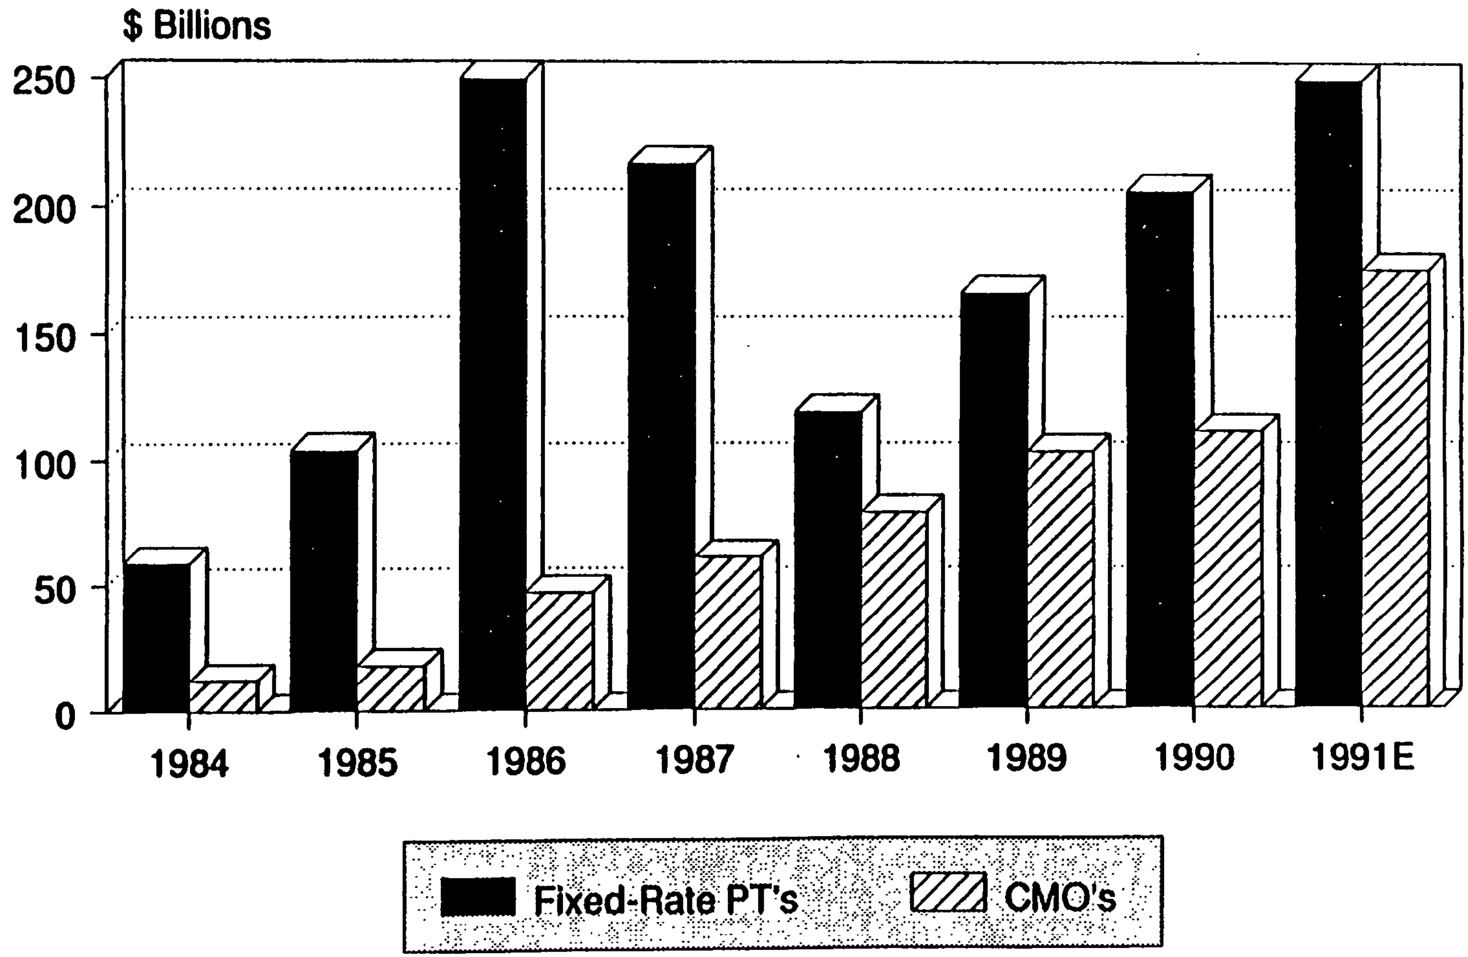

Any of these sources may be noneconomic from the point of reference of the end investor. That is, just because a mortgage is prepaid does not necessarily mean it is to the detriment of the investor to have that mortgage prepaid.

Having said that, I do feel that in most instances, any decision to prepay a mortgage by the mortgage holder is economic from his frame of reference. And it's important to understand that the economics of the decision from the borrower's point of view are not always contrary to the economics of the prepayment from the investor's point of view.

Here are some of the factors that you may want to consider in constructing a prepayment model, which helps explain why constructing a reasonably realistic model could be so difficult.

1. Seasoning -- This refers to the aging of a mortgage. There's very clearly an observable select-and-ultimate type of period. Somebody who takes out a new mortgage is less likely to prepay it in the early months than he is after some period of time.
2. Seasonality -- This refers to the differences in prepayments from winter to summer. I believe September tends to be the peak month of prepayment activity, and March tends to be the low one. I'm not quite sure why that is, but it's a very clearly observable trend. It may have something to do with the school year and cold weather.
3. Interest rates -- Interest rates and changing interest rates are probably the most obvious determinants of prepayment behavior. And that relates not only to the refinancing potential, but also to the general affordability of housing.
4. Selection -- This refers to the wide variety of mortgage products that are available to consumers today, which allow the borrower who expects to be in his home for a shorter period of time to select a mortgage option at a lower interest rate. He could choose an adjustable rate mortgage (ARM), or a mortgage with a balloon payment over some shorter period of time. He would expect one of these to reduce his ultimate payments.

   Therefore, the borrower of a traditional 30-year, fixed-rate mortgage is probably a little bit different today, on average, than he might have been 10 years ago when all these alternative ways to finance housing were not around. Many of the faster prepaying mortgage holders would not choose to take a 30-year fixed-rate mortgage in today's world.

5. Burnout -- This is a term that's commonly used to reflect the tendency of a pool of mortgages that have been exposed to a persistent refinancing opportunity to cease to have its prepayments dependent on mortgage rates. Let me give an example.

   I read the other day that there are still over $1 billion of GNMA mortgages outstanding, in which the homeowner is paying a fixed rate of 14.5% or higher. Now, these remaining mortgages are very clearly not interest-rate sensitive in the prepayment decision going forward. There seems to be no rational explanation of how they could persist.

And so, having said earlier that all prepayment decisions are economic from the perspective of the mortgage holder, my caveat would be that I'm not sure that all nonprepayment decisions are necessarily economic.

The last three factors affecting prepayments are more global macrofeatures:

1. Geography -- Warmer climates tend to experience different patterns than northern climates.
2. Economy -- General economic activity clearly affects housing turnover.
3. Demographics -- Even this may play a role. For example, as the bubble of baby boomers has come through the decade of the 1980s, there may have been a surge of people going from apartment to condo to starter home to larger, family-sized housing. That may not be symptomatic of what's likely to be happening in the next decade.

So there are all those factors that you really have to consider in developing an understanding of what causes prepayments.

One of the common older prepayment models, or prepayment experience studies, are the FHA experience studies. Those studies were based upon FHA mortgage loans. Two different prepayment tables are frequently cited (see Chart 5). One covers the period from 1957 to 1981, and the second covers the period from 1970 to 1984. I won't go into all the pros and cons of why this model is or is not appropriate. You can read that in the session we did in the 1990 Proceedings of the Valuation Actuary Symposium.

These tables do clearly indicate the aging phenomenon, that is, this tendency for mortgages to prepay more slowly in the early years than in the later years.

Incidentally, I'm not sure precisely what the relevance of the sharp upward spike towards the end is. I'm sure there's a lack of data there, but it may be that as mortgage balances get to be so small, there's a nuisance value associated with maintaining them, and they tend to pay off quickly.

This is clearly not a very convenient type of model to use to express prepayment behavior. It's not a very linear function. On the other hand, the prepayment functions in the first 20 years of those two models do resemble a piecewise linear, ramp-and-platform function which starts at about zero and grades up to 6% in about two and a half years, and then stays level at 6% thereafter.

This ramp-and-platform function has been chosen as a standard prepayment model by the Public Securities Association (PSA), which is a standard-setting association that deals with settlement processes and quoting processes for publicly traded securities.

A more specific description of the PSA model is as follows. First, I have some terminology: CPR stands for constant prepayment rate, or conditional prepayment rate. This is an annualized rate of decrement, which means that, if the CPR was 6% throughout a year, then at the end of the year the number of mortgages which are still outstanding is 6% lower than it was at the

CHART 5

Prepayment Models
FHA Experience

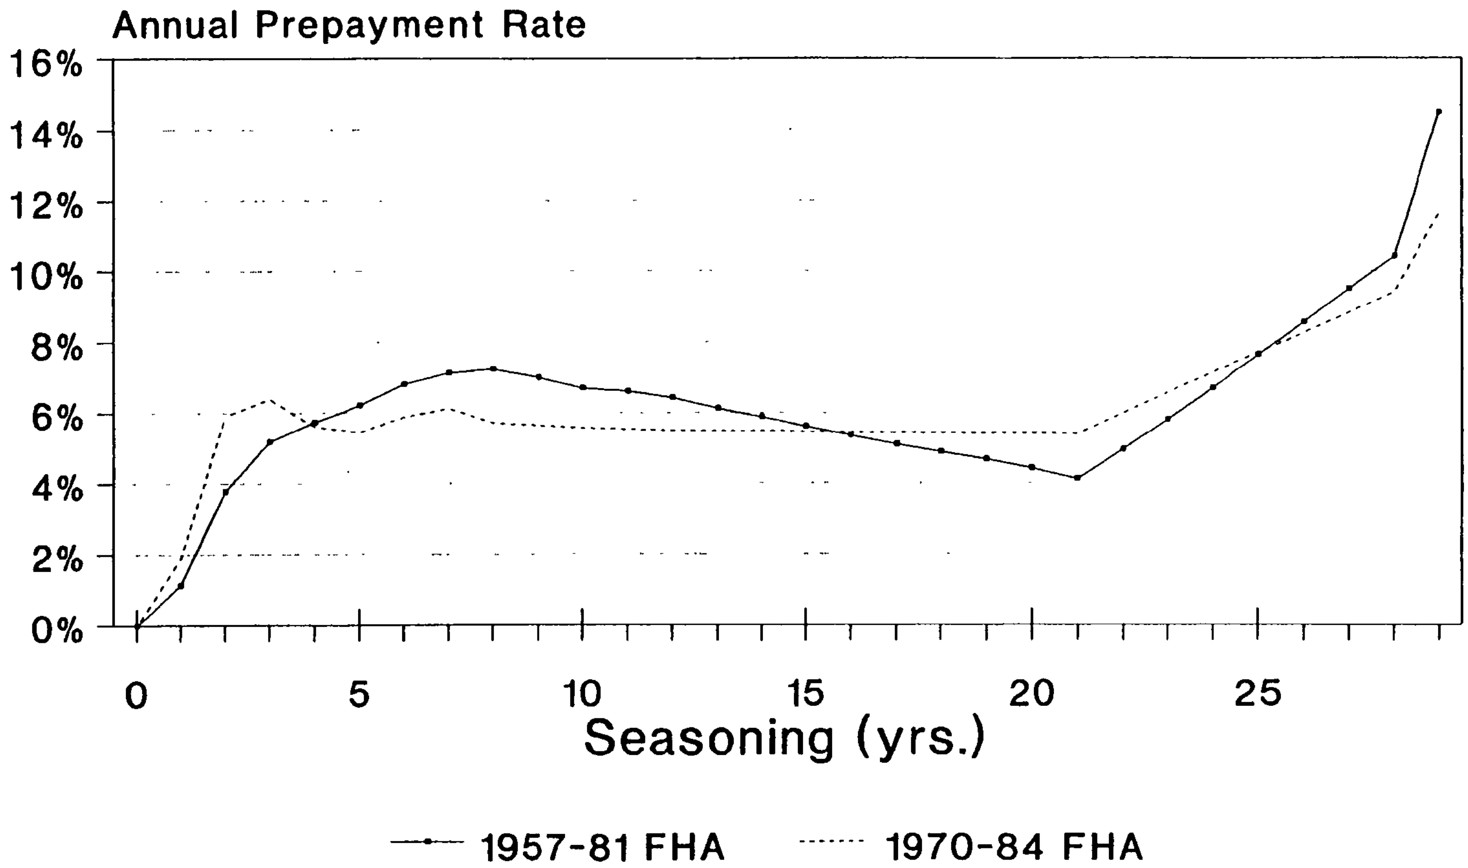

Source: US Dept. of HUD

beginning of the year. The PSA prepayment model grades up linearly from 0.2% CPR in month one to 6% CPR in month 30, and remains level at 6% CPR thereafter. This piecewise linear ramp-and-platform function seems to approximate the FHA tables. A graph of this PSA ramp-and-platform function is shown in Chart 6.

This standard PSA prepayment model is fairly important to understand because virtually all CMOs have their principal repayment characteristics stated in terms of the PSA model. Prepayment projections are commonly quoted based upon a multiple of this model. An example would be 150% PSA. This would be a prepayment table that grades from an annualized rate of 0.3% in month one up to an annualized rate of 9% in month 30, and then stays level at 9% thereafter.

Although the 100% of PSA model seems to be a fairly decent fit for the FHA experience, that does not mean that 100% PSA should be your baseline model. Very clearly, it has not been the baseline experience in recent time, and it probably will continue not to be. But the PSA model is what is commonly used, and it's important to understand it.

So, with all this information, how can we model prepayments? There are a great many factors to consider. Probably some sort of regression analysis would be the way to do it, looking at how all of these economic and demographic factors interplay. In fact, virtually all of the major Wall Street broker/dealers who are active in the CMO market have put their rocket scientists on this problem, to develop as complete a mortgage prepayment model as possible. So, given all of this brainpower and computing horsepower that Wall Street dedicates to this problem, you would expect that they would at least arrive at some sort of Wall Street consensus as to prepayment.

Chart 7 indicates that that's not true. It shows median, upper, and lower prepayment assumptions as a function of mortgage coupon, for all the major Wall Street broker/dealers. Clearly, there's a wide range of projections.

CHART 6

Prepayment Models
FHA Experience vs PSA Model

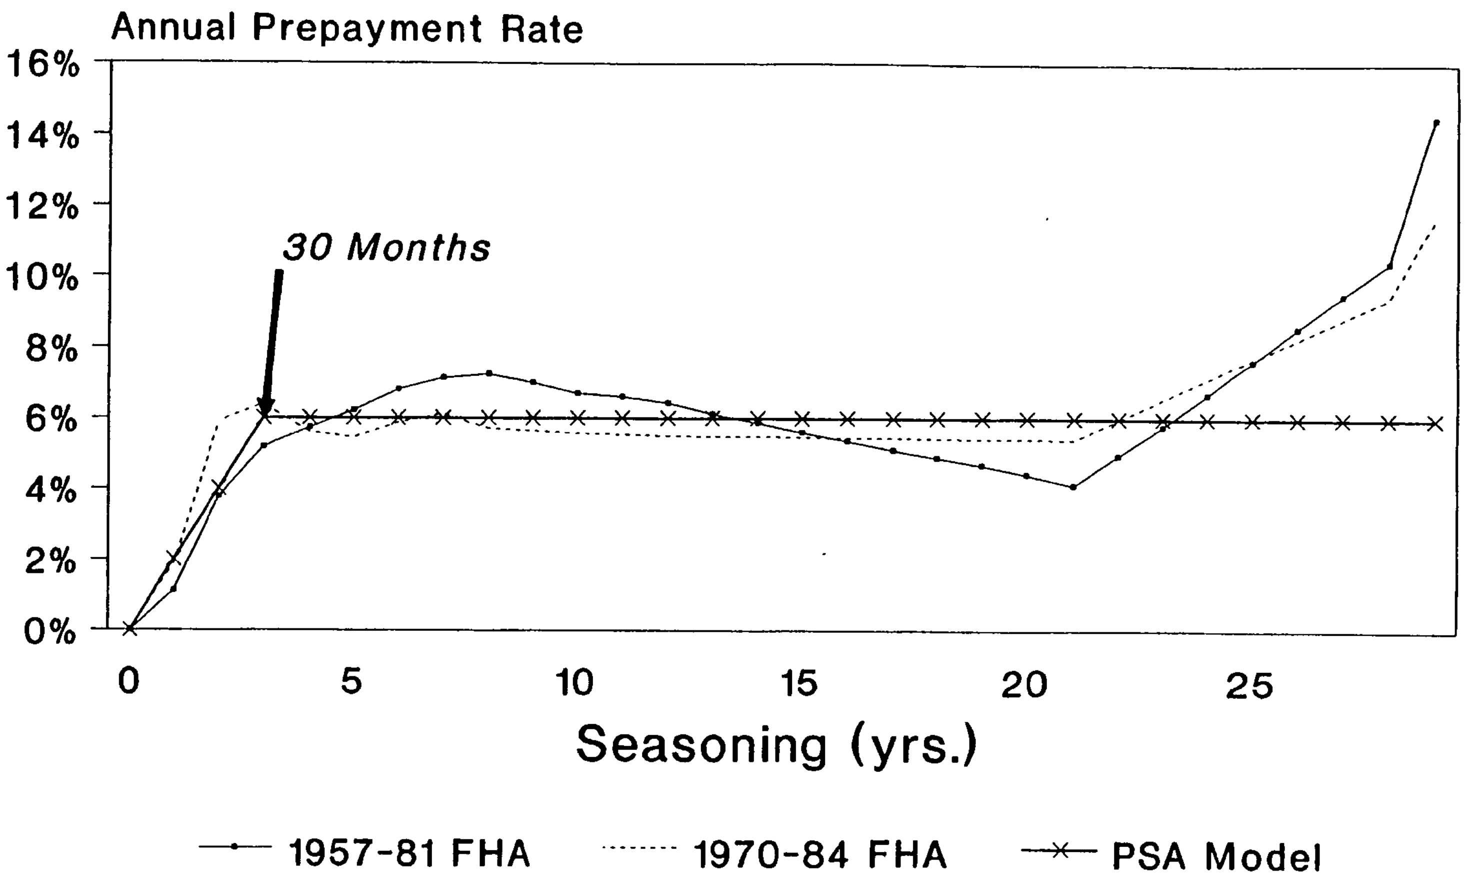

CHART 7

Dealers' Prepayment Estimates
FNMA/FHLMC as of 10/15/91

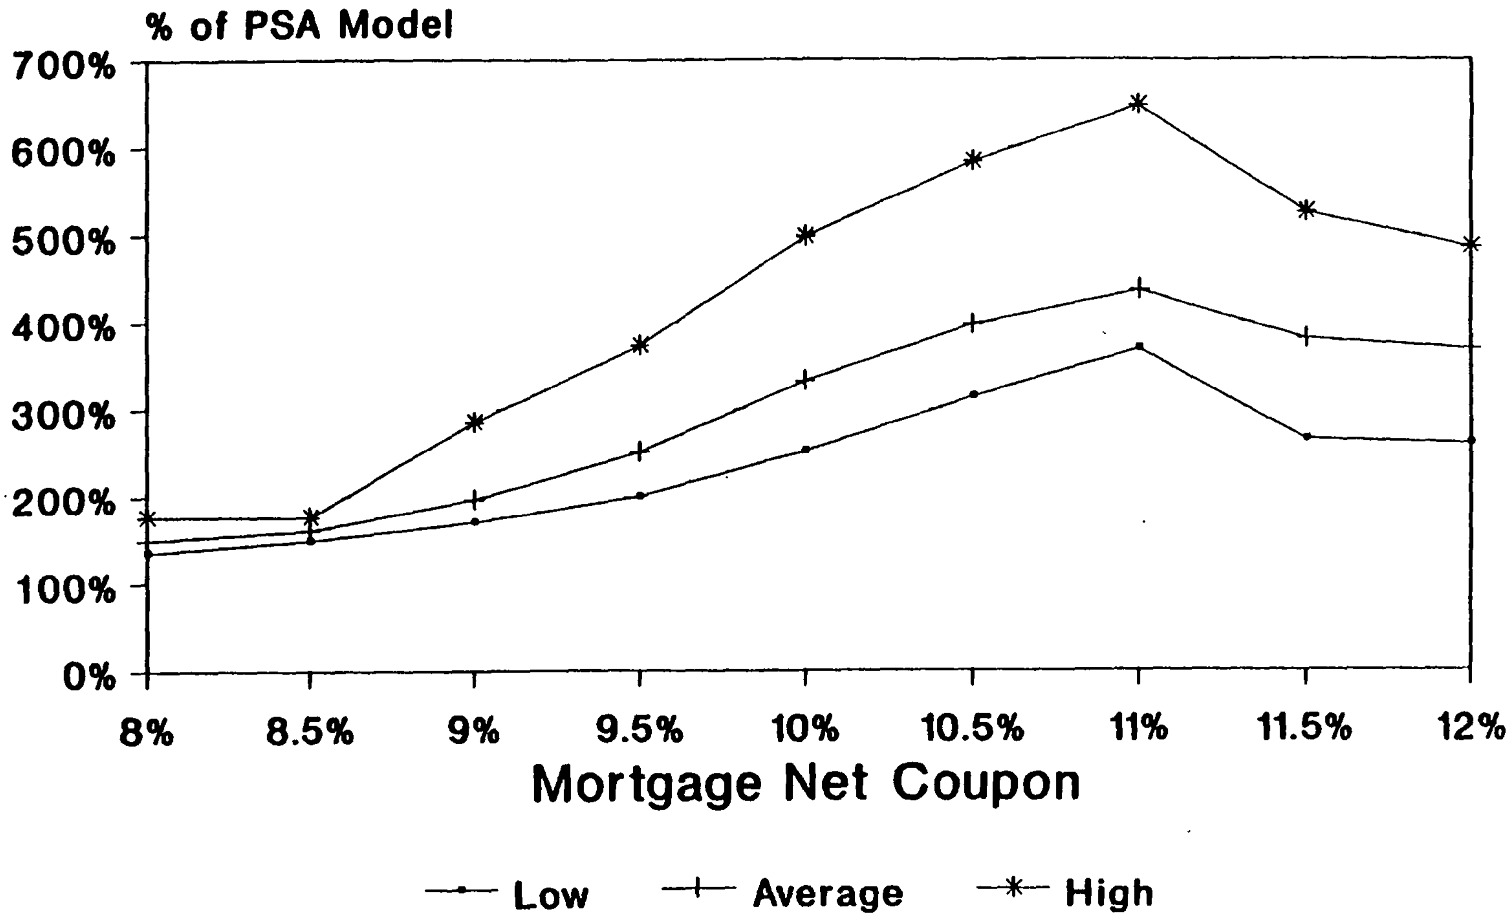

Source: Bloomberg Financial Markets

These are long-term projected prepayment speeds, on the assumption that interest rates remain constant at their initial level. For example, if we go to the middle of the table, for the 10% mortgage coupon, the average expectation would be 325% of the PSA model rate, over their life. However, there's at least one firm that is somewhere up in the 500% PSA range, roughly 50% higher. Similarly, there's at least one firm that's considerably slower; and in fact there's a wide range in between. These are not two extreme outliers; there's quite a wide range.

My point in showing this is that there is a wide variety of opinions as to what causes mortgage behaviors; and in spite of all the modeling effort that we do to try to understand that, we clearly haven't got it right.

This leads me to my final comment for this portion of the session, which is that models in general are based upon historical experience.

However, the world changes, and it changes very rapidly. So any model that is based on experience over the last five years is probably relevant to only those five years. And, very clearly, those five years are not accurate predictors of what is likely to be the case in the next five years. Consequently, although actuaries don't like to hear this, in order to be involved and comfortable with the CMO market, you have to get comfortable with prepayment uncertainty. You're not going to be able to model prepayments accurately, because nobody can. So you have to get comfortable with understanding what your range of outcomes is, given that you are very uncertain as to what the prepayment experience will be.

MR. BOUSHEK: A point that Dave just made brings up something we're going to spend a lot of time talking about at Session 8; namely, why static modeling does not work at all for MBSs, and why multiple-scenario modeling is absolutely essential.

Another term we need to define for you is weighted average coupon (WAC). The "gross WAC" is the weighted average interest rate on the loans in an MBS pool or CMO trust. The "net WAC" is the weighted average interest rate that is passed through by the loan servicer. The margin between the net WAC and the gross WAC represents the fee for servicing that's done by the original lenders.

Servicing fees vary by agency. For GNMAs, that fee is consistently 50 basis points. "GNMA 9%" refers to a GNMA pool of mortgages that will pay the investors a net-WAC coupon of 9%, which means that the pool is comprised of mortgages with gross-WAC loan rates of 9.5%. FHLMC and FNMA have more flexible guidelines. The general assumption might be that the difference between the gross and net WAC is 75 basis points, although it could range from 50 to 200 or more. In modeling the collateral, it is important to get a handle on not only the net WAC, but also the gross WAC.

Suppose we take the raw materials, which are the mortgage loans and the pass-throughs created from them, and put those into a CMO trust. There are basically two ways that we can slice the payments (principal and interest) going into the trust. We could slice them "horizontally," severing interest from principal, and create interest only strips (IOs) and principal only strips (POs). Or we could slice them "vertically" and create CMO tranches.

Chart 8 shows a division of the annual, year-by-year cash flows for a pool of 30-year mortgages (FNMA 9%). This example assumes 0% PSA; that is, there are no prepayments. In this case the annual cash flows are the same each year (since nobody is leaving the pool); but the portions of the mortgage payments that go to pay off outstanding principal are increasing over time, and the portions going to pay interest are decreasing over time.

The left-hand graph in Chart 8 shows the split of cash flows between interest and principal, where the dark bars on the bottom are principal and the cross-hatched bars on the top are interest. The right-hand graph shows the aggregate amounts of principal and interest received over those 30 years. Note that there is roughly twice as much interest paid over time as there is principal, based on this assumption (no prepayments). An IO or PO is simply the slicing and

CHART 8

Basic Design: Interest/Principal Stripping
IO/PO from FNMA 9 at 0% PSA

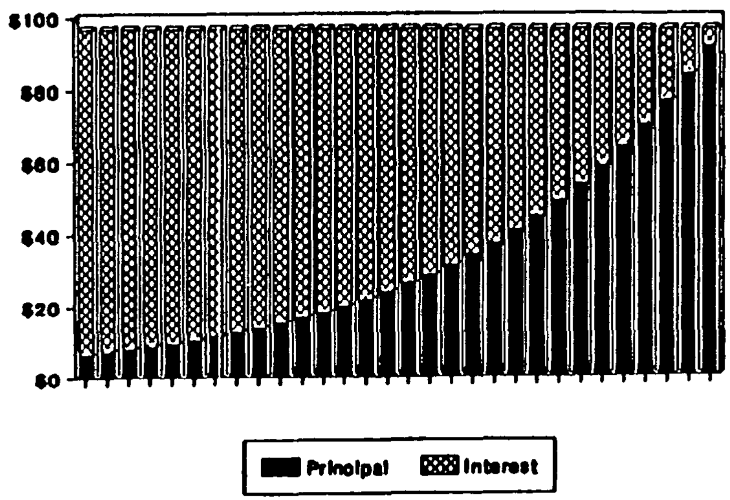

Annual Cash Flow

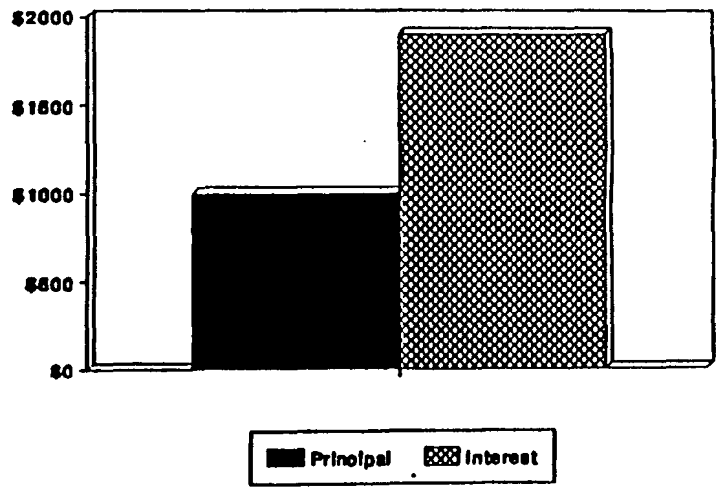

Total Cash Flow

severing of the interest from the principal, buying only the interest part of the cash flows (cross-hatched bars), or buying only the principal part of the cash flows (dark bars).

Now, instead of looking at this mortgage pool at 0% PSA, let's assume 100% PSA. This might be our base case assumption, based on certain projections of interest rates. This is shown on Chart 9. Here we see a dramatically different pattern of year-by-year cash flows, both principal and interest. (Note that these are on different scales from Chart 8.) In the aggregate, you can see that the amounts of principal and interest received are roughly the same. The total principal amount will be the same as in Chart 8, regardless of the PSA assumption. In the case of principal, it is the timing of the receipt that's uncertain. But you can see that the aggregate dollars returned from the IO is radically different. In the case of interest, it is not only the timing, but also the total absolute amount of the dollars received that is different.

Now let's say that we are looking at a scenario that calls for a significant decline in interest rates of up to 200 basis points. In that case, we might look at an assumption of 400% PSA. And you can see from Chart 10 what that does to the cash-flow pattern. We have significantly up-fronted the cash flows on this mortgage security, because of the rapid prepayment assumption. At 400% PSA, the total aggregate amount of interest paid is roughly half as much as the total principal paid.

The other way to slice the cash flows is to look at them vertically, as CMO tranches. The graphs I'm going to be showing you now are a little different. I want to explain them and how they differ from the earlier graphs. In Chart 11, the left-hand side shows only the principal that's being repaid in each year. CMOs are defined primarily as a subdivision of principal payments. They also involve subdivisions of interest, but to begin with I want you to think of them as simply a subdivision of the principal cash flows. On the left-hand graph, you see the principal cash flows of the pool described above, at an assumption of 0% PSA. (These are similar to the dark bars on Chart 8). On the right-hand side, you see the remaining principal outstanding, year by year.

CHART 9

Basic Design: Interest/Principal Stripping
IO/PO from FNMA 9 at 100% PSA

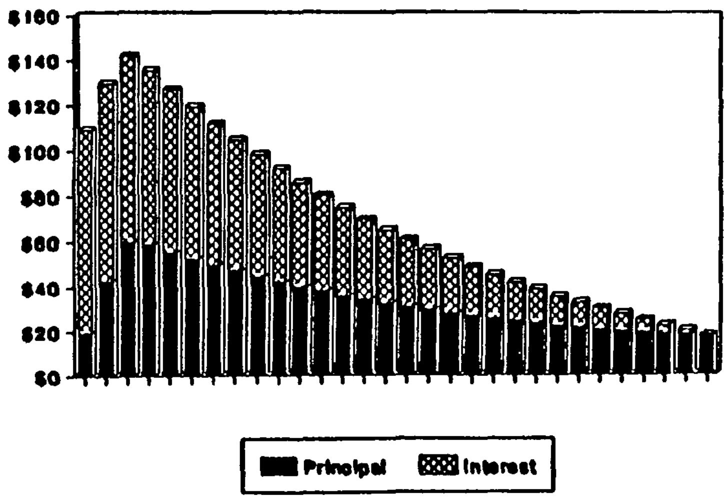

Annual Cash Flow

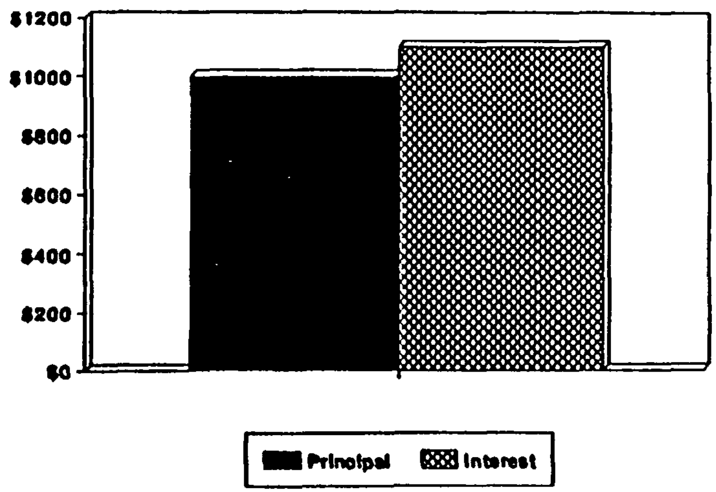

Total Cash Flow

CHART 10

Basic Design: Interest/Principal Stripping
IO/PO from FNMA 9 at 400% PSA

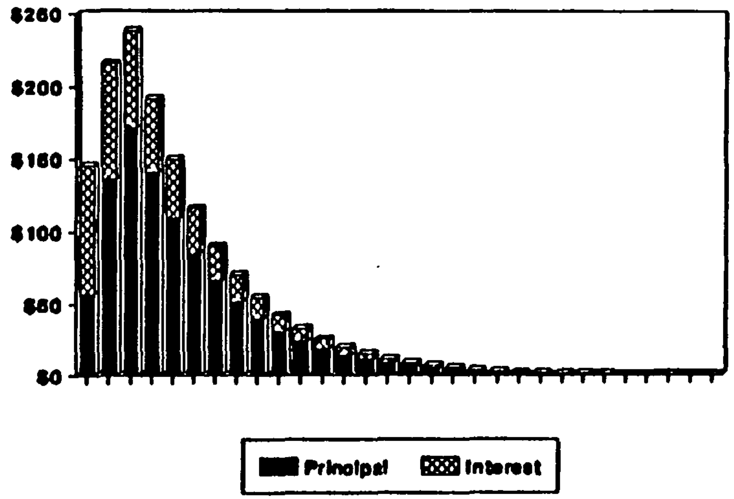

Annual Cash Flow

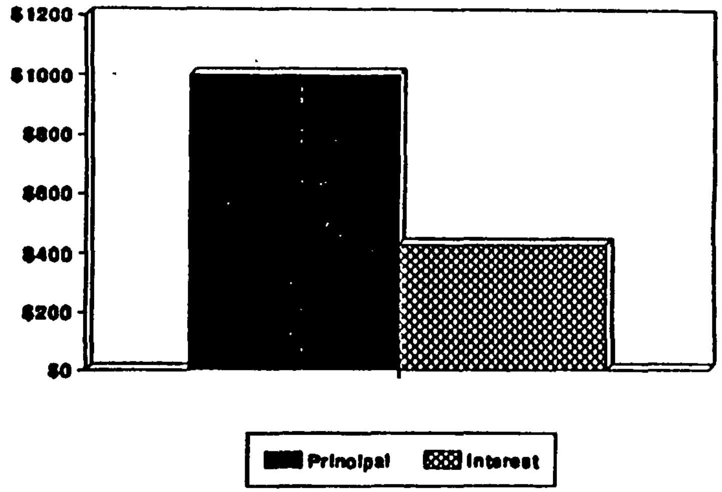

Total Cash Flow

CHART 11

Basic Design: CMO Tranching
Simple 3-Tranche CMO at 0% PSA

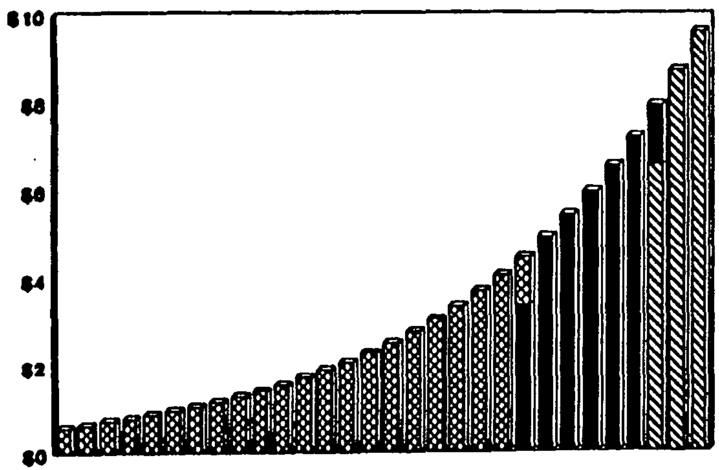

Principal Repayment

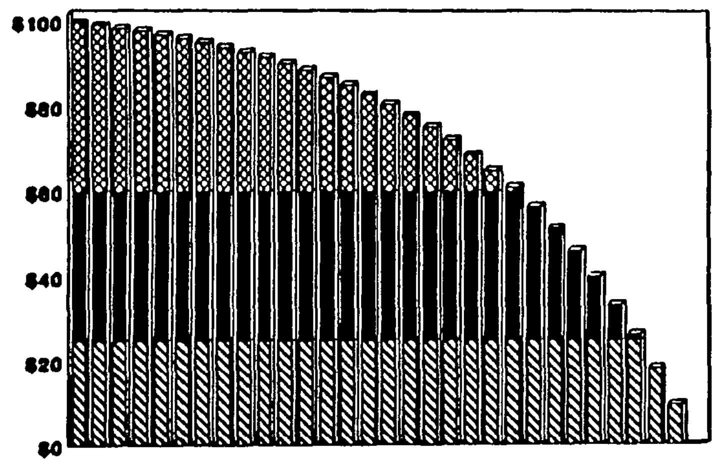

Principal Remaining

The earliest CMOs were subdivided on the basis of "sequential tranches." In a sequential CMO, the first so many dollars of principal repayments go into the first tranche, the next so many dollars of principal into a second, the next so many dollars into a third. In this particular example, the first 40% of the principal comprises tranche one (the cross-hatched bars). The next 35% comprises tranche two (the solid black bars). And the final 25% comprises tranche three (the striped bars).

Suppose that the tranche you are buying is the solid black tranche, tranche two. If the mortgages prepay at 0% PSA, then you will receive only interest until about the 21st or 22nd year, at which time your principal payment cash flows will kick in and continue for about six years. The right-hand graph shows that you will have a level amount of principal outstanding, generating a level amount of interest each year, until your tranche starts paying down in about year 21 or 22.

Now, this is not the way the bond was priced. If we look at this again at 100% PSA, we can see the pattern of principal repayments on the left-hand graph. (See Chart 12, which corresponds to Chart 9). You'll notice how significantly "shorter" your tranche has now become. The solid black bars have moved significantly to the left. Instead of buying an instrument with average principal repayments of about 25 years, you're buying an instrument with average principal repayments of about 10 years. On the right-hand graph, you can see a more rapidly declining principal balance outstanding.

If we take it one step further, and move to a 400% PSA assumption, you can see that your 10-year instrument has now become a four-year tranche. (See Chart 13, which corresponds to Chart 10). On the right-hand side, you can see the rapid amortization of the balance outstanding.

All CMOs are unique. There are no two CMOs that are created and look the same, and there are certainly no two CMOs that behave the same, because there are different pools of mortgages

CHART 12

Basic Design: CMO Tranching
Simple 3-Tranche CMO at 100% PSA

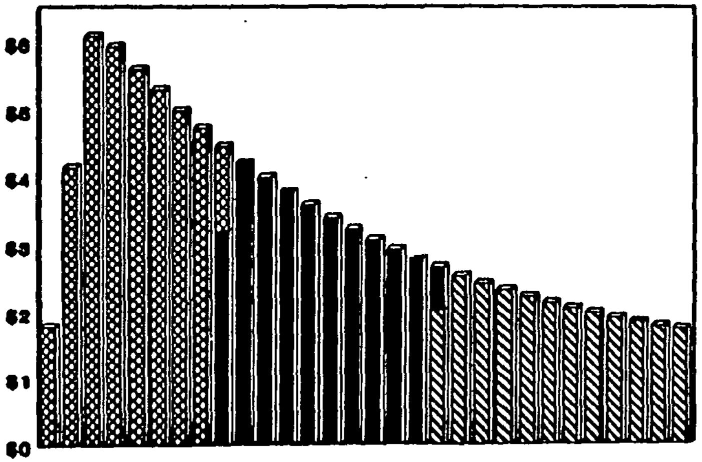

Principal Repayment

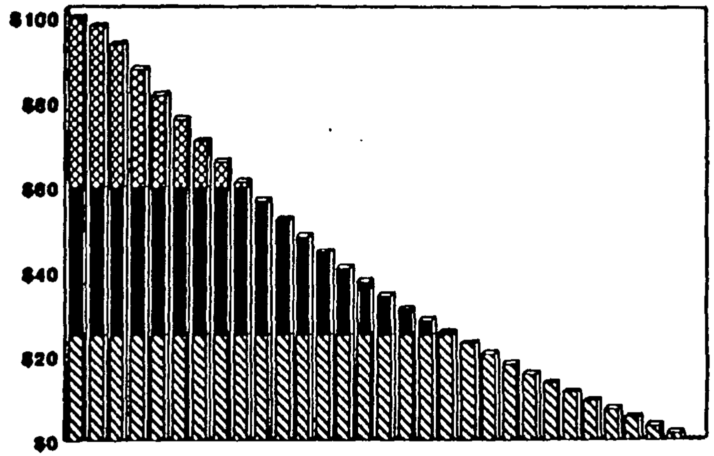

Principal Remaining

CHART 13

Basic Design: CMO Tranching
Simple 3-Tranche CMO at 400% PSA

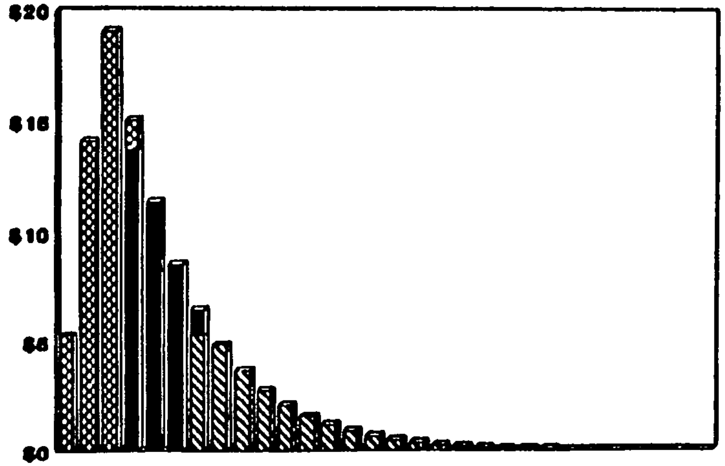

Principal Repayment

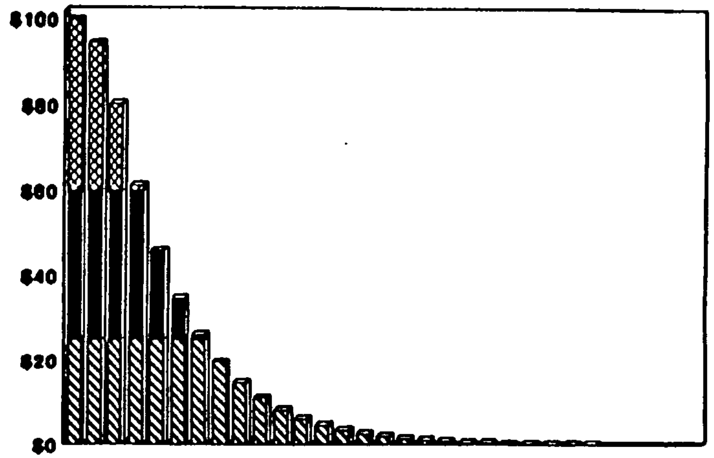

Principal Remaining

underlying each one. There are a number of key characteristics that differentiate CMOs from one another:

1. Packaging -- This comprises the types of loans in the trust. There are GNMAs (I or II), FHLMCs (old or gold), and FNMAs. You will use a different prepayment assumption for each, because they reflect different types of borrowers. FHA/VA loans (GNMAs), for example, pay down differently from conventional conforming loans. There are pass-throughs and CMOs that are based on 30-year loans; there are pass-throughs and CMOs that are based on 15-year loans. The 15-year loans pay differently from the 30-year loans. The loans may be graduated payment loans. There is a class of pass-throughs and CMOs referred to as tiered payment mortgages (TPMs), based on graduated payment mortgages. They have a different prepayment assumption associated with them. And balloon mortgages are a rather new development, but we are now seeing CMOs created out of FNMA five- and seven-year balloons.
2. Coupon -- The coupon makes a difference. You can have two seemingly identical FNMA 9% pools. If one of them has a 75-basis-point servicing fee (and hence a 9.75% gross WAC), that is going to prepay very differently from a FNMA 9% pool that has a 150-basis-point servicing fee (and hence a 10.5% gross WAC). The gross WAC is important in modeling the prepayments for a given CMO. In the earlier days of CMOs there was much more latitude taken in the combinations of coupons and maturities that were placed into a CMO. It is not uncommon to see some of the older CMOs comprised of everything from FNMA 8% pools to GNMA 12% pools. In the last few years, we've seen much more homogeneity in the pools that go into CMO tranches.
3. Seasoning -- This refers to how long the loans have been outstanding. The PSA schedule takes this into account somewhat, but not fully. There are other factors involved. And, more important, if you're looking at a secondary or seasoned CMO, how has it paid down so far? The way it was originally structured could differ radically from its current status, depending on how prepayments have actually emerged.

In addition to differences in the collateral that underlie CMOs, there are tremendous differences in the structure of CMOs: in the types of tranches, in the number and mix of those tranches, and in certain special features. Early CMOs often featured a simple A-B-C sequential structure, or possibly an A-B-C-Z structure. (A "Z tranche," or accrual tranche, is a tranche whose coupon interest is not paid out currently in cash, but instead is added or compounded to the tranche's outstanding principal balance, which will be paid down later). Now it is not uncommon to see CMOs with 20 or more tranches in them, convoluted and interrelated in any number of ways.

Table 1 contains a list of around 30 or so primary types of CMO tranches, each unique, each with its own characteristics. Dave will spend some time addressing these, just to show you the breadth of originality and creativity in the market. When Wall Street puts its best and brightest into the area of MBSs to create securities, they can come up with a tremendous variety of design.

TABLE 1

Basic CMO Design
Types of Tranches

Sequential	Floater	HIP
Simultaneous Seq	Inverse Floater	TAP
PAC	Super Floater	Stated Final
TAC	PO	VADM
PAC II	PAC PO	Bull
Companion Payer	TAC PO	Bear
Companion Z	Super PO	BRAIN
Clean Z	IO	TRAIN
PAC Z	PAC IO	TTIB
Jump Z	IAC	Residual
Sticky Jump Z	LOV	

There are a number of terms that are associated with certain types of tranches that you may come across, such as "window," "collar," "lockout," and so forth. These are spelled out for you in Table 2 simply for reference. Dave will make reference to some of those in his subsequent comments.

There are also some things that are unique to a CMO as a whole, instead of to a tranche. One example is the "clean-up call." It is not widely understood, but a lot of CMOs have a call provision that enables the trust to fully pay off that CMO once the outstanding balance reaches a certain amount, let's say 10% of the original outstanding. We are only now entering the phase where some of these CMOs are old enough and seasoned enough that this call provision could come into play. I think it will become a far bigger consideration in the future.

TABLE 2

Basic CMO Design
Special Features

Tranche Specific	Deal Specific
Window	Payment Frequency
Collar	Payment Delay
TAC speed	Clean-up Call
Lockout	
Stated Final	
Trigger	
Index	
Leverage Factor	
Accrual Factor	
Cap	

MR. HALL: As a way of thinking about the CMO market, I would like to classify CMOs by the way they allocate principal and they way they allocate interest. Allocation of principal can be classified into three different types:

1. There are sequential allocations, and the simplest form of those are the type that Randy was showing you, the "A-B-C," where one follows after the other, and tranche number two doesn't begin to pay down until tranche number one is fully exhausted.

2. There's a second category that is now prevalent in most deals (CMOs), which I've called a prioritized structure. In this, there are rules for allocating principal that give priority to certain tranches over other tranches within the deal.
3. The third category are those where the allocation of principal is indexed to something. Although that's not a very widespread type of security, they can be very interesting; and in some contexts they can be very useful for solving unique asset/liability problems.

For allocation of interest, there are two types:

1. Some tranches have a fixed rate of interest, or a fixed coupon.
2. Some tranches have a coupon that is indexed to something which varies.

I'll assume that you have at least some basic understanding of the "plain vanilla," or sequential, type of deal (CMO), which is structured only by the chronology of cash flows. So let's jump right to prioritized amortization structures. There are two principal classes: targeted amortization class (TAC) and planned (or prioritized) amortization class (PAC).

Let's start with the TAC, because it's really the simplest case of a prioritized amortization structure. The top curve on Chart 14 represents the total return of principal on a pool of mortgage securities projected at 165% PSA. It's a pattern that you saw in a number of Randy's graphs as well, such as the left-hand graph on Chart 9. The curve climbs steeply at first, as the PSA model seasons through 30 months; then it slopes down as prepayments overwhelm the increasing regular amortization on the mortgages that don't prepay, and the balance of the pool declines.

A TAC is a class under which a targeted sinking fund schedule is based upon a single multiple of the PSA model.

For illustration, I'll create a simple two-tranche CMO in which one tranche is a TAC. To define my TAC schedule, I'll start by assuming that the mortgage pool will prepay at a uniform

CHART 14

Simple TAC Structure
165% PSA TAC Projected at 165% PSA

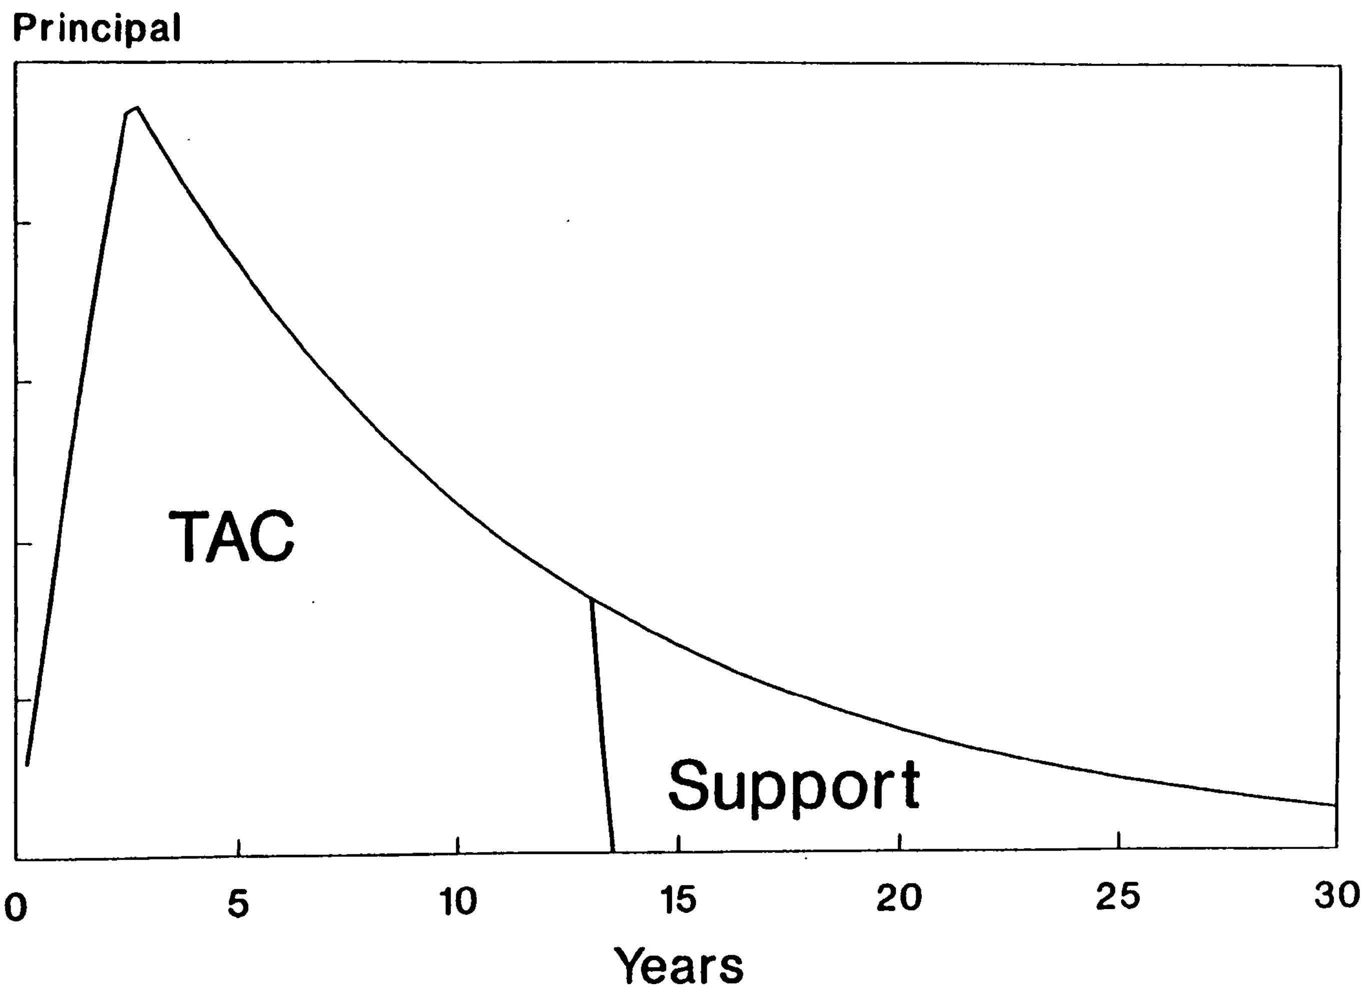

165% PSA throughout the life of the pool. The dollar amounts of the total annual principal repayments will look like the top curve in Chart 14.

Then I'll carve out a portion of this projected stream of principal repayments (the left-hand part of the area under the curve), write down the schedule of projected payments, and designate this schedule as belonging to the TAC tranche. (In this case the TAC piece comes at the start of the CMO, but it could have been placed somewhere in the middle.)

This TAC is going to have that particular schedule of cash flows, predefined at the issue of the deal, to define its priority of cash-flow allocation. The right-hand area, or the "support tranche," is whatever is left over beyond the scheduled terminal date of the TAC.

In this example, I have the TAC scheduled to extend from the first month to about 14 years. And if these mortgages actually do prepay at precisely this modeled speed, then the TAC tranche will receive exactly its scheduled amounts through the first 14 years, after which it will be gone and the support class will then get whatever is left over, and the world will all seem very orderly to all involved.

Of course, it is highly unlikely that prepayments will behave quite this stably. So let's look at how this type of structure will change as prepayments change.

Chart 15 shows the same CMO with the same TAC and support tranche; but in this case, I'm projecting them as though the mortgages underlying this deal are prepaying at 100% of the PSA model. Here they're paying much more slowly than under the original schedule. The solid curve is the total amount of principal payment coming into the trust. Notice how it has flattened out to the left, and thickened up to the right. The light dotted line can be thought of as a shadow of the prior graph. It helps you understand the degree of change that's occurring.

CHART 15

Simple TAC Structure
165% PSA TAC Projected at 100% PSA

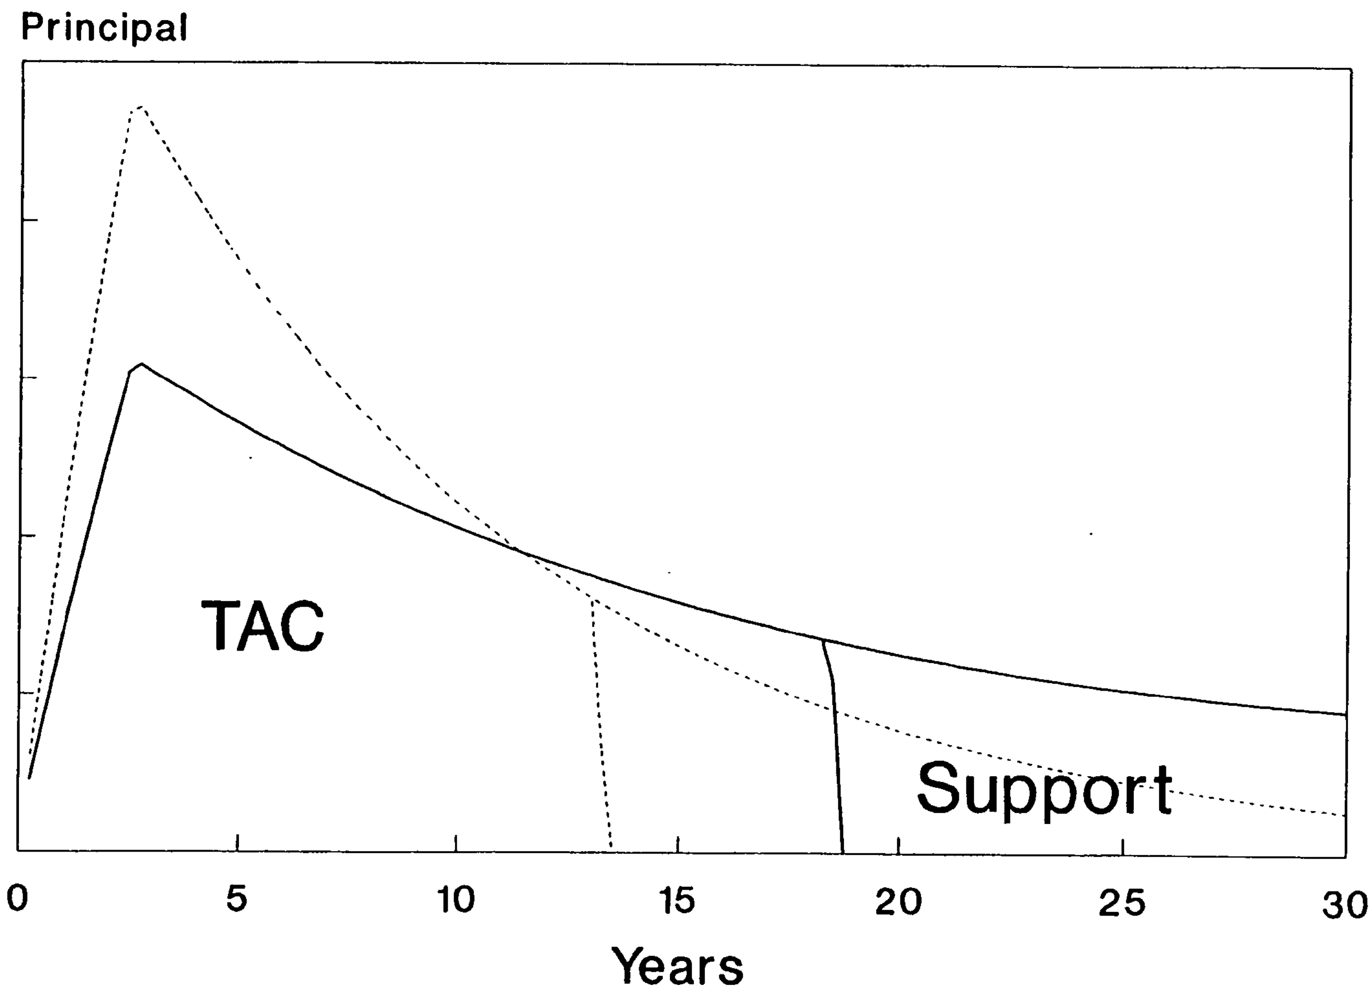

The TAC has priority to get paid up to its regular scheduled amount, which was defined under an assumption that the principal payments would be coming in faster than they actually are. Since there isn't enough principal coming in to pay those scheduled amounts, the TAC gets whatever principal payment is received. But it's obviously coming up a little short. And in fact, well beyond the 14th year, the TAC still remains to be paid down. So there's a period after the 14th year where it continues to be paid while it's getting caught up, until it's finally extinguished. The support tranche continues to wait for its paydowns.

So in this case, if prepayments are slower than the schedule at which the TAC has been constructed, the TAC will not pay off as quickly as it would have under schedule. It will fall behind and then have to catch back up. In this example, the support tranche doesn't even begin to get paid down until after the TAC tranche has finally gotten caught up.

Chart 16 shows the behavior if prepayments are much faster than expected. In this case, I'm using 325% of the PSA model. The top curve is the total amount of principal payment coming into the trust. Notice how it has heaped up to the left, and thinned out to the right. The light dotted line is again the shadow of the original schedule, and the lower solid line is the TAC.

In the early years, because there is more principal payment available than the schedule calls, the TAC is able to receive exactly those amounts which it is scheduled to receive. The support tranche picks up the excess of the actual principal payments over what the TAC tranche was scheduled to receive.

The most notable thing you'll see here is that the support tranche has shifted from the back end to the front end. It is supporting the TAC by picking up the excess cash flows in the early years. But at some point, there are no longer enough cash flows to continue to pay the TAC on schedule. So you can see, where the solid and dotted lines intersect, the amount of cash flow available is no longer sufficient to meet that month's scheduled principal payment for the TAC. The TAC tranche still gets as much principal payment as it can, but it stays behind

CHART 16

Simple TAC Structure
165% PSA TAC Projected at 325% PSA

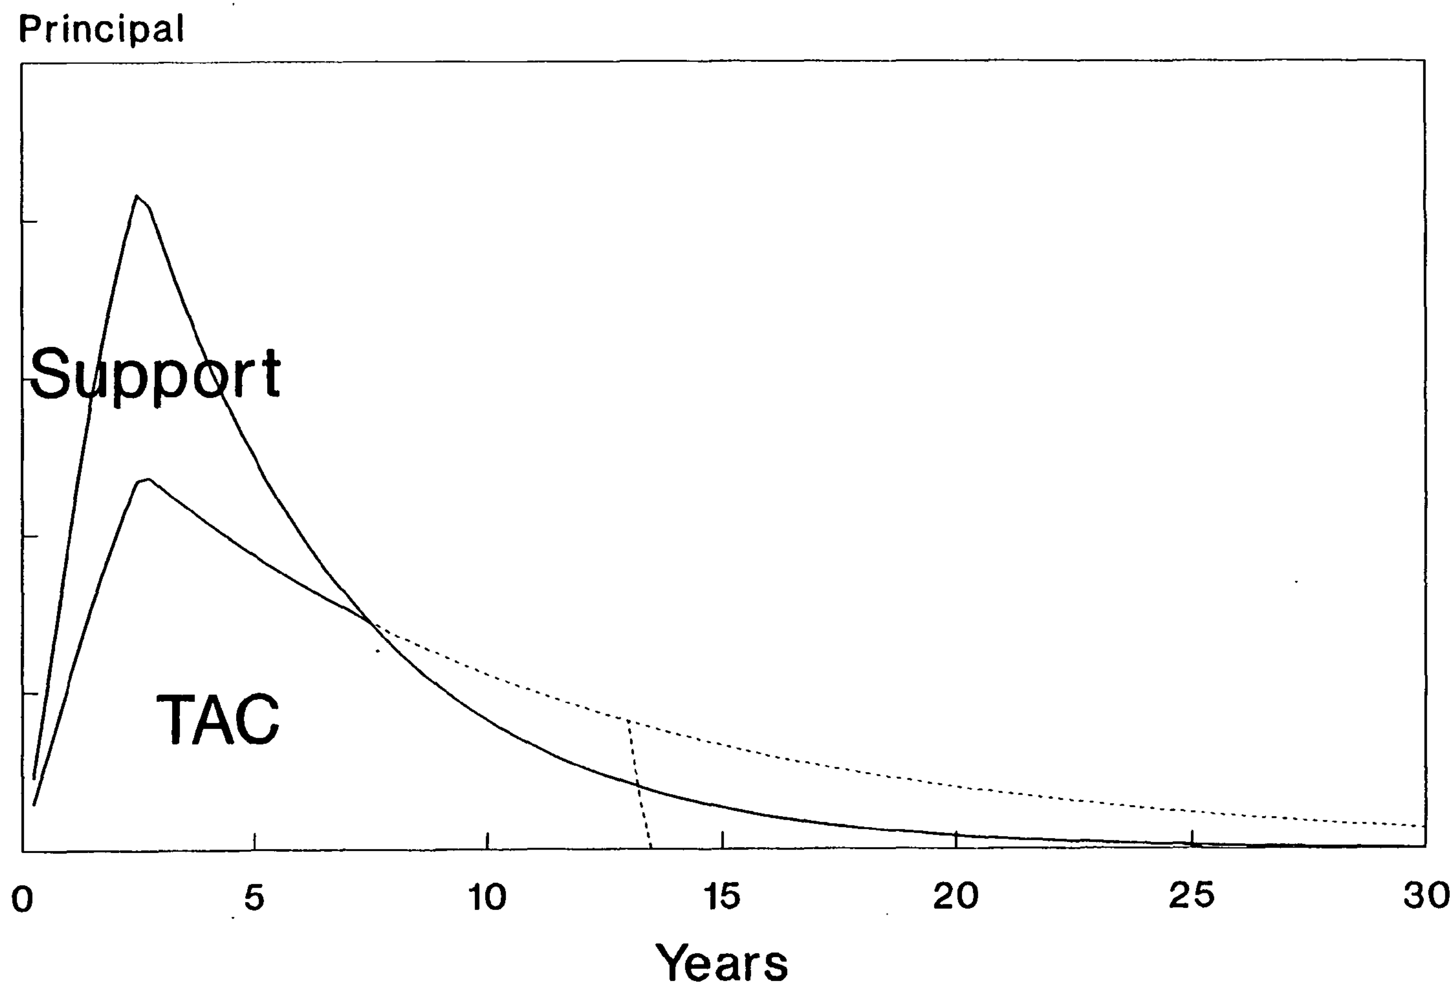

schedule for the balance of the 14 years, and then as you go all the way out to about year 25, it continues to catch up. It's hard to see it on the graph, but at about year 26 the TAC finally gets exhausted. Then there is still a little bit of support tranche left over, and that finally gets paid off in the last four years. Most of the paydown on the support tranche moved from the back end to the front end. But then there was still a tiny tail that had to wait until the end to get paid off. So under these different scenarios of prepayment, the TAC tranche has been much more stable than the support tranche in the timing of its principal return.

As I said, a TAC is a very simple example, because its schedule is defined by a single projected prepayment speed. TACs, however, are not included in many deals. A similar type of tranche, called a PAC, is more prevalent. The only difference between a PAC and a TAC is in the way their schedules are constructed. A PAC schedule is built from the intersection of the curves from a fast and a slow prepayment speed.

In the next example (Chart 17), you can see what the overlay of a "90-300% PAC" construction would look like. This is again an example of a two-tranche CMO, with the PAC being scheduled to start paying down immediately and the support tranche getting whatever principal repayment cash flow the PAC doesn't require.

The dashed line represents the principal return for the pool at 300% of the PSA model. The lower line represents 90% of the PSA model. And the solid line, which is the predefined sinking fund schedule of the PAC, is constructed as the lesser of those two amounts.

So a PAC also has a predefined sinking fund schedule, but its construction is developed by looking at the intersection of the curves for two different speeds.

Now we'll look at the cash flows from principal repayments in this CMO under four different prepayment speeds: 165% PSA (in the middle of the 90-300% collar), 90% PSA (at the low end), 300% PSA (at the high end), and 400% PSA (outside the protected collar).

CHART 17

Simple PAC Structure
Building a 90-300% PSA PAC Schedule

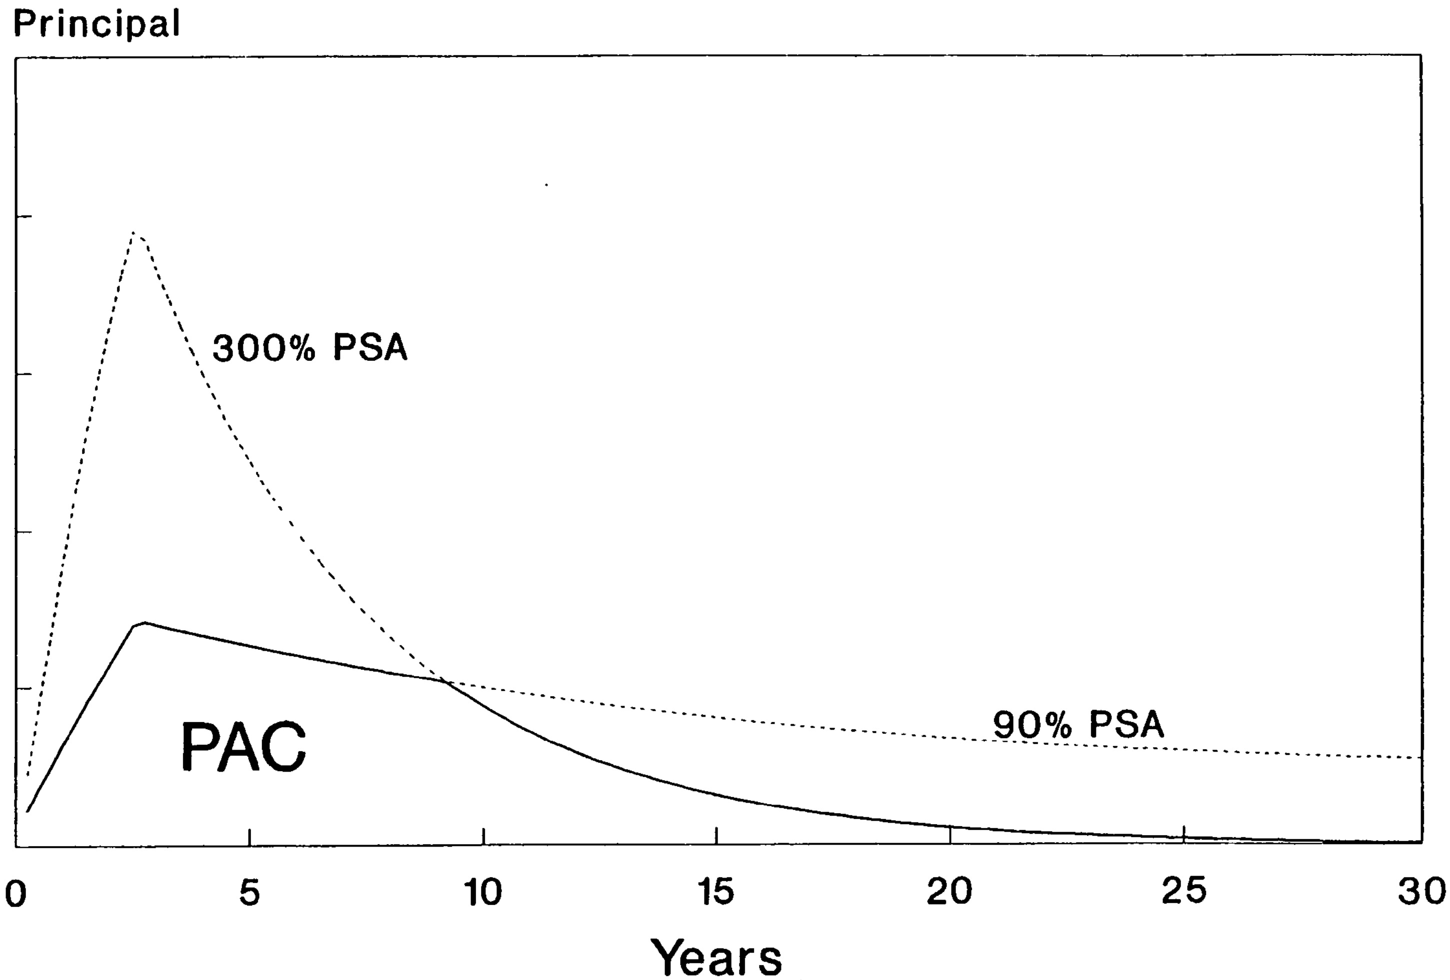

The principal repayments from the pool at 165% PSA and their split between the two tranches are shown in Chart 18. It's clear that the PAC is easily able to have its scheduled payments met throughout. Everything else, between the upper and lower solid curves, is going to the support tranche, and there's some excess being paid out to the support class throughout the term of the deal.

At the low (slow) end of the protected collar, at 90% PSA, the principal repayments from the pool and their split between the two tranches are shown in Chart 19. The light dotted line is again the shadow of the 165% PSA case, which was our base case. The lower solid line is the PAC, and that is also all of the cash flows that are available in the early years. What I'm showing here is that the PAC, even at a 90% PSA, continues to be paid out on schedule.

However, the support has now shifted: Instead of wrapping around the entire period, it's now moved to the back end. And in fact, it's not until that point of inflection (where the 90% and the 300% PSA intersect) that there is any cash flow available to begin to retire the support class. So the support tranche has shifted to the back, but the PAC has remained unchanged.

At the high (fast) end of the protected collar, at 300% PSA, the principal repayments from the pool and their split between the two tranches are shown in Chart 20. The light dotted line is again the shadow of the 165% PSA case, which was our base case.

Again, the PAC has not changed at all; but the support has moved completely to the front. At the point at which the PAC schedule changes from being based upon the 90% PSA level to the 300% PSA level, the support disappears (is completely paid down). Assuming that we continue to receive prepayments at 300% of the PSA, the PAC will continue to maintain its schedule.

CHART 18

90-300% PSA PAC Structure
Projected at 165% PSA

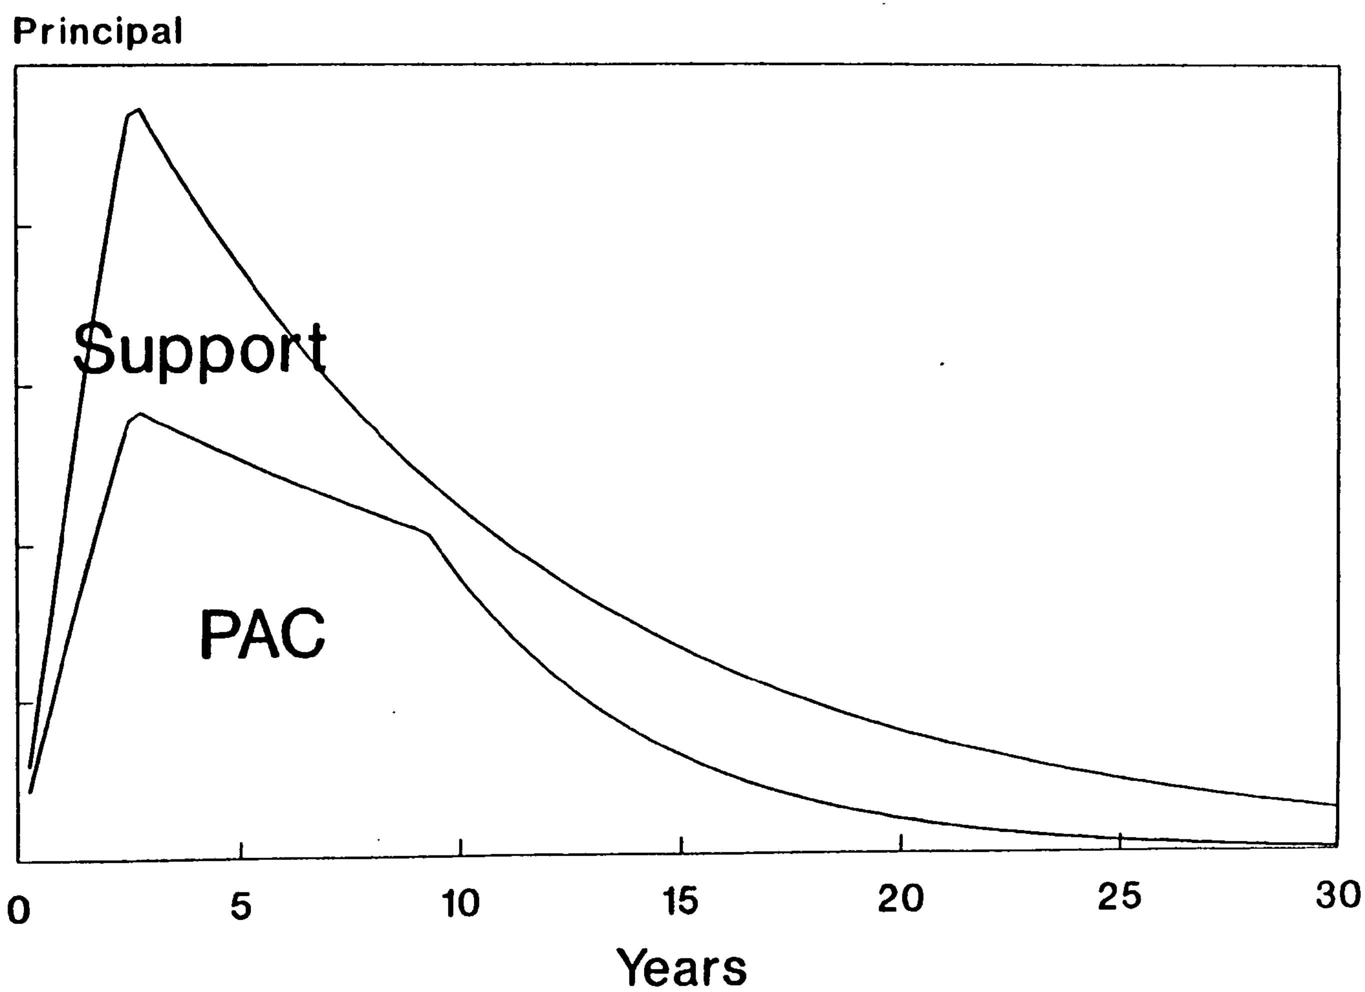

CHART 19

90-300% PSA PAC Structure
Projected at 90% PSA

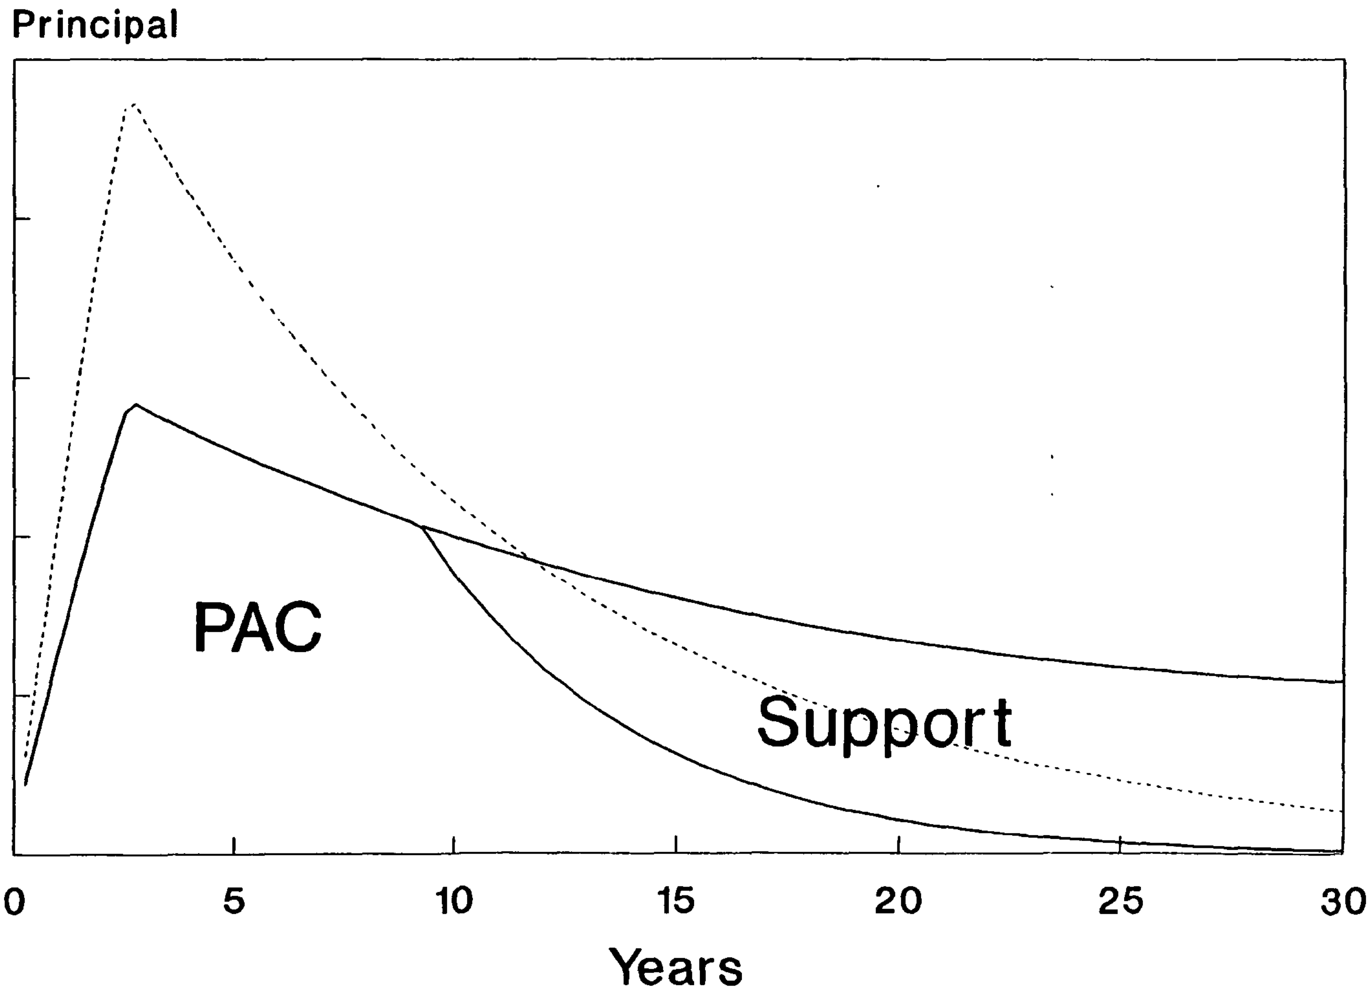

CHART 20

90-300% PSA PAC Structure
Projected at 300% PSA

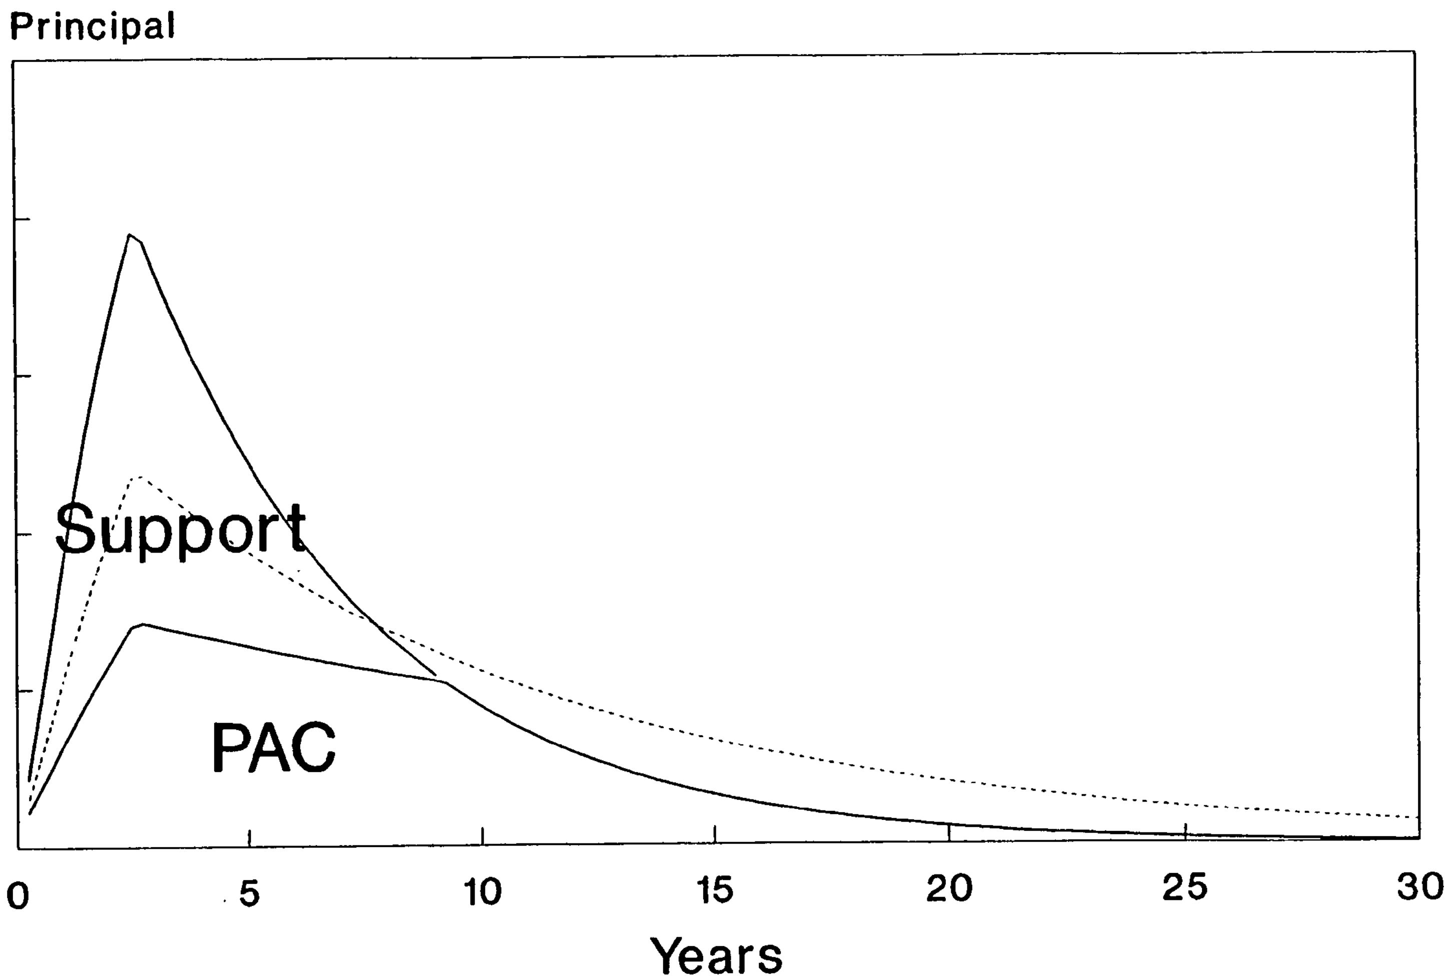

However, it's clear to see that from that point (about the 10th year) forward, there is no more support. So from then on, any variability in cash flows away from that 300% level (up or down) will have to be absorbed by the PAC tranche.

Now suppose the prepayment rate is even faster than the high end of the protected collar, such as 400% PSA in all years. This is shown in Chart 21. The PAC tranche is the lower solid line with a shark's fin sticking up, and the support tranche is the tall narrow piece sitting on top of the PAC tranche to the left of the shark's fin.

Again, with so much early money coming in, the support moves to the front of the line. But this time we run out of support even before we reach that point of inflection. At some point around year five we have exhausted our support, and again, all that's left is the PAC tranche. Then the PAC begins to accelerate, and it pays down faster than its schedule called for, since it is the only tranche left in this deal from this point forward. So although PACs retain a very significant amount of stability over a wide range of scenarios, there is still the possibility that a PAC may break its schedule under extreme conditions, and either come forward or move backwards.

This has been a very simple visual depiction of how scheduled sinking fund tranches work, or rather one of the most common forms of scheduled sinking fund tranches.

In my examples, I was showing you one large PAC or TAC tranche that started at the beginning of the period and ran, in some instances, through the duration of the security. The most common variation of this is a simple sequencing of PAC tranches. Most deals now will have subdivisions, sequentially aligned, within that PAC tranche. At one point, it was common to talk about "six PAC" tranches, where the PAC (as defined in a macro sense) would be carved into six different sequential tranches. You now find deals that have as many as 20 or 30 PAC tranches.

CHART 21

90-300% PSA PAC Structure
Projected at 400% PSA

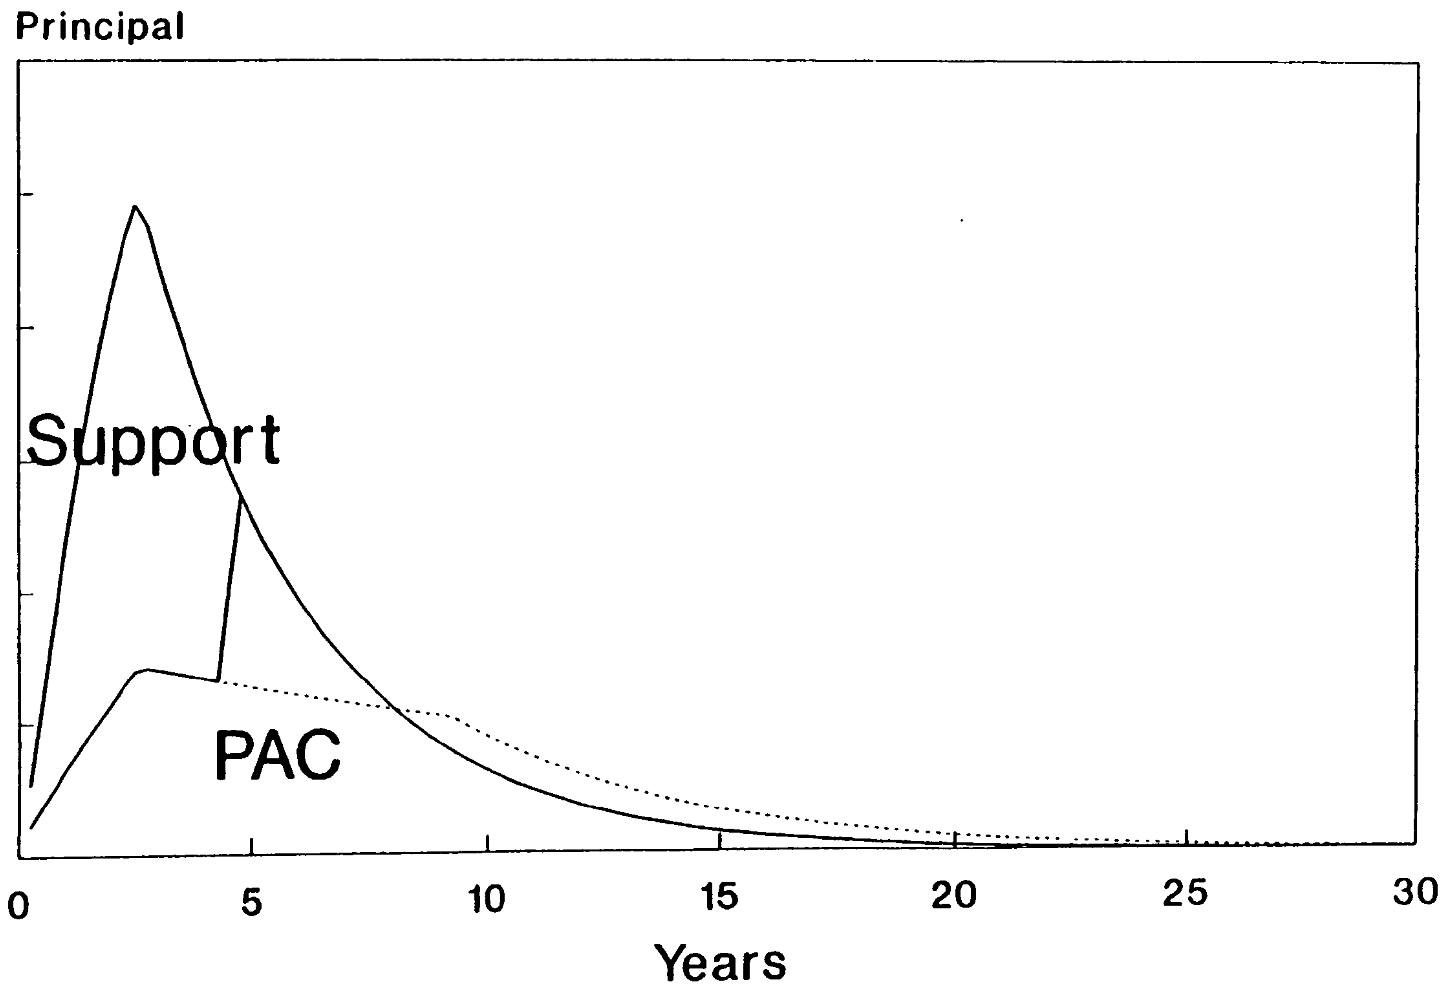

An accrual PAC is a PAC that is also an accrual tranche. An accrual tranche, or Z tranche, is one of the later tranches, and its interest, instead of being paid out as a coupon payment, is used to pay down the principal of an earlier tranche. This causes the accrual tranche to behave more like a zero-coupon bond, in that it's receiving no coupon payments for some period of time. But unlike a zero-coupon bond, its outstanding principal amount grows exponentially at the coupon-payment rate, and the amount of the outstanding principal, when it starts paying down, depends on how long it has been before it starts paying down. Accrual PACs have been very common over the last several years; as alternatives to zero-coupon bonds, I think they'll show up in many GIC portfolios.

PACs and TACs can also be split into IOs and POs. Again, the PAC or TAC defines a schedule of outstanding principal on a deal. You can carve PO tranches out of that; or you can even have IO tranches based upon that notional PO schedule.

There's something called a Type Two PAC (PAC-II). This is a scheduled tranche that, in terms of its priority for payment, falls in between a regular PAC and a support tranche. Although the PAC-II's cash flows have a predefined schedule associated with them, they are subordinated to a more senior PAC, or a Type One PAC (PAC-I), in case there is insufficient cash to meet the scheduled payments of all the PAC tranches in the deal. On the other hand, the PAC-IIs are still senior to the support tranches of a deal. So they're a tranche with an intermediate type of call (prepayment) protection: more stable than a regular support tranche, but less stable than a PAC, particularly as you get into extreme situations.

There are even Type Three PACs (PAC-IIIs), which just add one more layer of priority within a structure.

"Liquidity tranches" are special PACs for which the stated final payment date, even under a 0% PSA assumption, does not extend beyond some date, typically five years. They're called liquidity tranches because banks can use them in certain situations (because of some regulatory efficiencies caused by having that final date of maturity no longer than five years). In many cases, their schedules are supported only by the interest accruals of some later tranche.

These liquidity tranches are sometimes called VADMs. I believe that VADM stands for Very Accurately Defined Maturity. (I understand that these were also the brainchild of a Wall Street rocket scientist whose first name was Vadim. So it had a nice double meaning).

Floaters and inverse floaters can be carved out of PAC tranches, and we'll talk about that combination later on when we talk about coupon. I mention it here just to indicate that PACs and TACs can be floating rate or inverse floating rate.

And finally, there are what I call Ex-support priority differences. This is not a very user-friendly term, but here's what I'm referring to. In a deal with sequenced PACs, once all the support tranches have been paid down, the normal rule for allocation of principal is to go from front to back. That is, the shortest remaining tranche picks up the excess principal repayment first, and then the next shortest, working its way to the later tranches. But there are deals where that allocation process, once support is gone, is not sequential. In these deals, some PAC tranches that would normally occur at the back of a deal become like support tranches if prepayments are fast enough. So behavior outside of the protected PSA range sometimes can be different from deal to deal with all other things being equal.

Support tranches also have common variations. Again, the most obvious variation is sequencing. You can also have accrual support tranches. And you can have POs. A support PO is actually a special case, which is called a "super PO." A PAC-II, although a PAC, is also a support tranche. Support tranches can also be split into floaters and inverse floaters.
A "jump Z," which was a variant of a support tranche, was common for about a six to nine month period, ending about the time we developed the agenda for this session. It's an accrual tranche; but if prepayments exceeded some certain threshold, it would jump to the front of the

queue instead of following the normal sequencing of payment priority, and get paid down very fast.

There were "sticky jump Zs" and "toggle jump Zs." A "sticky jump Z" is one which remained in the front of the line if conditions were ever such that it was caused to jump. A "toggle jump Z" could jump back and forth, depending on whatever the relevant index was doing in that month.

Virtually all jump Zs are gone now, because prepayments have accelerated over the course of 1991. Just about every jump Z that existed has been triggered, and has jumped, and is gone, even if it was a toggle jump. I haven't seen any created lately, so it's probably a class of security that is of historical interest only. Since it's not particularly relevant, I won't spend any more time on it right now.

The third type of Principal Allocation Tranche are those that are indexed. The general rule here is that you take a normal type of tranche, such as a sequential, PAC, TAC or support class; then you split it into two companion classes. Then you define an allocation rule which operates strictly between those two companions, which is a function of some formula linked to a specified index or variable. These classes are usually built to provide for a very customized need, and therefore they're not in widespread use. But they can be very effective ways to deal with unusual problems.

This is probably most easily understood by way of an example. Let's consider a PAC PO, which has a PAC structure for principal payments but is a PO tranche with no coupon. Suppose this PAC PO is itself split into two classes that have equal principal at issue. Thereafter, the principal payments that would normally have been paid into the PAC PO as a whole are allocated between the two classes. The allocation rule for this combined PAC tranche will be a function of the prevailing 10-year constant maturity Treasury rate, at each prepayment date throughout the duration of the PAC tranche. This allocation of principal will occur until one of the two companion tranches is extinguished. At that time, whichever one is left continues to pay down as though it were a regular PAC.

An example of such an allocation rule is shown in Table 3. Let's first look across the middle line of the table. If the constant maturity 10-year Treasury yield is at 8% on a determination date, then Class A will get 50% of the principal, and Class B will get 50% of the principal. In fact, if it continues at that rate throughout the life of the classes, you will have a very uninteresting subdivision, because each half will pay exactly alike.

However, if the Treasury yield on any reference date is higher, Class A begins to get a higher proportion of the principal which would be allocated to the Combined Class A and B combination, offset by a lower percentage on Class B and vice versa. So the Class A security pays off more rapidly if interest rates rise (which is the opposite of the behavior of mortgage pools and callable bonds).

TABLE 3

Indexed Principal Example
Cash Flow Allocation Table

10-year CMT	Class A	Class B
10% (yield)	70%	30%
9%	60	40
8%	50	50
7%	40	60
6%	30	70

(For example, if the 10-year CMT is 10% on an allocation date, then Class A receives 70% of the principal which would normally be paid to the Class A and B combo.)

If the Class A security was purchased at a discount, then it behaves very well and yields more in a rising interest-rate environment, since it gives you a return of principal faster and so accelerates your discount booking. Therefore it offers some very interesting duration-shortening properties. This is often thought of as a "bear-ish" security (meaning it does well when "regular" bonds are losing value). Behaviorally, Class A is even somewhat like an IO, although in this case you're buying it at a discount, so you know you can't lose principal with this security. You know you're always going to get back the full amount of principal. Similarly, if rates fall, Class A will become a longer-duration security, and if bought at a discount will yield less.

Class B will be the complement. It will behave more like a regular PO, accelerating when interest rates fall, and slowing down when interest rates rise.

But a very nice feature of this Class B is that the behavior is directly linked to changes in Treasury yields. It does not depend upon prepayment behaving in a corresponding manner. so you don't need to get prepayments modeled exactly right. The PAC schedule is going to control principal amounts allocated to this combination (within reasonable bounds), and only changes in 10-year Treasury rates will affect the principal amounts lengthening or shortening.

Randy mentioned clean-up call provisions. Most deals now have 1% clean-up calls, which can be pretty much ignored. But some older deals have call provisions that could be triggered if the collateral reached 10% or even 20% of its original balance.

In some instances, there's a fixed call date. For example, at any time after December 1, 1999, whatever is outstanding is callable. There are even PAC tranches that have a call date that precedes the first scheduled principal date, which presents an interesting conflict.

In rare cases, securities are putable, as well. In most instances, these are direct issues of one of the government agencies, where the agency agrees to buy it back, typically at the end of five years, so that it can qualify for liquidity purposes.

Indexed coupon structures are similar to indexed principal structures. These come from pairs of companion classes, which could be combined again to form a fixed-rate class. Principal typically pays down identically on the two, but the coupon variability of one class offsets the variability of the other.

Common indexes for coupon indexing are the London InterBank Offered Rate (LIBOR); the 10-year constant maturity Treasury rate; or the "Eleventh District Cost-of-Funds Index," which is an average portfolio rate of banks in the 11th district (principally California and Nevada).

Let me show a typical coupon formula for an indexed tranche. (The general formula is shown in Formula 1, and a specific example is in Formula 2).

FORMULA 1

Floater/Inverse Floater
Typical Coupon Formula

$$C = (L * I) + K, \quad 0 <= \min <= C <= \max$$

where C = coupon rate
L = leverage factor
I = index
K = constant
min = minimum coupon rate
max = maximum coupon rate

FORMULA 2

Floater/Inverse Floater
Sample Coupon Formula

Given: Fixed rate 9.5% pass-thru collateral

Floater: C = LIBOR + .75%, .75% <= C <= 11%

Inverse: C = -6.33*LIBOR + 64.8825%, 0% <=C<=64.8825%

Par value of floater = 6.33 * par value of inverse

The coupon is equal to a "leverage factor" times the index, plus a constant. This coupon is bounded above by a specified maximum coupon rate, and below by zero or, in some cases, a minimum stipulated rate.

As an example, here's a combination of a floater and an inverse floater. The par value of the floater is equal to 6.33 times the par value of the inverse. The floating-rate class pays LIBOR plus 75 basis points. This coupon is subject to a minimum of 75 basis points (which would occur if LIBOR was zero), and a maximum rate of 11% (which would occur if LIBOR was 10.25% or higher). The companion class is the inverse floater, which has a leverage factor of 6.33 to one. Its coupon is (-6.33) times LIBOR plus 64.8825%. It is subject to a minimum coupon of zero and a maximum of 64.8825%.

Algebra will demonstrate that the average coupon of the two will never exceed 9.5%. This is a condition that is absolutely necessary in order for this floater and inverse floater to be fully supportable, if they are part of a 9.5% pass-through collateralized deal.

Residuals are a special class, and I can't begin to describe all the variations. These are definitely not for beginners. In the older deals, they were whatever was left over after everything else was paid down. They consisted primarily of excess principal and/or excess interest. In more recent deals, they tend to be structured to behave like regular tranches. However, they have some tax anomalies.

REMIC legislation provides that every deal has to have a residual that serves as a tax balancing item to prevent a REMIC issuer from allowing a tax arbitrage condition to create an incentive for doing a deal. The residual holder is generally the entity that controls the call option in the deal. All REMICs have to have one class designated as a residual.

I have some editorial comments to wrap up this fast exposition of some simple types of CMO structures:

1. You need to avoid the urge to view these structures generically. From deal to deal, there are quirks that affect the behavior of tranches, and you have to understand the entire structure of the deal to fully appreciate the behavior of each tranche.
2. You need to recognize that inefficiencies do exist in this market. One might suspect that if the market is efficient, that the relative value of different tranches within a class must all be the same. But I think it's been aptly demonstrated, over the last decade, that this is clearly not the case. The mortgage securities market, particularly through the development of CMOs, remains one of the most inefficient markets around. And there are opportunities to add value in this market.
3. You need to consider the impact of interest-rate variability on the behavior of your tranche. You cannot classify a tranche as a three-year tranche. Sometimes you can't even classify it as a three-year PAC or a three-year support. You need to think of how variability in cash flows affects the behavior of the tranche. This means that there's really no shortcut to evaluating CMOs.
4. You're really forced to think.

MR. BOUSHEK: I want to make one point about our discussion thus far. When we talk about agency CMOs (or agency-backed CMOs), we're talking about CMOs created out of GNMA or FNMA or FHLMC collateral. Those CMOs do not necessarily bear the name

"FNMA REMIC" or "FHLMC REMIC." Many are what we call "private label" CMOs. They are supported by agency pools, but they bear names like American Southwest Financial, or Collateralized Mortgage Obligation Trust, or Drexel Burnham Trust, or Morgan Stanley Trust, or any number of others.

When we talk about "whole-loan CMOs," we're talking about a completely different animal. These are non-agency-backed CMOs. If I can refer back to our opening discussion, there is another creation path a CMO can follow (Chart 22). As before, homeowners take out loans from originators; however, rather than selling those loans as packages to an agency, they (the originators) bypass the agency entirely and create CMOs directly out of the home loans.

The single largest reason for the creation of whole-loan CMOs, or whole-loan direct pass-throughs, has to do with size. The GNMA, FHLMC and FNMA programs all have maximum loan sizes for their pools somewhere below $200,000. For loans over $200,000, there is no agency to go to. The market for whole-loan CMOs has not grown as rapidly as it has for agency CMOs, until the last two years. Growth has come more recently as more players have become generally familiar with CMOs, and more willing to extend their horizons to credit instruments as well.

The differences between agency-backed CMOs and whole-loan CMOs can be classified into differences in (1) credit risk, (2) characteristics of the underlying loans, and (3) the impact of various forms of credit enhancement on cash flows:

1. Credit risk -- The unique thing about this type of instrument is that there is now also default (C-1) risk as well as prepayment risk. This is because there is no federal government agency standing in the middle, guaranteeing the cash flows.
   a. Variety of ratings (Unrated through Aaa) -- As to credit risk, there are a variety of ratings. You can buy whole-loan CMO tranches that are everything from unrated or below investment grade to Triple A.

CHART 22

Evolution of a Nonagency CMO

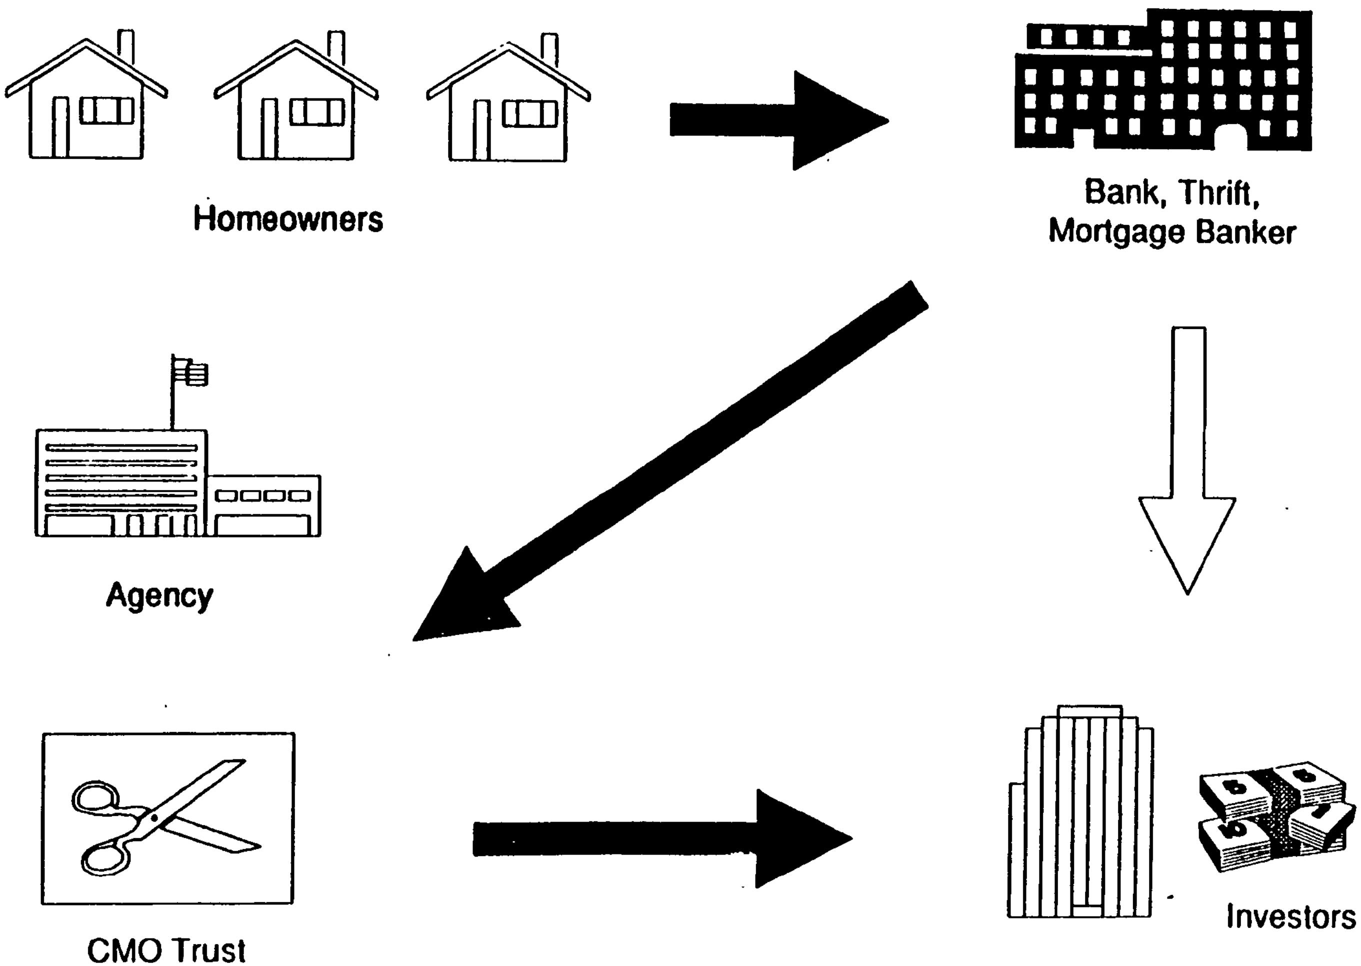

b. Variety of issuers (Citicorp, Pru-Home, RFC, Ryland, Chase, GE, RTC) -- These are some of the bigger names that may be in the market, but this list is certainly by no means complete.

c. Variety of credit enhancements (Pool insurance, letter of credit (LOC)/guarantee, subordination, reserve/accumulation account, tiered) -- Finally, there are a number of different ways that these CMOs or pass-throughs are enhanced to make them attractive to investors. Generally, some type of credit enhancement is added to somewhat mitigate the credit risk. This enhancement may take various forms, such as pool insurance or a guarantee.

What investors have come to prefer are structures that include subordinated tranches that absorb credit losses, or structures that include some kind of a cash account initially set up, or accumulated through time, via excess servicing. And there are actually now a number of structures that have combinations of these. It is not unusual to find a CMO that has what is called a super senior AAA tranche, a tranche subordinated to that tranche that is rated AA or even AAA, and then another form of credit enhancement, either pool insurance or a reserve account, underneath that.

2. Characteristics of the underlying loans.

a. Loan size, as discussed above.

b. Documentation/underwriting standards -- A number of the loans that do not qualify for FHLMC or FNMA fail because of their limited documentation. It is often high-wealth individuals who can go through limited, short-form approvals. This lack of documentation often disqualifies them from the agency programs.

c. Loan-to-value (LTV) ratios -- For agency pools, you could have LTV ratios as high as 95% for FHLMC and FNMA and 100% for GNMA. These loans would be for nearly the full value of the property, with little or no equity down.

In whole-loan CMOs, it is generally standard to look for an average LTV ratio under 80%, or in other words, an average of at least 20% equity in the deals. This is a form of credit protection to us as investors, because the homeowners have a significantly higher commitment of personal equity in the deal. That is one of the underwriting criteria we use in evaluating whole-loan packages.

d. Geographic concentration -- The whole-loan market is very heavily weighted towards California, in large part owing to the number and average size of mortgages there. That has a number of implications, not the least of which is diversification of risk, but also prepayment speed. This is because, historically, California has tended to be a rapid prepayment state. So there are implications of that geographic concentration. One of the underwriting criteria we use in evaluating these deals is to ask if there is enough geographic diversity in the underlying loans.

e. Gross WAC -- These are often higher-coupon loans. Because they're jumbo loans, they often carry a higher interest rate. So, you have a more interest-sensitive borrower, but you also have a higher threshold to refinancing. And those two work a little bit in offsetting one another.

3. Impact of various forms of credit enhancement on cash flows -- There are a few things that lead to cash-flow differences in the way you would model these types of CMOs.

   a. Slightly faster prepays -- Whole-loan CMOs generally display somewhat faster prepayments, due to concentration in the West Coast, larger loan size, and more sophisticated borrowers.

   b. Subordinated structures: default losses and shifting interest -- In the subordinated structures, there is another redirection of cash flows. Any prepayments that arise specifically from defaults may be split off from other principal payments and prepayments, and redirected into specified tranches. In a whole-loan CMO, a subordinated tranche is the first one to receive any prepayments that arise from defaulters. Now you have some prepayments that follow one rule, and some prepayments that follow a different rule. More significant is the shifting-interest structure that's common in whole-loan deals with subordinated tranches. In a shifting-interest structure, you might have a subordinated tranche that is initially 8%

of a CMO. However, the 8% actually grows (in percentage terms) over the first few years, because all prepayments other than defaults are directed away from the subordinated tranche.

So you actually receive a disproportionate share of prepayments if you are not the subordinated tranche. Therefore that accelerates your cash flow and again leads to faster prepays until you reach some point (such as five years) when you start sharing equally with the subordinated tranche.

c. Change in structure upon exhaustion of credit enhancement -- The last point is somewhat more on the edge. But it has the potential to alter the pattern of payments to your tranche, and you need to understand it in modeling cash flows.

A whole-loan CMO may burn through its credit support, whether that credit support is pool insurance, or a parent guarantee, or comes from a subordinated tranche. It may be protected for, say, 10% of losses.

Once it burns through that, three or four different things can happen: (1) all the remaining tranches will then share losses proportionally; (2) all losses will be borne by the first tranche that has to be paid off; (3) all losses will be borne by the last tranche that has to be paid off; or (4) all of the remaining tranches become pass-throughs at the time that the support is burned off. It's all deal-specific, depending on how the prospectus was put together.

It could happen that you own a tranche that was originally six years, and is now down to a two-year tranche, but all the credit support is gone. Then if there are additional losses, your tranche could become a 25- or 24-year pass-through, along with all the rest of the tranches in the deal.

That's somewhat of an extreme scenario, but it is something that you must understand, since it can happen to these types of securities.

The whole-loan market is a market that's growing, one that more and more companies are entering. The general problem is the same (the uncertainty of prepayment), but in this case there are other elements of risk and other factors that impact cash flows.